Atmospheric Chemistry and Physics Discussions

Atmos. Chem. Phys. Discuss., author comment AC3
https://doi.org/10.5194/acp-2021-924-AC3, 2022

Reply on CC2

Stephen E. Schwartz

Author comment on "Observation Based Budget and Lifetime of Excess Atmospheric Carbon Dioxide" by Stephen E. Schwartz, Atmos. Chem. Phys. Discuss., https://doi.org/10.5194/acp-2021-924-AC3, 2022

I thank Dr. MacCracken for his Comment and for his careful read of the manuscript. For convenience I repeat Dr. MacCracken's comments here with my responses interspersed in Bold.

General Comments (drawn out of the specific comments below):

G1. The first thing that would be helpful to have clearly documented is how the baseline redistributions of emitted carbon are split into the various reservoirs over time. So, what is the airborne fraction, what fraction is going into the ocean and what fraction (net of uptake minus deforestation source) is going into the terrestrial troposphere. The Global Carbon Project prepares its analysis each year based on a range of data that they gather. And then, once done for the present, does the model represent the integrals over the historic period, etc. It is nice to base various fluxes on observed information, but the real test is whether, when everything is put together, the fluxes redistributing carbon into the three closely coupled reservoirs (the atmosphere, ocean mixed layer, and living terrestrial biosphere). It would also be helpful to have a comparison with how the MAGICC does this redistribution in that it has been calibrated to existing models—this is not to suggest that the detailed models that attempt to build up the carbon cycle from fundamental processes as much as possible are correct, but to understand what differences do exist and what that might show.

G2. The results of the paper are, as is acknowledged in the paper, dependent on how the transfers of carbon between the mixed layer and deep ocean. In that this relationship is determined by the circulation of the ocean, it would seem essential that the model for the interaction represents the ocean circulation and its effects. It is well-established through climate simulations that the representation used in the Schwartz model simply does not lead to the observed temperature distribution of the ocean, namely a relatively warm upper ocean and a cold deep ocean. The observed distribution is a result of accounting for the downwelling of cold water in polar regions and its gradual upwelling over much of the area of the lower- and mid-latitudes. And it is also the case that simply representing the upper ocean as a wind-mixed layer without accounting for transport along isopychnals from near the surface in mid-latitudes down to several hundred meters in the lower latitudes is not likely to be adequate. These ocean influences can all be represented with a quite simple parameterization without having to include a globally finely resolved ocean representation. The proposed ocean model basically ignores the well-established understanding of the ocean circulation and instead includes only a one box deep ocean.

Were this the approach used in a climate model, the equilibrium result would be a deep ocean and mixed layer having the same temperature—this is not only not the case in the real world in terms of the actual temperature, but even if trying to model only the perturbed temperature, one would not expect the change in the upper and lower temperature to be the same. And since CO_2 is transported in the ocean by the same ocean circulation as the heat, on would not expect that the CO_2 is being correctly being moved around in the ocean.

Despite this situation, the author goes to great ends to defend his one box approximation for the deep ocean, citing papers not really relating to the type of use he is putting the model, basically saying global models are too uncertain. Rather than what has been done in simplifying down climate models, which calibrate themselves to the complex models, this author basically claims to be sufficiently correct. Given all the time that has been put into this paper, it would have been quite easy for the author to actually investigate if his approximation and a slightly more complex model would give the same result. He could readily put in a one-dimensional ocean model with a polar downwelling pump as is done in the MAGICC climate model (and has been done in many other simple climate models). Another useful test that could be done is to test the dependence of the results on the assumed depth of the ocean. Detailed models and observations both suggest that downward mixing of heat in the ocean goes only as deep as near the depth of the thermocline (so maybe down to 750 meters total or so, including the mixed layer) and cannot go deeper because of the upward ocean movement in low latitudes as a result of the polar downwelling.

With radiation transport not a process in the ocean, it is the circulation that carries the heat around—and so could also be used to carry the CO_2 around. If such a big claim is going to be made about the quite short lifetime of a CO_2 perturbation if emissions are suddenly halted, there is just no excuse for not really doing the extra modeling to prove his choice is justified.

I first reiterate that the two-box ocean model used in the present study that Dr. MacCracken disparages has successfully represented the uptake of heat by the deep ocean and more broadly the response of global mean surface temperature to abrupt perturbations in forcing. The reader is referred to the Response provided earlier to the Comment by Dr. Halpern.

Here I briefly recapitulate the process by which the piston velocity representing the rate of tracer transport between the ML ocean and the deep ocean was developed in the manuscript: The AR5 assessment of the rate of heat uptake by the world ocean was separated into two components: uptake by the ML ocean (taken as 100 m) and transfer from the ML ocean to the deep ocean (DO). The rate of the latter was evaluated from the difference between total heat uptake (as measured by soundings, over time periods, 1971-2010; 1993-2010) and the difference in Global Mean Surface Temperature over the same time periods. As described in the manuscript other, independent estimates of the piston velocity are obtained from the difference in $^{14}CO_2$ between the ML and DO using the lifetime of ^{14}C as a clock. The concept of a piston velocity between the two compartments is thus seen to be well established.

I also note that the two compartment ocean model does not attempt to represent the temperature in each compartment but rather the difference in temperature relative to preindustrial. The end state of a pulse input of heat into the upper ocean would not be a uniform ocean temperature (equilibrium) but rather a difference in temperature between the two compartments that is equal to the difference prior to the heat input (steady state). The same would be said for an input of DIC (dissolved inorganic carbon) following emissions of a certain

amount of CO_2 into the atmosphere, the situation being examined here following abrupt cessation of emissions.

As Dr. MacCracken notes, the MAGICC (Model for the Assessment of Greenhouse Gas Induced Climate Change) model that he advocates was developed as an emulator of more sophisticated greenhouse gas and climate models; for a historical overview and description of the current version of the model see Meinshausen et al. (2011, 2020). I would emphasize that application of that model requires specification of some dozens of parameters, an approach that differs greatly from that taken in the present study, which rests on a minimal number of observationally constrained parameters. Following Dr. MacCracken's encouragement to examine the consequences of abrupt cessation in the MAGICC model, I used an "off-the-shelf" version of this model that is available for use with default parameters at https://live.magicc.org/ to compare the results from the model developed in the present study (present model, hereinafter) with those from MAGICC for abrupt cessation of anthropogenic emissions. The resulting time dependence excess of CO_2 following cessation of emissions in 2015, Figure R1, indeed shows much less rapid decrease than that determined with the present model (cessation in 2017) for the three sets of parameters used in Figure 1 of the manuscript. The question is thus raised over the reason or reasons for the difference.

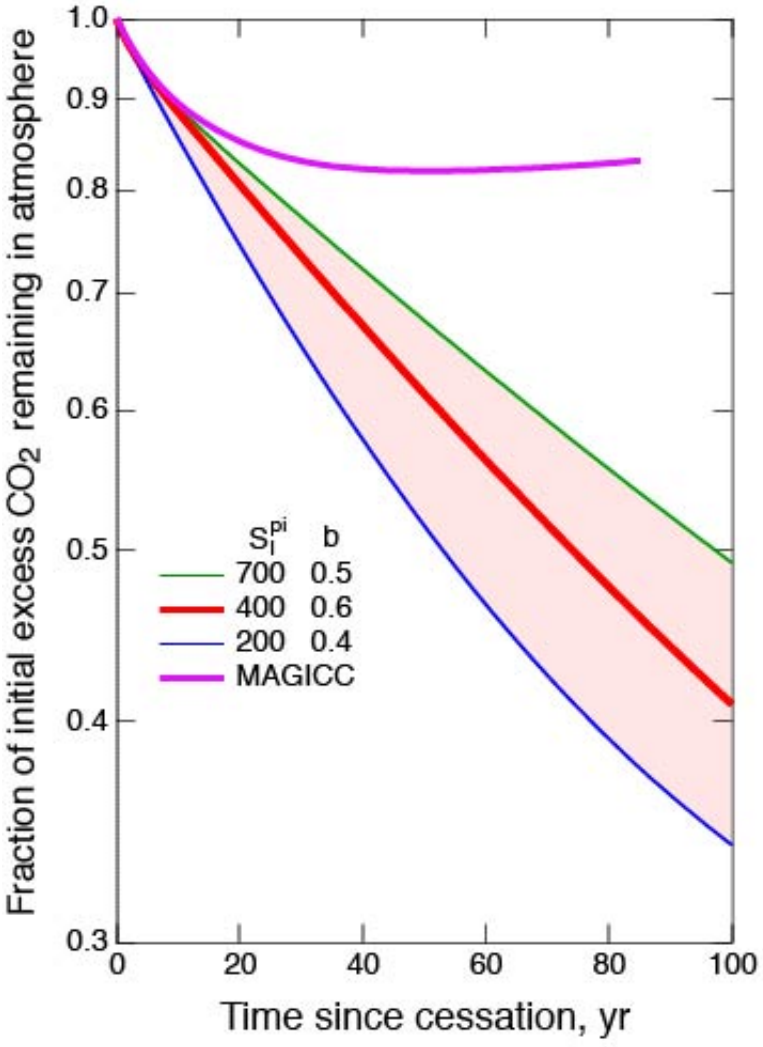

Figure R1. Decay of excess atmospheric CO_2 and corresponding adjustment times over the initial 100 years following abrupt cessation of emissions as determined with the present model and by the MAGICC program (85 years). CO_2 amounts are plotted as fraction of amount in excess of preindustrial amount relative to value at time of cessation of emissions on a logarithmic scale versus time subsequent to cessation as in Figure 1 of the manuscript. Results from the present model are shown for three pairs of parameters (preindustrial stock in labile biosphere, S^l_{pi}, and fertilization exponent b) bounding the range of rate of decay, as described in the manuscript.

The annual change in the stock of CO_2 in the atmospheric compartment (AC) ΔS_a and fluxes from the AC into the ocean compartment (OC) and terrestrial biosphere (TB) are shown in Figure R2 as a function of time subsequent to cessation of anthropogenic emissions of CO_2 as calculated with the present model and with the MAGICC model. Prior to cessation the two models exhibit close agreement in the annual rate of change in the atmospheric stock ΔS_a. However, consistent with Figure R1, after cessation these rates of change differ markedly, with S_a decreasing much more slowly in the MAGICC results, and actually increasing slightly towards the end of the record, in contrast to the relatively large rate of decrease obtained with the present model over the 85-year time period shown in the Figure. Examination of the net fluxes from the AC into the receiving compartments, the OC and the TB show much greater differences between the two models throughout the time record. Throughout the record, the flux from the AC to the OC is much lower in the MAGICC results than in the results from the present model, and correspondingly the flux from the AC to the TB is much lower in the MAGICC results than in the present model results. (The sums of the two fluxes are roughly the same for both the models, as constrained by the difference between emissions and increase in S_a.). It would be valuable to understand the reasons for this difference in apportionment of sinks, even for this period prior to cessation of emissions.

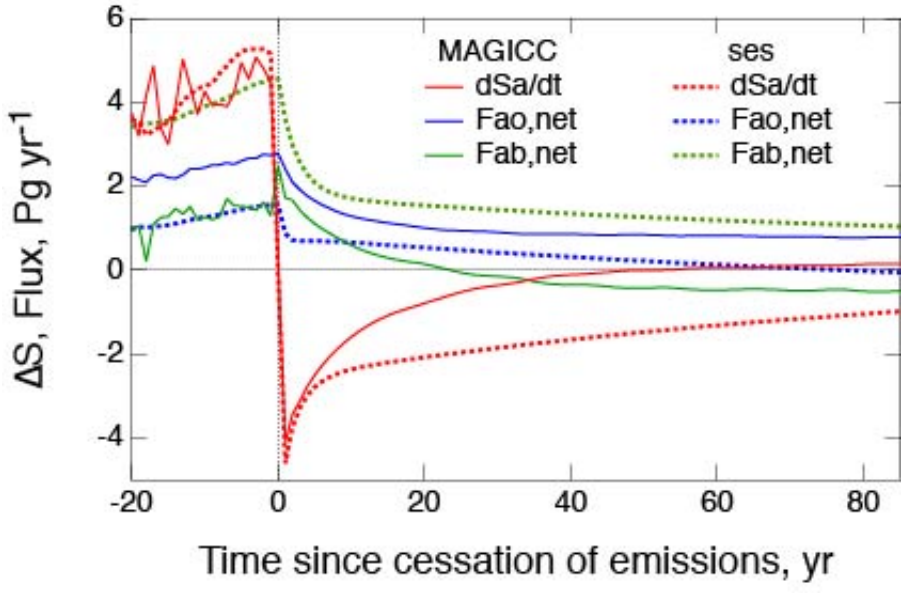

Figure R2. Annual change in stock in atmosphere ΔS and flux from atmosphere into ocean and terrestrial biosphere, F_{ao}, F_{ab}, respectively as function of time subsequent to cessation of CO_2 emissions as calculated with the present model

(preindustrial stock in labile biosphere S^{l}_{pi} = 400 Pg; fertilization exponent *b* = 0.6.) and with MAGICC model.

There has, of course, been much examination of this apportionment, from both the modeling and measurement perspective. Here attention is called to Figure 15a of the manuscript, reproduced in Figure R3 with the results from the MAGICC calculation. This comparison shows that the net anthropogenic flux calculated with both the present model and MAGICC is comparable to measurements, that with the present model a bit lower, that with the MAGICC model a bit higher. (Here the total atmosphere-ocean flux is taken as the total flux augmented by 0.6 Pg yr^{-1} to account for negative preindustrial flux from atmosphere to ocean due to riverine flux from the terrestrial biosphere to the ocean, as described in the manuscript.) Likewise the net flux between the atmosphere and the terrestrial biosphere calculated with MAGICC is slightly lower than the results from the present model calculations and from other current models.

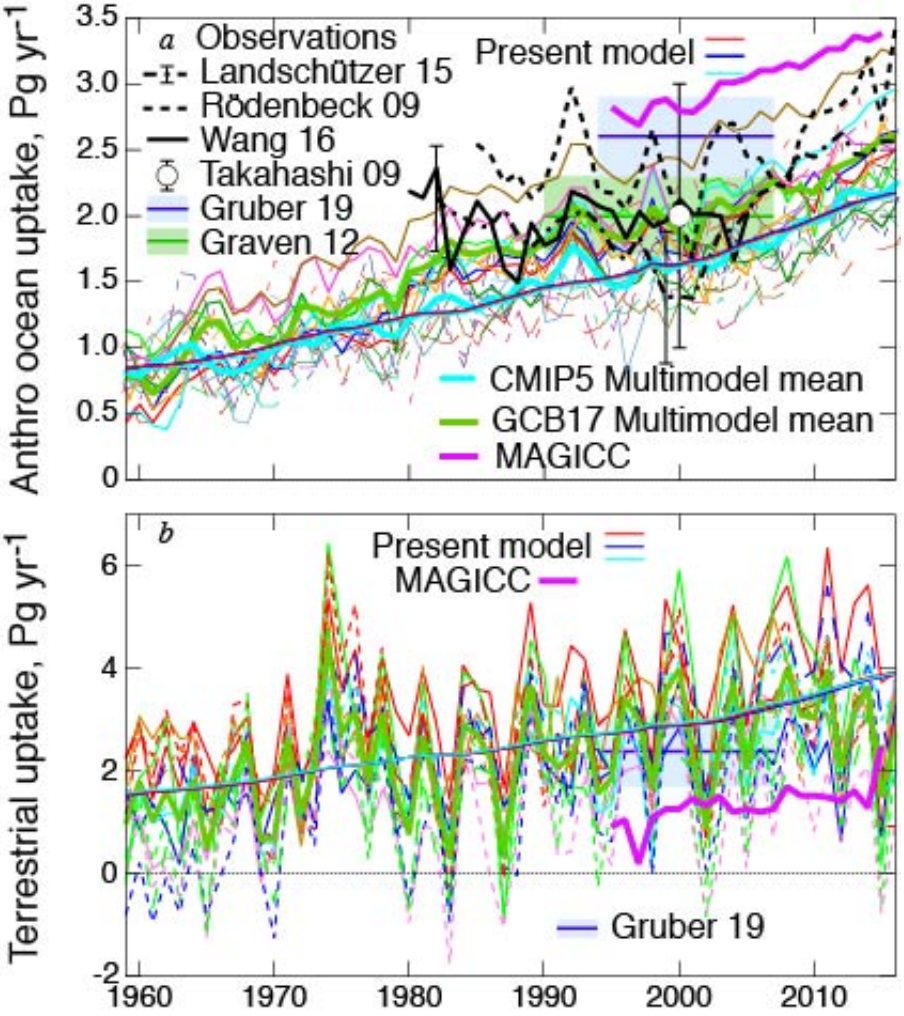

Figure R3. Figures 15*a,b* of manuscript showing also, a, net anthropogenic ocean uptake, (augmented by 0.6 Pg yr^{-1}, as explained in text) and *b*, net terrestrial uptake from MAGICC calculation for years prior to cessation of anthropogenic emissions.

The picture differs markedly subsequent to cessation of anthropogenic emissions, Figure R2, which shows a rapid decrease in the stock of the AC, ΔS_a, in the MAGICC results than in the results from the present model. This decrease is seen to be due largely to the decrease in the net flux from the AC to the TB, with the net flux calculated by MAGICC dropping sharply in the first two decades after cessation, and even becoming negative (i.e., terrestrial biosphere becoming a net source of CO_2 to the atmosphere). The decrease in this net flux must almost certainly be a consequence of increase in transfer rate from the TB to the AC, as the transfer rate from the AC to the TB would be expected to decrease very little, being proportional to the stock in the AC, which decreases only by about 15% in the MAGICC results, Figure R4. Such an apparent increase in the transfer rate from the TB to the AC after cessation of anthropogenic emissions in the MAGICC calculations would seem difficult to explain.

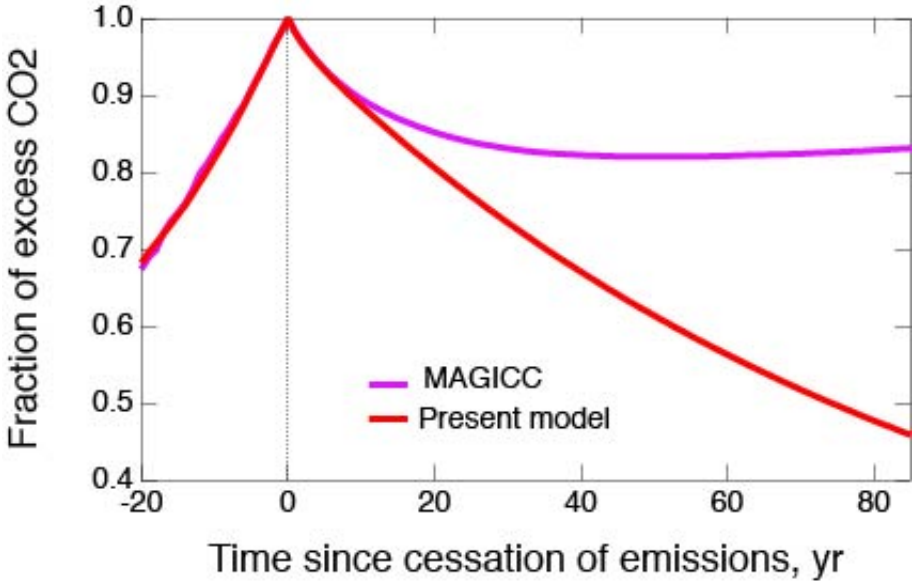

Figure R4. Fractional decrease of anthropogenic CO_2 following cessation of anthropogenic emissions as calculated in the MAGICC model and the present model (preindustrial stock in labile biosphere S^l_{pi} = 400 Pg; fertilization exponent b = 0.6).

In summary, the brief examination of the results from the MAGICC model presented here would seem to attribute the much lower rate of decrease of excess CO_2 following abrupt cessation of emissions obtained with that model compared to that obtained with the present model (and the resulting increase in adjustment time) mainly to treatment of the terrestrial biosphere in the MAGICC model, and raises questions about that treatment.

I turn now to Dr. MacCracken's specific comments.

- Unless one reads very carefully, it is hard to find out that really what is happening after emissions stop is a redistribution of the total amount of injected fossil-fuel CO_2 among the various reservoirs, leaving the atmosphere with a percentage of its peak excess concentration persisting essentially indefinitely—the atmospheric perturbation due to fossil fuels would not naturally go back to zero until all the emitted CO_2 is taken up in

ocean sediments or permanent land carbon. The IPCC has a simple representation of the decay in the atmosphere reservoir that involves, I think it is, five (or maybe four) separate exponentials, each one accounting for the time it takes for the carbon to mix into each reservoir and the coefficient being the share that would remain in each reservoir out to quite long times. The paper does not present the IPCC approximation nor explain how its results compare for each reservoir, and it would be informative to do this.

The sum of decaying exponentials (including a constant term) presented in the several IPCC reports (e.g., AR4, Forster et al., 2007) that Dr. MacCracken refers to does not result from a treatment of this decrease that is governed by a set of eigenvalues or other model properties, nor do the coefficients represent the fractions that would remain in each of the several reservoirs, but rather simply represents a fit to the decay function of a pulse emission obtained with a specific model, the revised version of the Bern Carbon cycle model superimposed on a background CO_2 concentration value of 378 ppm. This decay of a pulse emission, yielding the fraction of a pulse emission at time 0 remaining at time *t*, is convenient for calculating global warming potentials and the like. However it is not directly pertinent to abrupt cessation, for which the fraction remaining at time subsequent to cessation would, in a linear system, be given by the convolution of the pulse emission fraction with emissions prior to cessation. The fraction of atmospheric CO_2 at the time of cessation remaining in the atmosphere evaluated in this way is substantially greater than the impulse response function, Figure R5, the reason being the amount present at the time of cessation would already have been substantially diminished by the rapid, early decrease in amount according to the impulse function response and applying this rapid early response again to the amount at time of cessation results in much greater reduction in amount than would be the case if the impulse response were applied to each year's emissions and then summed, as is evaluated by the convolution.

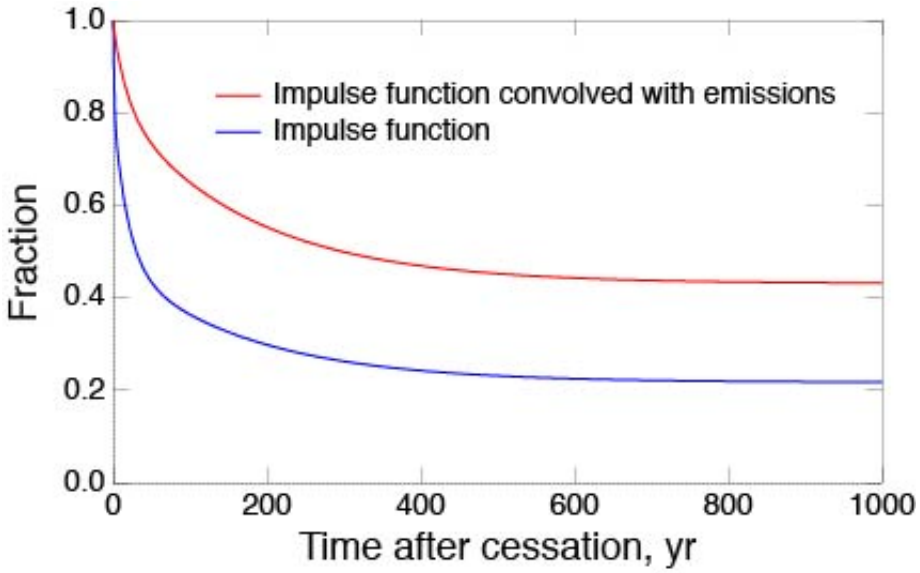

Figure R5. Impulse response function of Forster et al. (2007) and convolution with anthropogenic emissions used in this study (Supporting Information, Figure S4.1; data table, sheet 3).

- It is strange in Figure 7 that the amount of carbon in the terrestrial biosphere looks to

remain constant at the amount that is in these reservoirs at the time of the peak CO_2 concentration, such that, apparently, a reduction in the atmospheric CO_2 concentration does not lead to a reduction in the terrestrial carbon reservoir. This seems very strange given that the equilibration is quite rapid as the CO2 concentration is increasing (FACE experiments seem to show that equilibration for shrubs and grasses is within a few years; trees would take longer). This needs to be checked. Normally, it is thought that as the CO_2 concentration drops, CO_2 will be, in essence, exhaled by the biosphere back into the atmosphere. It would also be worth looking at the transfer from the living biosphere to the long-lived biosphere—this is normally thought to be a pretty slow process. Perhaps I missed it, but it would be interesting to have indications of how much carbon end up in each of the model's 5 reservoirs in the decades and then centuries after the ending of emissions.

It seems that Dr. MacCracken is referring here to Figure 7a, which shows stocks. In the present model there are two compartments within the TB, designated labile biosphere LB and Obdurate Biosphere OB; the latter is diminished by land use emissions, so there are three curves in that panel of the figure that refer to stock in the TB, the curve denoted S_l, stock in the LB; the curve denoted S_o, stock in the OB; and the curve denoted as the integral of the net anthropogenic flux between the atmospheric compartment AC and the OB. As Dr. MacCracken notes, carbon in the LB exchanges rapidly with that in the AC, exhibiting near a steady state relationship (Sec. 8.4 of the manuscript); consequently as S_a increases prior to cessation of anthropogenic emissions S_l (thin brown curve) likewise increases; unfortunately in the figure prior to cessation the thin brown curve is hard to discern because the stock is nearly the same as the stock in the deep ocean S_d, thick blue curve. Subsequent to cessation S_l immediately begins to decrease, again in parallel with decreasing atmospheric stock S_a; (the deep ocean stock continues to increase). The stock in the OB, S_o, light green curve at the bottom of the panel, initially decreases on account of land use emissions; subsequently this stock is augmented by transfer from the AC, so this stock reaches a minimum and then increases. The curve denoted as the integral of net transfer from the AC to the OB, dark green curve, increases throughout the model run as the net transfer continues to increase as a consequence of the increase in the atmospheric stock. As the excess stock in the AC decreases toward the end of the time period shown in the panel, the stock in the long-lived OB compartment levels off toward a constant value. I hope that this explanation satisfactorily resolves the question raised by Dr. MacCracken.

- It is also not clear if there should be a limit on how much the terrestrial carbon can build up in the terrestrial biosphere given that the areas of growth are finite, the supply of nutrients is finite, trees have genetic limits on how large they can become, and one cannot just squeeze in more and more trees—density is limited.

Dr. MacCracken raises an interesting question. In response I would say that the parameters of the model are based on observational constraints over the industrial period. The fertilization exponent so determined (range 0.4 to 0.8) is well within the range of current observation- and model-based estimates (Appendix C, Figure C1). Subsequent to cessation atmospheric CO_2 would decrease and whatever constraints there are on the rate of uptake of carbon by the TB would likely become less restrictive, and fertilization would diminish. However, if atmospheric CO_2 continues to build up, the constraints noted by Dr. MacCracken might become increasingly important. That said, it would seem that that concern would not apply to the consequences of cessation of anthropogenic emissions examined here.

- It would be interesting to see model results compared for more cases than simply a

sudden cutoff of peak emissions to zero. So, what about a gradual reduction in emissions? What about if emissions continue? This is important because emissions are not going to suddenly go to zero, but will decrease slowly and so there will be very little adjustment right at the zero point as the adjustments have been going on all along as the emission levels drop.

I concur with Dr. MacCracken that the consequences of such gradual future emissions reductions would be interesting. The model developed here is well suited to such a profile of emissions. This is readily done with the present model; just specify the emissions profile and re-run the model. As for the consequences of continued or increased emissions, the resultant increase in atmospheric CO_2 and indeed in carbon in the several other compartments would gradually take the model out of the regime in which the parameters of the model were determined, as for example the limits on fertilization that Dr. MacCracken notes in his Question 3.

- As near as I can tell, the model looks at the perturbation of the carbon cycle due to humans. It would be interesting to know if one simply put the amount of total carbon in the various reservoirs in the preindustrial period if it would distribute to the reservoirs as observed and if this would be a steady state condition. And then, again from baseline amounts, it would be interesting to know what the distribution might be for some larger amount of carbon. Given that the end point of the study described in the paper is the redistribution among reservoirs, is the model set up so the observed distribution would result were the preindustrial total loading of C being the initial condition—right now it appears that the preindustrial distribution is prescribed rather than used as a test of the model representation.

The model is quite stable if emissions are set to zero. However, the time constants associated with the two long-time-constant reservoirs are quite long, half millennium to multiple millennia, so taking a fixed amount of carbon and perturbing the distribution and then waiting for steady state to ensue would take a long time in the model. I have not yet tried that. In response to the last sentence, the preindustrial distribution is indeed prescribed, based on observations; the fact that it is stable in the model in the absence of perturbation of emissions is a consequence of the transfer coefficients having been determined based on that set of preindustrial stocks.

- Note—I have not reviewed the appendices nor examined all the equations in detail, generating my comments based on the text that is presented; and this was all done pretty quickly.

Specific Comments and Thoughts/Suggestions (noted as they come up, and sometimes repeated if points arise again):

Lines 23-24 (and elsewhere): My understanding of the current approach is that, in representing the decay time in one equation as opposed to a model with each process separately represented, there are five or so time-decaying exponentials, so separately into the biosphere, surface ocean, deep ocean, sediments, soils, etc. and so characterizing the time as if it is one number is just not a correct characterization of how the number is currently viewed. That is, there will early on be a short time constant for emissions into the atmosphere to redistribute some into the mixed layer and living biosphere, but that once this occurs, the time constant will be much longer for the redistribution into the deep ocean and long-term terrestrial biosphere.

There are five differential equations governing the evolution of the stocks in the five compartments of the system. Because of conservation of matter, the system

actually reduces to 4 differential equations. If the system were linear there would be four eigenvalues. See, for example Raupach (2013; reference in manuscript). Because of carbonate-bicarbonate chemistry the system is "slightly" nonlinear. With such a system of equations one can solve for time-dependent eigenvalues that can be related to the time constants governing evolution of the system. I have not done that with the current model but am rather confident that that would be the outcome.

Lines 28-29: A bit strange to be quoting Ramanathan when it was Revelle, I think it was, who talked about this as a 'great geophysical experiment.' I don't know reference for this, but it might be the 1965 report of the President's Science Advisory Committee (PSAC) in the chapter (or annex) on climate change that Revelle chaired or in his paper with Harmon Craig, I think it was, on ocean uptake of C-14 and/or other species.

Touché; I will change the attribution.

Line 34: I don't really like the idea of saying just "anthropogenic CO_2" as if the actual molecules are different than the CO_2 in the air. Actually, the adjustment that is occurring is based on all of the CO_2/carbon that is in each reservoir, and one cannot just do a difference of the anthropogenic CO_2 amounts. I think this is important to be careful of because, on the other hand, the C-14 generated by the nuclear testing is so dominant in amount that it is, as I recall, most of the C-14 in the atmosphere, and there is also radioactive decay going on of the amounts over time.

Dr. MacCracken is quite correct that one should not refer to "anthropogenic CO_2" or anthropogenic stocks in any of the compartments (unfortunately I lapsed into that terminology, although would seem reasonable to refer to the CO_2 emitted from fossil fuel combustion as anthropogenic CO_2). Better terminology would be the anthropogenic perturbation of the stocks and fluxes. That said, I think it is reasonable to take the difference in a given stock or flux between the unperturbed (or preindustrial) state of the system and that of the system that has been subject to the perturbation and denote that perturbation as the anthropogenic perturbation. Thus for present CO_2 mixing ratio 420 ppm and preindustrial 278, it seems perfectly reasonable to refer to the difference 142 ppm as the anthropogenic perturbation (increase) in the atmospheric mixing ratio. I note that in the caption to Figure 2 I wrote "perturbations resulting from anthropogenic emissions in red."

The rate of decrease of the perturbation of atmospheric $^{14}CO_2$ due to weapons testing in the 1950's and early 1960's, and then abruptly halted, is unfortunately often confounded with the rate of decrease of the anthropogenic CO_2 that would ensue in the absence of emissions.

Lines 40-41: Are there really estimates that suggest the lifetime of the perturbation is only a few years or a few tens of years? Those seem very, very low if one is talking about the lifetime of the full increase in the anthropogenic loading. I do understand that some experts do say that there would be a fast initial drawdown, but this would not take out anywhere near all of the overall anthropogenic loading, as I understand things. So, is this a comparison of apples and oranges? I see paper does cover this somewhat in the following sentences, but it seems to me confusing to be giving a time constant based on an initial slope (based I would imagine on how the airborne fraction allocation is going on) that does not lead to all of the perturbation being removed.

Archer et al. (2009) wrote: "There are rival definitions of a lifetime for anthropogenic CO_2. One is the average amount of time that individual carbon atoms spend in the atmosphere before they are removed, by uptake into the

ocean or the terrestrial biosphere." My own view is that that is a bit of a straw man; I don't think anyone would really take that as the lifetime of excess CO_2. But it leads to use of the rate of decrease of excess $^{14}CO_2$ after cessation of weapons testing as a measure of lifetime, and that has been invoked numerous time.

As for a few decades, see the works cited at lines 63-64 of the manuscript.

Line 54: CO_2 does not really "decay"—it appears as part of different compounds or in different forms, but does not really disappear the way C-14 does. Strange word choice.

I concur; better "decay of excess CO_2"; applicable for example in characterizing the curve of Allen et al. in Figure 1 at line 67 as showing "essentially no decrease of CO_2"; I should have stated "essentially no decrease of excess CO_2"; I shall look carefully for other instances of this. Precision in language is important.

Line 64 (regarding Figure 1): It appears that the approach does not account for the increasing time it takes to mix into the deep ocean and then into the sediments, so that the atmosphere gets to its adjustment time quite quickly. This seems incomplete because the time constants for spreading through the longer-term reservoirs do not seem to show up as does appear in other model estimates. This seems a bit strange.

I think that Dr. MacCracken is referring here to processes that occur on millennial to multimillenial time scales. As explicitly stated in the manuscript at line 188 this study focuses on the time scale over which, in the absence of atmospheric emissions, the amount of CO_2 in the atmosphere would recover substantially to its preindustrial value. This is entirely in keeping with the qualitative definition of the lifetime of anthropogenic CO_2 advanced by Archer et al. (2009) quoted at line 172.

Lines 83-84: On the relatively rapid response of surface temperature, that is really because of the relative magnitudes of the heat capacity of the various reservoirs. So, the atmospheric heat capacity is equivalent to less than three meters of ocean water, and so with the upper ocean being roughly 100 meters deep, the atmospheric temperature tends to follow the upper ocean temperature. The deep ocean is of order four kilometers deep and the flush through time of the upper ocean is of order 25 years, but for the deep ocean is more like 1500-2000 years. And then there is the fact that the IR emission at the surface of the ocean will shorten the temperature adjustment time as well, so not surprising. It is not clear to me that this can serve as an analog for the carbon cycle adjustment time in that the relative amounts of substances in the various reservoirs is more even and there is no real loss term the way there is an IR loss term for the energy.

Yes, the top 3 m of the ocean are going to accommodate rapidly to perturbations in radiation balance, but this heat is rapidly mixed throughout the mixed layer ocean, rapidly in an absolute sense (seasonally; see e.g., Sutton and Roemmich, 2001) but rapidly also relative to much slower perturbations in forcing. As to analogy of heat loss by infrared radiation, loss of carbon to the obdurate biosphere would seem a fair analogy. Another major difference between CO_2 and heat is that heat is much more "soluble" in ocean water than is CO_2, in the sense that virtually all the excess heat is in the ocean, whereas the distribution of anthropogenic CO_2, Figure 2, is only 10% of the amount in the AC and ML is in the ML. That is the reason that transfer of excess heat from the upper compartment of the two-compartment system to the lower compartment is much more rapid than for CO_2: a time scale of 10 vs 100 years.

Line 103: As I think I have said before on this approach, I think the two-layer approach to

representing the upper and deep ocean has been shown to be fatally flawed in representing the ocean—basically, over time, the two boxes would tend to have, for example, the same temperature as opposed to a warm ocean on top and a cold ocean below. That is why the way the simple models are set up is with a one column ocean with multiple layers and slowly rising waters, and then a pipe from the upper ocean directly to the deep ocean to represent polar downwelling of cold dense waters (see the explanation in the documentation of the MAGICC model). And for carbon in the main column with its multiple layers, one would have to represent carbon going down a bit by biospheric action and then dissolving on the way down and being carried back up. I also wonder how you are (or are not) representing the compensation depth (so the depth at which CO_2 tends to dissolve from sediments, and that is being affected by ocean acidification). So, just to note, I am already now very suspicious of the approach. Given that such an additional representation would not add substantial time or complexity to your model, I'm surprised that you have not done this.

This point is discussed above. But I would add that the difference between the results with the present model and with the MAGICC model seem to be much more due to differing fluxes into and out of the TB than deep ocean.

Lines 115-117: Just to note that virtually all of the observations to be used for constraining the model are for a situation where the CO_2 is increasing, and assuming reversibility seems to be a rather significant assumption (e.g., waters are thermally stratified—and advection is not really reversible, etc.). So, if you had an upwelling-diffusion model (so downward polar pipe and upwelling column), there would not be the same reversibility that the model now seems to have. And, I might ask, is the living biosphere flux reversible—will a lower CO_2 level lead to a lower amount of biomass, as one would expect?

A premise of the present model is that the processes affecting CO_2 transport (and heat transport) are the processes that have been affecting these transports from preindustrial times up to the present; that the excess CO_2 or heat is piggybacking onto the processes that have been responsible for such transport in the unperturbed preindustrial world or minimally perturbed anthropogenically influenced world, the argument ultimately resting on the net forcing over the industrial period of, say, 2 W m^{-2}, being small compared to the net absorbed solar irradiance of 260 W m^{-2} that drives the whole system.

Lines 163-165: I agree there is much confusion, especially about the difference between the lifetime of a particular molecule of a substance and the lifetime of the perturbation to a concentration in situations where molecules of the substance are going in both directions across an interface. There is also difficulty when the substance is spread among multiple reservoirs that each have particular and quite different exchange times (hence the IPCC's five or so decaying exponentials).

Both points addressed above.

Lines 172-178: Well, that is a start at the problem—there are then multiple components and varying types of transfer processes, not just diffusion, etc.

The reference is to the quotation from Archer et al. (2009) defining the lifetime of interest as the time scale over which atmospheric CO_2 would substantially return to its preindustrial value in the absence of anthropogenic emissions. I concur in the complexity of the processes governing the decay of the long tail of excess atmospheric CO_2, but the focus of the present study is the time scale over which excess atmospheric CO_2 would decrease substantially from its value at the time of cessation and the processes governing this decrease.

Lines 226-227: But is the labile exchange time the same in both directions? It is quite fast if one has new plants growing, but once created, decay can take time and I'd think there can be rather long lag terms if one grows actual trees—which might keep growing even though the CO_2 is down a bit. I'm just not sure that I agree there is an equivalent exchange in both directions—yes, for leaves, etc., but I'd not think that the case for new wood that is created, which could have a hundred-year return time.

The response of atmospheric CO_2 to an abrupt upward perturbation in emissions is a very interesting question, in model world and in the real world. Although this question is not addressed in the present study, it could be readily examined in the present model. Clearly in this model, and I think pretty clearly in the real world, the ML ocean would rapidly (time scale of a few years) reach near equilibrium in response to such an upward perturbation in emissions. Likewise it would seem that the LB would rapidly reach a new steady state with the atmosphere. Processes governing longer-term response would be transfer from the ML to the DO, governed by the transfer coefficient k_{md}, de-rated by the fraction of the excess CO_2 that is in the ML upon the sum of excess carbon in the AC + ML + LB. Likewise, at least in model world, the longer time response would be governed by the transfer coefficient k_{at}, here de-rated by the fraction of the excess CO_2 that is in the AC upon the sum of excess carbon in the AC + ML + LB. These would yield the same time turnover time of the stock in the AC + ML + LB as obtained in the present analysis, and seemingly the same adjustment time to an abrupt upwards perturbation. At least in the present model the adjustment time would be decreased by increased rate of uptake of CO_2 by the TB, as governed also by the fertilization coefficient, Appendix C, as constrained in this study to the range 0.4 to 0.8 by comparison of model results with observations. However response to an upward perturbation in CO_2 emissions would be going outside the domain that is constrained by observations. Here it might also be noted that the perturbation in GPP over the Anthropocene, as given by the range of fertilization exponent determined here, 0.4 to 0.8, is 17% to 36%, a not insignificant increase, so in the real world there might be limitations on future uptake of CO_2 by the TB that have not been seen thus far.

Line 272: I'm just not up to trying to work through all the equations, so will be responding to the text. I'd just note that the labile cycle of uptake as leaves and wood would be different, as would the times of decay—I'm just not at all convinced such a simple model will be sufficient (so a caveat I am holding, waiting to hear about the tests being run). What is normally done, as for MAGICC, is to calibrate the simple model versus complex models. While I understand you want to calibrate directly, I think there is really the need to explain why the difference with the model is occurring and whether the assumptions made in gong to the simple model capture this.

The equations Dr. MacCracken refers to are for several different measures of adjustment time and are independent of the mechanism governing the decay of excess CO2 following abrupt cessation of anthropogenic emissions. Dr. MacCracken then turns to specific removal processes. As Dr. MacCracken notes the MAGICC model is calibrated against even more complex carbon cycle models (and more broadly against complex climate models). As seen in Figures 14c and 15b of the present manuscript and elsewhere there is enormous intra-model spread in the extent of uptake of carbon by the TB in response to historic emissions, to abrupt cessation of anthropogenic emissions (MacDougall et al, 2021, replotted in Figure 1c of the present manuscript) or to an emissions pulse (e.g., Joos et al., 2013). Such spread in results would seem to make these models a weak reed upon which to lean. As noted above, it seems that the results from MAGICC (or at least from the version of MAGICC available for online use) for the response of the TB to an abrupt cessation of emissions seems

fundamentally wrong. Consequently it would seem that comparison of rates or extent of changes in stocks in the several compartments and in the fluxes between these compartments in response to perturbations, such as abrupt cessation of emissions examined here, would be extremely informative of the reasons why the models differ so greatly.

Lines 280-285 or so: The problem with using a half-life or 1/e times is that even small amounts of the perturbation affect the climate—there really is not a tolerance level.

This comment refers to the effects of the long tail that would persist long after the much more rapid substantial decrease of excess CO_2, 80-85%, that would, according to the model developed in this study, take place n a time scale of 100 years or so. The long tail is not the subject of this study. This study is focused on the time period over which the excess atmospheric CO_2 (and that in the mixed layer ocean and the labile terrestrial biosphere, which are tightly coupled to the atmosphere) would decrease substantially from the value at the time of hypothetical abrupt cessation of emissions.

Line 291: Given how much discussion you have here I'm surprised that you don't present the IPCC decay function, so the sum of five (or so) exponentials for comparison—does your model agree with the decay times the IPCC has for each of the terms and for the fraction going into that reservoir?

This is responded to above, Question 1.

Line 317/Line 3010: Just a note on Figure 2 that the number 120 for F_{al} has the 1 sort of lost due to some sort of cropping so looks like 20.

Good catch.

Line 360: Ah, interesting, and good to hear of the separation.

The reference is to distinguishing the short- and long-lived components of the biosphere. There is precedent for this, references in text. An important study by Naegler and Levin (2006) examining decrease of atmospheric radiocarbon after cessation of atmospheric weapons tests used three compartments, but as stated in the text, line 1600, those investigators found that with two well-mixed compartments the simulated stocks agreed rather well with the observations and that the agreement was somewhat improved for three or four compartments but raised concern that with increasing number of compartments the parameters became increasingly underdetermined.

Lines 423-424: Treating everything below the mixed layer as the deep ocean seems to me to be a serious oversimplification, given how the isopychnals really control downward transport in middle and low latitudes, allowing a good bit of horizontal mixing from surface ocean to depths down to several hundred meters, but not down below 750 meters or so due to the upward motion that is balancing polar downwelling.

As stated above, response to General comment G2, the two-box model works rather well for heat transport and thus should be expected to work similarly well, with observationally determined transfer coefficients, for dissolved inorganic carbon (DIC).

Line 446: I agree this is a curious anomaly. I started drafting a note for AGU/EOS many years ago but never got to completion. One of the interesting thoughts I heard about it had to do with possibly the spreading of weeds, etc. over croplands as a result of so many

men from farms being pulled into the military (as I recall, the flattening started early in the war years), and then perhaps a change after World War II as weeds were cleared and then in the amounts of cropping of C3 versus C4 plants and relative carbon uptake. This all got to be more than I felt that I could get into and so I sort of abandoned the paper—but it seems to me finding an understanding of this pause (which did not seem to be in emissions) might be very insightful. Perhaps with your model, that could be something to look into. [And I should note that I think there is a paper or more looking at this in detail, perhaps by an Australian author or two.]

The reference is to the flattening of the curve of atmospheric CO_2 with time in the 1940's. There is general concurrence in the literature that this is a curious anomaly (Trudinger et al., 2002; Bastos et al., 2016; Rubino et al., 2019, cited in the manuscript).

Lines 467-468: I really don't like this notion of a piston velocity to represent the links between the mixed layer and deep ocean. I'd really suggest trying a different representation of the ocean and seeing if your conclusion holds up.

The piston velocity is simply an alternative means of expressing a transfer coefficient in a two box model, in favor of which arguments are presented above.

Line 476: On the flushing time of the deep ocean, 650 years seems lower than what Broecker, etc. have talked about, which I think is more like 1000- 1500 years or so. And I'm not sure that flushing time is the right way to think about it instead of as more like a pipe that it takes 1000-1500 years or so to pass through, so the increased uptake of CO_2 does not become available to the atmosphere until after that much time [though isopychnal mixing does lead to some higher amount of CO_2 down several hundred meters (or even more) in the lower latitudes]. Again, I think that the ocean component of the model needs to be upgraded. And I'd also note that I don't think ocean processes are simply reversible—time constants on that will apply.

The present results are quite insensitive to the return time from the DO to the ML. As shown in Figure 1, the flux from the DO to the ML for central estimate of the transfer coefficient (or piston velocity) is 0.2 Pg yr^{-1}, relative to the downward flux of 1.8 Pg yr^{-1}. Doubling or zeroing this return flux would change the net uptake by the DO from 1.6 Pg yr^{-1} to 1.4 or 1.8 Pg yr^{-1}. A much greater contribution to the uncertainty in this flux is that associated with the piston velocity, which is responsible for the 33% uncertainty in this flux.

Line 486: So, if you change the ocean to a deep ocean with a pipe going down from the surface in high latitudes and then the water spreading out, this provides the basis for slow upwelling to occur as new amounts of downwelling water push underneath and lift it up. So, the downwelling flow is a result of dense, cold water sinking and the CO_2 amount is based on how much CO_2 can be held at equilibrium by the cold water with the increased atmospheric concentration. This downward transport of CO_2 would then increase only slowly if the overturning circulation stays the same, even with the atmospheric concentration going up. And the water coming up into bottom of the thermocline layer (at 750 meters) would be staying at about the same loading as preindustrial due to the long circulation time. This would all be very different than how things work with a piston type approach to representing the downward flow. An addition advantage of actually representing the over-turning circulation would be that you could then do experiments changing the amount of the ocean overturning (with some suggestions that this overturning amount would be caused by Atlantic surface warming and reductions in the amount of rejected dense brine water as sea ice freezes—and some indications this change is already occurring). Were the overturning circulation to stop, the net uptake of CO_2 from the atmosphere and downward transport would necessarily go down a lot (also

less upwelling of nutrients into the mixed layer to feed the biospheric pump of C to the deep and intermediate ocean). With your piston approach, there would not be such a reduction in the flux—it would just keep going (wouldn't it?). I guess what I would recommend is to apply your model framework to the energy in the system and see if the resulting vertical temperature distribution would result—basically, the transport of CO_2 and heat occurs in the ocean in the same way (there is no radiation term to have energy jump from one layer to a much different layer), so whatever structure you have should work for both—and getting the deep ocean to be cold won't occur with a piston velocity.

This comment seems to give a more nuanced justification for a more sophisticated description of fluxes of heat and DIC between the ML and DO. This point is discussed above. I concur that it might be of great interest to examine the consequences of such a more sophisticated description. Perhaps that is reflected in the MAGICC model that Dr. MacCracken advocates. Here I would note only that the nominally more sophisticated approach of that model (at least in the web-based version, the results of which for the abrupt cessation calculation are shown in Figure R2 above), yields a net transfer rate of CO_2 from the atmosphere to the ocean that is 0.5 to 1 Pg yr^{-1} greater than that obtained with the present model, i.e., in the direction of a shorter adjustment time in the MAGICC model than in the present model. The principal reason for the shorter adjustment time in the present model than in MAGICC is not ocean transport; it is uptake by the TB. That is not to say that better understanding and more accurate model representation of the processes that control transport from the ML to the DO is unimportant. But the principal source of difference between the two models contributing to the shorter adjustment time found in the present model is uptake of carbon by the TB.

Lines 520-523: Here I totally disagree (or misunderstand—perhaps you are just referring to the amount of wind-driven exchange that is going on—and not the net transfer). The transfer from the atmosphere to the mixed layer depends on the gradient in the CO_2 concentrations and I don't understand how you can say this is constant given the annual increase in the CO_2 concentration has changed over time due to changes in emissions (is not what you are doing assuming instant equilibration of ocean and atmosphere instead of allowing a gradient to form and drive the flux into the ocean?). Also, with the colder (saturated) water rising up in lower latitudes, it emits CO_2 as it warms and then in high latitudes, as the ocean waters cool, they take up CO_2. So, I just don't understand how you can say the gradient will be constant over time. And as the overturning circulation changes, there will be a change in the flux—so how can you make an assumption about it being constant? On the mixed layer to deep ocean fluxes, etc., there are the downward flux flows (in the downwelling waters and the biological pump) and then the upward flux flow in the slowly rising waters that went down long ago and are now rising with the CO2 burden of the past. I think each of these fluxes needs to be kept track of separately. Also, given that the biological activity is dependent on the amount of nutrients that are carried upward—if the overturning circulation changes, then the flux of nutrients change. I get the sense that your model would work reversibly whereas with a more complete representation, this would just not be the case—things are not just instantly reversible.

Dr. MacCracken raises (or reiterates) several points here that require response. In general the processes that control the change in transport and uptake of anthropogenic carbon in the several reservoirs depend on the differences between the time-dependent amounts of carbon in these reservoirs and the preindustrial amounts. (The same can be said for heat.) To stay with heat for a moment, it should not be inferred that the present treatment requires that. at steady state characteristic of the preindustrial world, temperature, the intensive variable that drives heat transport is required to be uniform throughout the world ocean, but rather that under the perturbed condition the net flux of heat

from the upper to the lower compartment is driven by the difference in the departure of temperature from the preindustrial value in the two compartments, ΔT_u - ΔT_l, suitably averaged over space and time. Similarly net transport of CO_2 from the atmosphere is driven by the difference in the departure of fugacity (the intensive variable) between the two compartments from that in the preindustrial state, suitably averaged over space and time, i.e., globally and annually. Dr. MacCracken notes that "with the colder (saturated) water rising up in lower latitudes, it emits CO_2 as it warms and then in high latitudes, as the ocean waters cool, they take up CO_2." This phenomenon takes place in the base case and in the perturbed case; what is not so much the spatial or temporal differences in fluxes, but rather the departures of these fluxes from their preindustrial values, which give rise to the net uptake of CO_2 by the world ocean.

To elaborate, locally and instantaneously the rate of transfer of CO_2 (or heat) depends on the wind-speed driven transfer coefficient and the local difference in fugacity (or temperature) between the atmosphere and ocean. That is the wind driven exchange that Dr. MacCracken refers to. It is this local wind–speed- and fugacity–difference-driven flux that governs the gross transport of CO_2 between the atmosphere and the ocean, estimated as 70 Pg yr^{-1} at preindustrial conditions, 102 Pg yr^{-1} at present conditions, as given in Figure 2 and Table 1 of the manuscript based on multiple references. This gross flux is much greater than the net flux, which is driven by the globally averaged fugacity difference and transfer coefficient. Because equilibration between fugacity in the two phases is rapid (roughly a year), the net uptake of CO_2 by the ML ocean is insensitive to the value of the transfer coefficient. Another comparison is the net global average CO_2 flux from atmosphere to ocean of 1.6 ± 0.9 Pg C yr-1 (Takahashi et al., 2009), evaluated from 3 million determinations of local (4° x 5°) and instantaneous (monthly average) fluxes, evaluated from wind speed and fugacity of CO_2 in ocean water. This global average flux, which corresponds to 0.046 g m^{-2} $month^{-1}$, is two orders of magnitude less than the range of the fluxes in individual latitude and longitude boxes in given month, -9 to 9 g m^{-2} $month^{-1}$, and, given that this range is for monthly (4° x 5°) averages, the actual range would be expected to be substantially greater.

Lines 586-587: Nice to use a transfer coefficient for heat, but just to note that your ocean would not lead to a cold deep ocean, and so it would really be inappropriate to be using heat influx for transfer into the deep ocean—you really need to do a better representation of the ocean.

Just to reiterate, the preindustrial steady state of the world ocean does not require uniform ocean temperature. And the departure from steady state is not in temperature *T* but in departure of temperature at a given time relative to preindustrial temperature ΔT.

Line 655: On the issue of CO_2 fertilization, there are also limits imposed by the supply of nutrients and of water, and so there would seem to be a need to be very careful about this, especially as the area of the relatively dry subtropics is growing.

The issue of CO_2 fertilization is a major uncertainty in the real world and in model world. Figure C1 of the manuscript plots gross primary production GPP vs CO_2 mixing ratio x_{CO2} as if x_{CO2} were the sole controlling variable. But, as Dr. MacCracken notes, undoubtedly there are other phenomena controlling GPP, so a plot of GPP against x_{CO2} is aliasing these other phenomena into x_{CO2}. In the present model the uptake of CO_2 mixing ratio x_{CO2} as if x_{CO2} were the sole controlling variable. But, as Dr. MacCracken notes, undoubtedly the obdurate biosphere is treated as dependent on the fertilization exponent *b*, which is the

slope of a graph of log(GPP) vs log(x_{CO2}). Use of the model for future times requires assuming that the aliasing of the variables that control GPP into dependence on x_{CO2} would continue into the future. In application of the model to examine the consequence of net cessation of emissions, the CO_2 mixing ratio stays within the range of values used to constrain the model parameters, lending some justification to application of the model within this range. Extending the model to higher values of x_{CO2} would become increasingly uncertain.

Lines 688-691: So, have you run your model from preindustrial concentrations into the future assuming no emissions, such that the model holds the CO_2 concentration constant? I would think that this would be something to explain at the start of this subsection as proof that the model is stable in the absence of emissions. Again, I'm not really clear on how you are calculating the fluxes into the ocean and biosphere. So, are these fluxes being driven by the gradient created each year by the emissions—if so, then if emissions go to zero, there would no longer be a gradient, and so how would the flux continue to be the same? Or is this turnover time you calculate based on the current flux rate driven by the annual emissions. I just don't see how this number of 44 years accounts for the fact that as emissions go to zero, the gradient driving the flux would go down and so then would the flux? So, shouldn't the net flux from the atmosphere to the mixed layer go down exponentially over time, and do so quite rapidly as you are assuming a quick adjustment time of the atmosphere and mixed layer? Now, it might be that in your model the mixed layer to deep ocean flux would stay the same and so this would keep pulling down the mixed layer concentration and so then sustain the atmosphere to mixed layer flux. If so, then, again, making sure that you have the ocean exchanges properly represented is critical, and as I've said, I don't like the ocean circulation that you have. Basically, what will happen over a thousand years or so is the amount of C would re-equilibrate so that the total fossil fuel burden is spread through the upper and deep ocean (and as that process pulls down the CO_2, it will pull CO_2 back out of the labile biosphere reservoir). So, on your rate, is that an initial rate? I don't see how that can persist as the atmospheric concentration is pulled down as the time constant for mixed layer and deep ocean is so long.

In this regard, it would also be interesting to know what sort of steady equilibrium occurs with different total amounts of C in the combined set of reservoirs. So, for the amount in the preindustrial world, the equilibrium is 278 ppm (say), so would your model come to that equilibrium at that level with the total amounts of C in the non-sediment reservoirs? What about with 50% more? Or are your equations all based on the departure of the system from the 278 ppm base level, so you will inevitably come out at that if the distribution among reservoirs is not started as it was observed to be (so, what would happen if you put all the C in the deep ocean at t = 0 and then ran to equilibrium; what about if started with all C in the atmosphere—would the preindustrial CO_2 level result?

Dr. MacCracken raises multiple questions here.

First, In the absence of emissions the model is quite stable. This is not surprising as the coefficients governing transfer in the preindustrial state are obtained from observed preindustrial stocks.

Second, the rates of transfer of CO_2 between reservoirs depend on the transfer coefficients and on the stocks in the leaving compartments, but not directly on the rates of change of those stocks, only as the rates of change affect the stocks in the leaving compartments. For example, with respect to transfer from the AC to the ML and between the ML and the DO, as illustrated in Figure R6, the net transfer rates are only slightly changed across the cessation; the abrupt downward jump in $F_{amNetAnt}$ results from the stock in the AC no longer being increased by anthropogenic emissions; the magnitude of this jump is roughly

10% of annual emissions, consistent with the ratio of annual increases in ML and AC stocks shown in Figure 2 of the manuscript. The changes in signs of the slopes before and after cessation result from transition from increases to decreases in the stocks.

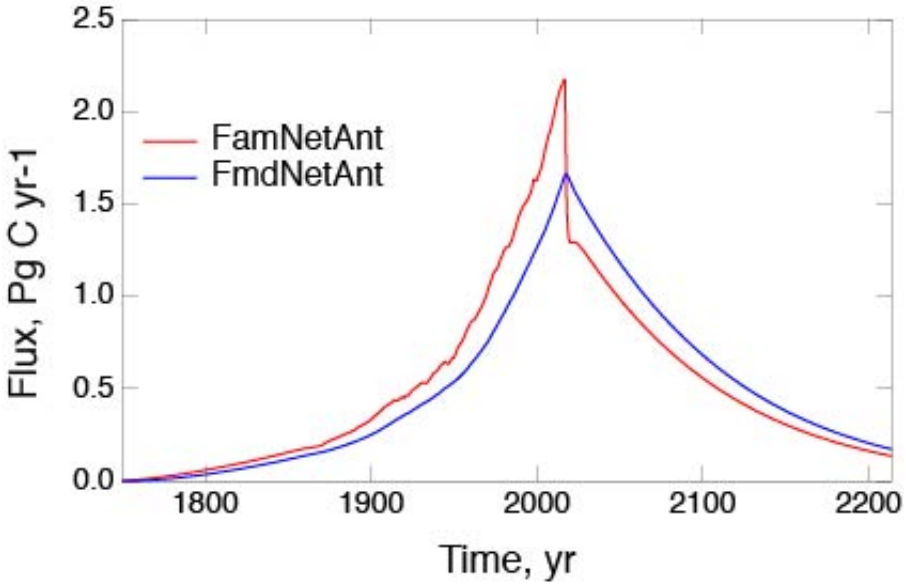

Figure R6. Net fluxes from the atmosphere to the ML ocean and from the ML ocean as calculated with the present model, preindustrial stock in labile biosphere S^{l}_{pi} = 400 Pg; fertilization exponent *b* = 0.6. Cessation of emissions is in year 2017.

Third, with respect to turnover time, as emphasized in the manuscript, multiple turnover times can be defined, and it is therefore essential to identify the turnover time pertinent to the system. This identification is based time constants characterizing exchange of perturbations in stocks between compartments to ascertain whether the compartments should be considered as coupled. As stated in the manuscript, the turnover time for the atmospheric compartment by Eq 5.1, as if the AC were not coupled to any other compartment, 44 yr, leads to a substantial underestimate of the adjustment time.

That said, Dr. MacCracken describes fairly accurately what would happen subsequent to cessation. Indeed the sink from the ML to the DO sustains the flux from the AC to the ML, although at a decreasing rate that more or less resembles exponential decay. Also the coupling between the stocks in the AC and the ML is characterized by an equilibration time (not an adjustment time, that term referring to the rate of decrease of stocks subsequent to cessation of emissions).

After again briefly bringing up the need for accurate representation of ocean circulation, Dr. MacCracken returns to the subject of what would happen over a thousand years or so subsequent to cessation. Once again I would stress that that the residual long small tail of excess CO_2, while not unimportant is well beyond the time scale over which substantial reduction in excess CO_2, the principal subject of this study, would have taken place. As shown in Figures 12 and 16 of the manuscript, the rate of decay of excess CO_2 does indeed greatly decrease on the long time scales Dr. MacCracken refers to, but by that time following cessation of emissions, the great majority of the excess atmospheric

CO_2 would have been taken up by the deep ocean and the obdurate component of the terrestrial biosphere.

Lines 700-701: Okay, so one exponential is the time for the atmosphere and mixed layer to equilibrate. This does not get rid of the overall perturbation.

I think Dr. MacCracken meant to say one eigenvalue or one time constant.

Line 708: Missing word—should be "it is useful"

Yes.

Line 726: Typo—should be "the situation"

Yes

Line 736: Just to note that the water coming up can be supersaturated in CO_2 as a result of the dissolution that occurs as the biological pump is working so that CO_2 keeps getting recycled up. Presumably, as the CO_2 concentration goes down, this will be affected as well. Again, I think there is a real problem in an overly simplified ocean model.

I should probably be more careful in language. Better "transfer of anthropogenic DIC" than "transfer of DIC". All statements about sinks should be made relative to the anthropogenic component, the difference between the total quantity and the preindustrial.

Lines 750-751: Does this large flux back and forth allow for the seasonal build-up and release of C by the TB as is indicated by the annual cycle of the CO_2 concentration as seen in the Mauna Loa record? Might that variation suggest you need to subdivide the LB box?

The model is dealing with annual quantities. In the real world there is a seasonal cycle, but not in the world of the present model .

Line 756: This just can't be the case for trees growing? The FACE experiments might suggest there is a few-year equilibrium for weeds and grasses, but I don't see how this could be the case for wood as it would take much longer times for a full forest to come to equilibrium. So, again, perhaps the LB needs to be subdivided.

In the present model, the rate of transfer from the atmosphere to the LB was taken as the GPP (as increased from PI by the fertilization effect); the corresponding transfer coefficient was used only for estimating the time to reach steady state in the system by eq 5.6, with range 1 to 3 years, more or less consistent with the value Dr. MacCracken cites for weeds and grasses, which would be in the LB compartment. A much longer time constant, 125-525 years was found for the obdurate biosphere, which would be comprised of trees.

Line 806-808: So, you have no removal to the sediments, and so there is no ultimate sink of the C, all there will be is a redistribution among reservoirs—is that correct? If so, you will never get back to preindustrial if you have added C emissions. So, then, you might have a decay rate, but the level will never get back to preindustrial? Is that correct?

Yes, no treatment of removal of excess deep ocean DIC to sediments, and yes, never get back to preindustrial. All this takes place on a time scale much longer than examined here.

Is there any limit on how much the obdurate reservoir can build up? Basically, there is

only so much land for buildup to occur unless you have a sink to peat and eventually back to fossil fuels. So, how much can really build up (are there limits due to nutrients, etc.? Should you have a return term based on wildfires? etc.?)?

There is a return flux from the OB to the AC, corresponding to the 125-525 year time constant noted just above. In the example in Figure 7 for preindustrial stock in the LB 400 Pg, and OB, 2400 Pg (total 2800 Pg) at long time (183 yr) after cessation of emissions the net increase in the OB (uptake minus integrated land use emissions) is 100 Pg, hardly taxing the limit of the terrestrial biosphere. Wildfires are not explicitly included; the transfer coefficient characterizing the flux from the OB to the AC is based on historical observations.

And, using the piston approach, the vertical distribution of the excess CO_2 will be wrong (well, in fact, the single box does not have a vertical distribution—but, in essence, any amount taken up will instantly be creating a back flux, which is just not how the ocean works.

As in the response at line 476, the return flux from the DO to the ML is minimal even at present. This back flux is proportional to the integral of the forward flux, and thus starts from zero and increases only slowly.

Line 816: So, in setting a $S_a$eq, does this mean this is a perturbation model and is not based on gross amounts in a reservoir, such that you will return to the preindustrial value even if you through emissions add a large amount of C to the system that would, one would think, end up as distributed among the various reservoirs? I also don't understand why the difference would be with respect to the equilibrium value and the difference term in the fourth term in the equation—again, why is the difference done with respect to the concentration in the previous time step (so the emissions increase the atmospheric concentration and then this increases the mixed layer loading, etc.—what is the equilibrium value and why is it used?).

If the equations were linear it would not matter whether the quantities to be solved for in the set of ordinary differential equations (ODEs) were the total quantities or the perturbation. For the nonlinear quantities, the atmospheric stock and the ML stock, the transfer rate is expressed in departure from equilibrium, just as the oceanography community calculates the net flux on partial pressure (fugacity) difference.

The ODE solver uses the Bulirsch□Stoer method with adaptive step sizing, which uses a convergence test for solution at much higher time resolution than the one-year intervals used for input and output. Line 817: Why is not the flux simply based on the gradient between the atmospheric and the mixed layer values—what are these equilibrium values?

The equilibrium between stocks in the AC and ML is calculated with the widely used CO_2SYS program (Lewis and Wallace, 1998) with details presented in the Supplementary Information, Sect. S2, S3)

Line 823: Something is missing—"in" what?

Should read in Eqs 6.1 and 6.5. Thank you.

Line 903/line 3066: In Figure 7a, I do not understand how the amounts in the TB stay identical once the emissions stop. Over time, C will redistribute to the deep ocean, and this will lower the atmospheric concentration, and this will pull the labile C down pretty quickly and eventually the obdurate C—how is it that these stay identical over time?

Perhaps there is some confusion over the line codes. The upper two curves in Figure 7a, which reach constant values at time of cessation of anthropogenic emissions, denote both the integrated emissions (total, top curve; and fossil only, second curve) and for the top curve the sum of anthropogenic stocks in all the compartments plus integrated land use emissions; and for the second curve only sum of anthropogenic stocks. Neither the darker green curve denoting integrated net transfer from AC to OB, nor the lighter green curve, denoting net change in OB stock, taking into account LU emissions, becomes constant at the time of cessation.

Line 1135 or so: Have you also run comparisons of the models into the future with emissions continuing to see how you match or don't? This would seem interesting to see as well.

Figures 14 and 15, which are presented to compare model output with measurements or with results from other models over the measurement period, show only the time period up to cessation of emissions, 2017. It would be readily possible to run the present model out into the future for a given set of, say, RCP emissions, with the caveat that the parameters of the model are determined by prior observations.

Lines 1186-1188: So, this is not really clear from the earlier discussion, namely that the atmospheric level, for example does not decay to zero, but to 16-22% of its peak value. That would have ongoing climatic effects and so the temperature would not return to preindustrial. Basically, there is a distribution of the fossil fuel carbon among the reservoirs, which is just how IPCC represents things in its five-exponential (I think it is) equation. I'd like to see how your model results for each reservoir compare to the term in the IPCC equation for that box. I'd guess the ultimate distribution might be similar—well, except I just don't like how your ocean is represented.

The long-term non-zero floor is explicitly mentioned in the Abstract, lines 17-18. As stressed throughout this Response, that non-zero floor is of only secondary importance to this study. The constant term in the AR4 representation of the response to a pulse emission, 0.217 (Forster et al., 2007), is essentially the same as the long term-floor found in the present study. The influence of a long tail is also explicitly noted in the definition of measures of adjustment time, line 269. But perhaps I should stress that more.

Just a note here that I would like to see how your model performs versus carbon cycle models assuming that global emissions go down over several decades to zero (so say, to net-zero by 2050). Are the results similar or different?

This calculation can be readily done with this model. Stay tuned.

Lines 1347-1349***: Given this result, it would seem that the interactions and links between the mixed layer and deep ocean need to be done in a much more representative way than is done in your 5-box model. We know, for example, that such a simple model would not explain the cold deep ocean and warm surface—and that this feature is a result of how the ocean circulation is represented (there is not radiation transport, which along with convection is needed to explain the atmospheric structure). In that CO2 is carried along by the circulation (save for the biospheric pump), I just do not think it appropriate to use a formulation that just does not appropriately represent the ocean circulation. There is just too much of a chance that the result you get, being different than the more complete models, is a result of the inadequate representation of the ocean circulation to accept your results as a serious challenge to the models.

Dr. MacCracken keeps flogging the same horse. But one point to be noted here is that one cannot expect a two compartment model to represent the vertical temperature structure of the ocean. Rather it represents the departures from the preindustrial steady state in the two compartments that would ensue from a perturbation, here of heat introduced into the upper compartment.

Lines 1457-1459 and following: Might this difference be because your model leaves the amounts in the terrestrial biosphere constant in time, which just seems wrong. Might there be a term missing to reduce those over time as the atmospheric concentration declines? Going back to Figure 7, your model seems to have the TB indefinitely holding a very large fractional increase in amount as a result of the fossil fuel emissions—and the fraction seems so large it is just not clear there is enough land for this—or perhaps all trees have to grow as tall as redwoods or something. A persistent increase of the amount shown just does not seem plausible to me. Maybe what your biosphere model needs is to have the age of trees limited to some number and then those trees decay and new trees grow up with a lower CO_2 concentration; one just cannot keep having such high values (so is your algorithm based on the net mass of the trees or the fact that there is an ongoing exchange going on all the time, with average tree life perhaps, say, 40 years or so as many trees that try to grow die off as others succeed, etc.)?

I think that Dr. MacCracken has misread Figure 7a, as discussed above at Line 903/line 3066.

Line 1590: EXACTLY.

The text that Dr. MacCracken applauds reads "In view of the simplicity of the present model the question might be raised whether the exchange of CO_2 with the terrestrial biosphere and the ocean can be accurately represented by two compartments each in the TB and the ocean, as here." However the text goes on to cite recent studies with global compartment models that have represented the TB by one (Graven, 2015), two (Raupach et al., 2014), or three (Meinshausen et al., 2011) compartments and notes that Raupach et al., showed good agreement between atmospheric CO_2 over the Anthropocene as calculated with their model and as measured. The manuscript then goes on to cite the results of Naegler and Levin (2009) regarding the degree of improvement in representation of the decrease of atmospheric 14CO2 in the post-bomb era, noted above at Line 360.

Line 1594: But note that MAGICC has a much better representation of the ocean!!

In response I would again note the substantially greater net atmosphere to ocean flux obtained with the online version of MAGICC than in the present model, Figure R2 above. This would lead to a shorter adjustment time in that model than in the model developed here, except that that greater atmosphere to ocean flux is offset by the much lower rate of uptake by the TB, Figures R2 and R4 above.

Line 1601: But you are going much further out toward equilibrium, so not sure a C-14 result is really applicable. With gradients, most will surely be in the ocean, so what is so hard to figure out about that.

The reference here is to the work of Naegler and Levin (2009) that showed substantial improvement in modeled $^{14}CO_2$ in the atmosphere and the ocean (and implicitly in the TB by difference) when the number of terrestrial components was increased from one to two, but lesser improvement from two to three. The other side of the coin is that increasing the number of uncertain parameters increases the uncertainty in the model output.

Line 1614-1615: I don't recall the models in the papers cited to know how they dealt with the oceans. Very clearly, one cannot get the baseline ocean temperature distribution with the way you have done a two-box model, and it would be simply incorrect to think that the end equilibrium temperature distribution doing it with just the perturbed amount of heat to think that the temperature anomalies in the mixed layer and deep ocean would be the same. I just don't think you can justify this model for the time period you are talking about. It seems to me you just can't claim this—you really need to justify this by putting in a better representation to get results and then show that the simpler model is adequate—just saying so and citing a few papers with models being used for a different substance is just not compelling or convincing to me. It would not be hard to put in an ocean representation that is better (could likely be done in a time a lot shorter than reading your quite long paper)—so I think you should just do it. At the very least, make your deep ocean box only 650 or so meters depth so it only includes down to the thermocline.

I would reiterate my statements above regarding representation of ocean transport.

Lines 1618-1622: It is one thing to start with a complex model and work to simpler model, checking as you go that the results match. It is quite different to start with a simple model and just sort of claim it works.

The reference is to the use of a two-layer emulator of more complex models. I would concur with Dr. MacCracken provided the underlying model were shown to provide an accurate representation of the processes under examination. However as shown in Figure R3, the spread among the models does not, in my judgment, lend a lot of confidence to this approach at least at present. In fact that is much of the reason that I embarked on this study.

Lines 1627-1628: You have chosen a model form that is known not to yield a good representation of the temperature distribution or of the perturbation to temperature that the models get. Even though there are uncertainties, your ocean model is just known to not represent the effects of the ocean circulation, which is what is essential to be representing.

Lines 1657-1665: I'm sorry, but given the ocean representation you've chosen which does not adequately represent ocean circulation, this is just not convincing to me (and there seems to be some problems with the terrestrial biosphere representation as well).

The above two comments speak once more to ocean transport.

Line 1706: At this point, I'm going to pass on going over the appendices—this paper is really quite long, especially Section 8, and quite a challenge to get through.

In conclusion to this Response, I am enormously indebted to Dr. MacCracken for his careful read of the manuscript and for the probing questions he has posed. His Comment certainly stimulated my use of the online MAGICC model and the comparisons in this Response. At the end of the day I would hope that this Response has clarified his understanding of my manuscript, has allayed many of his concerns, and has identified important directions of future investigation.

References (to works not cited in original manuscript)

Forster, P., et al. 2007: Changes in Atmospheric Constituents and in Radiative Forcing. In: Climate Change 2007:The Physical Science Basis. Contribution of

Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change [Solomon, S.et al. (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA

Meinshausen, M., Raper, S.C. and Wigley, T.M., 2011. Emulating coupled atmosphere-ocean and carbon cycle models with a simpler model, MAGICC6–Part 1: Model description and calibration. Atmospheric Chemistry and Physics, 11(4), pp.1417-1456.

Meinshausen, M.et al. 2020. The shared socio-economic pathway (SSP) greenhouse gas concentrations and their extensions to 2500. Geoscientific Model Development, 13(8), pp.3571-3605.

Sutton, P.J. and Roemmich, D., 2001. Ocean temperature climate off north□east New Zealand. New Zealand Journal of Marine and Freshwater Research, 35(3), pp.553-565.