Reviewer response and marked-up manuscript

We include a response to the reviewer's comments and a marked up version of the changed manuscript below. We do not include a list of changes due to the comprehensive restructuring of the manuscript as requested by the reviewer.

Response to reviewer

We thank the reviewer for a very thorough and useful review of our manuscript, which has greatly improved the manuscript. We answer the reviewer's specific points in line below.

1. Even after reading the paper, I do not really understand the last sentence of the Abstract.

We have removed the final sentence of the abstract in order to make it more focused on the core messages and analysis in the paper.

2. I found the Introduction generally good, but maybe the need for the study could be made clearer by noting that the results from Joos et al. (2013) clearly show that: (1) the airborne fraction for a given pulse size depends upon the background state of the atmosphere; and (2) the airborne fraction for a given background state of the atmosphere depends upon the pulse size. Conversely, the standard IRF is state independent...

We have included these points more prominently in the introduction.

3. Page 2, line 16. Which "simple climate-carbon cycle models"? MAGICC? Standard IRF? Depending upon the answer, stating that these models have not been "evaluated in terms of their pulse-response behaviour" could be wrong. Similar comment for page 2, line 26.

We have been clearer on this point in order to rule out those simple models that were included in the Joos et al. 2013 paper and now give DICE as a specific example of an untested IAM carbon cycle model.

4. Page 2, line 25. In my opinion Davis and Socolow (2014) did not really evaluate the "required energy-system transitions that are needed to limit warming to below particular thresholds"; they rather estimated emissions from existing infrastructure.

We have updated our use of the reference to better reflect the content of this paper.

5. Page 3, line 15. The 4-exponentials IRF is for mathematical convenience and does not really correspond to the time dynamics of actual mechanisms, so the

“processes” provided in Table 1 are just ‘guiding analogies’. This should be noted here instead of below and made clearer.

Changed in the revised manuscript.

6. Page 3, line 18. Replace “i = 1, 4” with “i = 1, ..., 4” or “i = 1–4” or “i = 1 to 4”.

Changed in the revised manuscript.

7. Page 3, line 25. Mention that the authors decided to give a finite (1x10^6) value to \tau_0, as this differs from Myhre et al. (2013). Also mention (here or elsewhere) that for PI-IR, the a_i come from Table S2 of Joos et al. (2013).

Included in the revised manuscript. We have updated the PI-IR parameters to fit the multi-model mean of the Joos et al PI100 experiment, the parameters of which were not included in the original paper. We have stated this is the provenance of the PI-IR parameters used in the revised text.

8. Page 3, line 25-28. Defining ECS and TCR (and TCRE when mentioned much later in the paper) would probably be worthwhile.

Included in the revised manuscript.

9. Page 3, lines 25-28. The link between the c_j and both ECS and TCR is not clear. The reference to Millar et al. (2015) does not really clarify this point as the latter study presents c_j and d_j in their Supplement only, and under a mathematical form that differs from Eq. (3) considered here. To address this and other comments below, the authors need to add an Appendix or a Supplement in which they: (1) clearly show the mathematical link between the c_j and d_j as appearing in Eq. (3) and both ECS and TCR; and (2) clearly explain how they obtained c_1 = 0.46 and c_2 = 0.27.

We have provided information about the link between ECS/TCR and the model parameters with equations 4 and 5 in the revised manuscript. We do not believe this requires an appendix or supplement and hope that detail provided is sufficient to allow readers to understand how to invert the equations to solve for q_1 and q_2.

10. Page 4, Table 1. Provide the units for c_1 and c_2.

Units provided. We have changes the notation for c_1 and c_2 to q_1 and q_2 to avoid having both upper and lower case “c” variables within the paper.

11. Page 5, line 10. iIRF is not the “average airborne fraction over a period of time”, but the product is this average fraction with the length of the integration period.

Corrected in the revised manuscript.

12. Page 5, line 21 to page 6, line 2. The text is poorly structured; I think the following order would help. First, explain how FAIR uses Eqs. (4) to (6) at each time step, along with Eq. (3) for the temperature. Second, explain how were determined the specific values of r_0, r_C, and r_T (was it simply through trial and error? until finding what?). Third, state that these values work well but could be tuned even more (i.e., the text that currently appears on lines 21-27). Fourth, address the iIRF_100 > 100 years issue (would it really occur in the runs if iIRF_100 had not been limited to 95 years? if yes, the authors need to discuss the implications of this issue later on in the text).

We have corrected the order of explanation as suggested by the reviewer. We have also expanded on the selection of r_0, r_C and r_T and the effect of a maximum iIRF100 within the text.

13. Page 5, line 29: “This means the iIRF_100 is only exactly reproduced [...]”. Why?

We have added to the text to explain that this expression is only valid if \alpha is assumed to be invariant in time, a limit that would be reached for an arbitrarily small pulse.

14. Page 6, line 3. Adding one or two new methodological subsection(s) is required to clearly explain the simulations performed, how FAIR parameters were modified (uncertainty analyses), give the sources of input data (RCP, etc.), etc. Much of the text from the Results should be transferred here and expanded. Below, I refer to these new subsection(s) as “2.3”.

We agree that a methods section would aid comprehensibility and readability of the manuscript and have included one as requested.

15. Page 6, line 3. Results are discussed as they are presented, which I think is appropriate in this paper. Therefore, the section should probably be named “Results and discussion”.

Changed in the revised manuscript.

16. Page 6, lines 11-13. These two sentences are not necessary.

Removed in the revised manuscript.

17. Page 6, lines 13+. Two undiscussed elements stroke me when looking at Fig. 1. First, one would expect PI-IR to end up with lower atmospheric CO2 than historical observations (because PI-IR CO2 sinks work with pre-industrial

efficiency throughout) but this is not the case; why? I think this is because PI-IR was obtained by Joos et al. (2013) under a pulse of 100 GtC, whereas historical annual emissions were much lower and therefore initially had less impact on CO2 sinks efficiency. The authors should provide this explanation (if they agree with it) as it addresses the issue and strengthens their point about the inadequacy of state-independent IRF model. Second, PI-IR CO2 sinks are less efficient than FAIR CO2 sinks until about year 2000 (Fig. 1c). This seems mathematically impossible when looking at parameter values in Table 1 and the different equations... unless \alpha in FAIR has a value < 1. I think the authors should explain this here, and also give in the Methods the initial value of \alpha (about 0.16, right?) when C_acc and T are still equal to zero (i.e., when iIRF_100 is equal to r_0). The way FAIR is introduced, I initially thought \alpha would always be > 1 and got confused.

We agree with the reviewer regarding their explanation of the PI-IR curve in Figure 1a (along with the inability of the temperature independent AR5-IR and PI-IR models to capture temporary reductions in iIRF100 due to naturally forced cooling) and have explicitly stated this in the revised manuscript. We have included the value of \alpha in the pre-industrial state for reference in the revised manuscript.

18. Page 6, lines 16-17. “The AR5-IR displays a too large [...]”. This sentence is a poor description of Fig. 1c, as no single model is really “consistent with the observations”. The authors can only state that AR5-IR is always higher than FAIR and that both are much more stable than observations.

We have removed these phrases and comparisons from the revised manuscript.

19. Page 7, line 6. Mentioning the social cost carbon one time in the Introduction is OK, but coming back to this concept throughout the paper seems out of place (has the paper been written for another journal?) and pointless (a model of CO2 dynamics needs to give good results to be useful for any application, not just the social cost of carbon); please remove. Similar comment for page 2, line 13; page 6, lines 15-17; and page 13, line 13.

Removed in the revised manuscript.

20. Section 3.1. The text should refer to the results in Fig. 1d, Fig. 2c, and Fig. 2d or these panels should be removed. With a current total of 28 Figure panels, less would probably be better.

All panels shown are now discussed in the text. We have also streamlined the number of figures in the text to help communicate the main points of the paper.

21. Page 7, lines 15-19. This is Introduction-type text, not for the Results.

Now included in the introduction.

22. Section 3.2. The authors often mention iIRF_100, but this variable is not shown in the Figures. The authors should present airborne fraction results instead or add a Table with iIRF_100 results.

A table of iIRF100 results is now included.

23. Page 7, lines 20-26. The majority of this text belongs to 2.3, along with the explanations about how “fully-coupled”, “biogeochemically-coupled”, and “radiatively-coupled” results were obtained. I also suggest removing Fig. 3a, which is more ‘methodological’ (i.e., diagnosed emissions required to reach a particular CO2 level) and not really interesting in itself.

This material is moved to the new methods and the figure panel removed.

24. Page 7, lines 31-32. All models show a rapid temperature increase followed by a relatively stable value, not just the “fully-coupled” FAIR model.

Changed in the revised manuscript.

25. Page 8, lines 1-3. Cumbersome sentence; please rephrase.

Phrasing changed in the revised version of the manuscript.

26. Page 8, lines 18-26. The majority of this text belongs to 2.3, where the decision to maintain the same ratio between r_T and r_C needs to be justified. But in fact, I even suggest removing Fig. 4 from the paper as I do not believe it adds much value.

Text moved to methods section and the choice of fixed r_T : r_C ratio (as the fully-coupled PD100/PI100 experiments do not constrain the balance between r_T and r_C) is discussed. We choose to keep figure 4 as we believe that it shows that FAIR is capable to simulating the carbon-cycle responses in the full range of ESM and EMIC models from Joos et al in both the PD100 and PI100 experiments using just a single set of parameters for each model, an important demonstration of versatility for policy relevant use where spanning ranges of ESM responses is important. This analysis was also explicitly asked for by a reviewer in the previous round.

27. Page 8, line 27. Actually, the FAIR model is able to “successfully capture much of the response” only for the well-behaving models (i.e., no major year-to-year variability); this should be noted.

We note this caveat about not being able to simulated complex model inter-annual variability.

28. Page 8, line 31 to page 9, line 8. The majority of this text belongs to 2.3, with possibly some elements to the Introduction.

Text moved in the revised version of the manuscript.

29. Page 8, lines 5-12. The authors apparently misunderstood Zickfeld and Herrington (2015). The issue with the results of Ricke and Caldeira (2014) is not so much that they did not "account[] for feedbacks on the carbon cycle and fail[ed] to capture the plateau of CO2-induced warming" as that they did not account for the effect of the pulse size on the shape of the temperature response (because they used a state-independent IRF model): for very large pulses, there is no longer an early warming peak followed by a plateau. FAIR is able to capture this change of shape from small pulses leading to an early warming peak (Fig. 8d) to large pulses without an early warming peak (Fig. 5a), whereas the standard IRF model is not (Fig. 5a). Although this outcome further illustrates the scientific value of FAIR, this change of shape occurs for pulses > 1000 GtC (also see Zickfeld and Herrington, 2015) and is therefore of little practical relevance for real emission scenarios. Given the amount of results presented, I thus suggest removing Fig. 5 from the paper as I do not believe it adds much value. If the Figure is kept, the text describing its results should be made accurate.

We have revised our discussion of Zickfeld and Herrington to focus the discussion on the dependence of maximal warming on the pulse size and have attempted to clarify and tighten the discussion around figure 5.

30. Page 9, lines 14-31. The majority of this text belongs to 2.3, where the 'stopping rule' for the different %/yr simulations should be given (until quadrupling initial CO2?). The long sentence on lines 23-26 adds little value.

Text moved to the methods section and stopping rule specified. We have moved the sentence from page 9 line 23-26 to the conclusions. We choose to retain this sentence as we do believe the integrated analysis of mitigation and solar radiation management requires a model that correctly distinguishes the effects of warming and carbon on the carbon-cycle and that this is a worthwhile point to make. However this point perhaps fits better in the conclusion section.

31. Page 10, line 3. Why does a "constant airborne fraction necessarily give[] an approximately quadratic increase" and what is "approximately quadratic" (e.g., an exponent of 1.8)?

This statement, which the reviewer correctly identified as confusing has been removed from the revised manuscript.

32. Page 10, lines 6-8. Cumbersome sentence; please rephrase.

Rephrased in the revised manuscript.

33. Page 10, lines 8-10. Specify that these results are for “radiatively-coupled” experiments.

Changed in revised manuscript.

34. Page 10, lines 10-14. Cumbersome sentence; please rephrase.

Rephrased in revised manuscript.

35. Page 10, line 15. The “cumulative airborne fraction” could be more clearly defined: it is the fraction of all past emissions that are still in the atmosphere.

Changed in revised manuscript.

36. Page 10, lines 29-34. This is Introduction-type text, not for the Results.

Moved in revised manuscript.

37. Page 11, lines 1-5. Much more details about what was done and how it was done (i.e., which was the range of c_j values used) must be provided in 2.3 and possibly the new Appendix/Supplement. Figs. 7a and 7b should also be better introduced.

We now provide the range of q_j used to sample the TCR and ECS range. As mentioned above, we also provide more detail in the model description about how to link TCR/ECS and q_j. We try to introduce figure 7a and 7b better in the revised manuscript.

38. Page 11, lines 8-13. This sentence is confusing: it seems to imply that the increasing airborne fraction was due to changes in FAIR parameters (r_0, r_T, r_C, and c_j) *through time*. I rather assume that the increasing airborne fraction results from the structure of FAIR that leads to an ever-increasing \alpha, with changes FAIR parameters being responsible for the blue shading.

We have tried to make it clearer that the r_0, r_T and r_C are indeed constant over time and that the uncertainty corresponds to different values of these parameters in the revised manuscript.

39. Page 11, lines 14-23. FAIR results in Fig. 7d do not show a “straight-line relationship between cumulative carbon emissions and human-induced warming”, the downward curvature being obvious starting from 500 GtC at least (not only at high cumulative emissions as mentioned). Please specify that the TCRE value provided is valid for 1000 GtC only.

We have revised the text to indicate the downward curvature is apparent across the range of cumulative emissions shown. We include a statement indicating that the value of TCRE is only exact for the first 1000GtC. Whilst recognising that there is some curvature to the relationship we choose to refer to it as 'approximately linear' to reflect the terminology widely used to refer to this simulated relationship in the literature.

40. Page 11, lines 24-26. This is Introduction-type text, not for the Results.

This text has been moved in the revised manuscript.

41. Page 11, line 28 to page 12, line 10. This text is not really related to FAIR and could easily be deleted. If not, it should be moved to a new subsection 3.5.

We believe that this section makes some important points that underlie our ranges for sampling TCR and RWF in the paper. We therefore keep this text as part of the new methods section.

42. Page 11, line 30. Gillett et al. (2013) highlighted the policy relevance of TCRE, not TCR.

We have altered this text (now in the methods section) to reflect that Gillett et al showed that thermal response uncertainty (therefore TCR) is the dominant uncertainty in the CMIP5 simulated range of TCRE.

43. Page 12, lines 10-16. The majority of this text belongs to 2.3, where explanations must be much improved. In particular, the c_j (not TCR and ECS) were changed in FAIR; how?

This text has been moved to the methods. See previous responses for our changes regarding how TCR and ECS are related to q_j.

44. Page 12, lines 16+. The text should refer to the results in Fig. 8b or this panel should be removed.

Panel 8b has now been removed from the revised manuscript.

45. Page 12, lines 17-23. The majority of this text belongs to 2.3, where the IPT equation and the distribution (shape and values) of d_1 must be justified. I also suggest combining what remains of this paragraph with the next one, as they are logically linked (i.e., there are no results for d_1 uncertainty only, right?).

We now include this text in the new methods section. We remove the IPT equation as our update the thermal response timescales (to reflect the mean to Geoffrey et al as opposed to the values used in IPCC-AR5) mean that is it no longer valid. Paragraphs are now combined in the revised manuscript.

46. Page 12, lines 24-27. The majority of this text belongs to 2.3, where the values chosen must be justified, the 300 random draws approach must be explained, and the “median” shown in Fig. 8 must be defined (is it the median of the 300 draws or the result obtained with the median values of the parameter distributions?).

We have moved this text and tried to justify our approach more in the revised manuscript.

47. Before the Conclusions. A short discussion of the following point is missing: the idea behind FAIR is to adjust the standard IRF time constant based on iIRF_100, which is by definition for a time horizon of 100 years. Do you expect FAIR to perform well for time horizons of 1000 years and more? Would this require expanding the time horizon of iIRF?

We include this discussion in the revised version of the manuscript as part of the conclusions, which felt a more natural fit for these points.

48. Page 13, line 10. Strictly speaking, the authors did not show that “including both explicit CO2 uptake- and temperature- induced feedbacks are essential”. One could hypothesize that FAIR may give similar results working with temperature-induced feedback only but using a higher value of r_T (or with CO2 feedback only but using a higher value of r_C).

We would disagree with the reviewer on this point. We believe that we have shown the including temperature-induced carbon-cycle feedbacks are essential in order to replicate the ‘radiatively-coupled’ experiments under the prescribed concentration increase experiments. As mentioned in the text, the correct partitioning of the feedbacks will be important in scenarios with future differences in the relative contribution of CO2 and non-CO2 forcings on the climate-carbon-cycle system.

49. Page 15. The same study is referenced under both Meinshausen et al. 2011a and 2011b.

Corrected in the revised manuscript.

50. The Figures along with the related legends and captions are very poor. Here is only a subset of all the issues I saw.

- Fig. 1 caption: these results are for the historical period, so why the “RCP scenarios”?

We have updated the figure caption to give a more-complete statement on where

the non-CO2 forcings come from.

- Many panels lack the x-axis title.

All panels now have an x-axis title in the revised manuscript.

- There is often an overlap between axis scales (e.g., Fig. 1d bottom left) or with the panel letter (e.g., Fig. 5a bottom left).

Corrected in revised manuscript.

- Axes are not labelled consistently, for example “CO2 concentration” vs. “CO2 Concentrations” vs. “Atmospheric concentration”, or “Airborne Fraction” vs. “Airborne fraction” vs. “Fraction of CO2 impulse remaining”.

Corrected in revised manuscript.

- Results are often ‘cut’ because the scale is not large enough (e.g., Fig. 2c).

Corrected in revised manuscript.

- Why aren’t MAGICC results shown on Figs. 2c and 2d; I presume they are available?

We have removed figure 2c and 2d from the revised manuscript.

- Fig. 2c shows results with very different scales, while Fig. 2d shows too many results. Fig. 2 should have 6 panels: a) as currently, b) as currently, c) temperature for RCP8.5, d) temperature for RCP2.6, e) temperature vs. emissions for RCP8.5, and f) temperature vs. emissions for RCP2.6.

We have removed figure 2c and 2d from the revised manuscript.

- Fig. 3 caption: the shading is grey, not black.

Corrected in revised manuscript.

- The legends of Fig. 6a and 6c should both appear in 6a as they apply to all panels of Fig. 6. The legend of Fig. 6b should be ordered logically.

Corrected in revised manuscript.

- Fig. 6 caption: replace “cumulative total carbon uptake” by “cumulative ocean and land carbon uptake”.

Corrected in revised manuscript.

- There is no purple bar in Fig. 7a, in contradiction with both the legend and caption. There is no purple bar in Fig. 7d, in contradiction with the caption.

A purple bar was included in Figure 7a. We have altered the colour of the bar to aid visibility.

- Why aren't AR5-IR results shown in Figs. 7a and 7b?

AR5-IR results are not included in Figure 7a and 7b as these models share a thermal response model and these simulations are conducted with prescribed concentrations and not emissions, so would produce identical resutls. We have included a statement in the manuscript to make this point clear.

References

Davis and Socolow (2014). Environmental Research Letters 9, 084018.

Gillett et al. (2013). Journal of Climate 26, 6844-6858.

Joos et al. (2013). Atmospheric Chemistry and Physics 13, 2793-2825.

Millar et al. (2015). Climatic Change 131, 199-211.

Myhre et al. (2013). IPCC AR5, WG1, Chapter 8.

Ricke and Caldeira (2014). Environmental Research Letters 9, 124002.

Zickfeld and Herrington (2015). Environmental Research Letters 10, 031001.

A modified impulse-response representation of the global response to carbon dioxide emissions

Richard J. Millar[1,2,3], Zebedee. R. Nicholls[1], Pierre Friedlingstein[3], and Myles R. Allen[1,2,4]

[1]Department of Physics, University of Oxford, Oxford, UK
[2]Oxford Martin Net Zero Carbon Investment Initiative, Oxford Martin School, University of Oxford, Oxford, UK
[3]Department of Mathematics, University of Exeter, Exeter, UK
[4]Environmental Change Institute, University of Oxford, Oxford, UK

Correspondence to: Richard J. Millar (richard.millar@physics.ox.ac.uk)

Abstract. Projections of the response to anthropogenic emission scenarios, evaluation of some greenhouse gas metrics, and estimates of the social cost of carbon ~~,~~ often require a simple model that links emissions of carbon dioxide (CO_2) to atmospheric concentrations and global temperature changes. An essential requirement of such a model is to reproduce the behaviour of more complex Earth System Models as well as an ability to sample their range of response in a transparent, accessible and reproducible form. Here we adapt the simple model of the Intergovernmental Panel on Climate Change 5th Assessment Report (IPCC-AR5) to explicitly represent the state-dependence of the CO_2 airborne fraction ~~and reproduce~~ . Our adapted model (FAIR) reproduces the range of behaviour shown in full and intermediate complexity Earth System Models under several idealised carbon-cycle experiments. We find that a simple linear increase in 100-year integrated airborne fraction with cumulative carbon uptake and global temperature change is both necessary and sufficient to represent the response of the climate system to CO_2 on a range of timescales and under a range of experimental designs. ~~Quantified ranges of uncertainty (analogous to current assessed ranges in Equilibrium Climate Sensitivity and Transient Climate Response) in integrated airborne fraction over the 21st century under a representative mitigation scenario, and an assessed range in how much this quantity may have changed relative to pre-industrial conditions, would be valuable in future scientific assessments.~~

1 Introduction

~~Future emissions of~~ In the long term, future climate changes will largely be determined by future cumulative CO_2 ~~over the remainder of the century~~ emissions (Matthews et al., 2009; Allen et al., 2009; Meinshausen et al., 2009), but the timing and magnitude of emissions are uncertain and a strong function of future climate policy (Van Vuuren et al., 2011). ~~Future climate changes, and their associated impacts, will largely be determined by future cumulative carbon dioxide emissions (Matthews et al., 2009; Allen et al., 2009; Meinshausen et al., 2009), but linking~~ Linking specific CO_2 emission scenarios to future transient climate change requires a model of the interacting climate-carbon-cycle system. Comprehensive Earth System Models (ESMs) explicitly simulate the physical processes that govern the coupled evolution of atmospheric ~~carbon~~ CO_2 concentrations and the associated climate response (Friedlingstein et al., 2006). However, such models are typically highly computationally intensive and can therefore only be run for a few representative future emission scenarios (Taylor et al., 2012).

For analysis of arbitrary emissions scenarios, as required for the integrated assessment of climate policy~~and calculation of the social cost of carbon~~, a computationally efficient representation of the Earth system is required (Marten, 2011).

Simplified representations of the coupled climate-carbon-cycle system take many forms (Hof et al., 2012). A key test for simplified ESMs is whether they correctly capture the physics of the co-evolution of atmospheric CO_2 concentrations and global mean temperature under both idealised settings and under possible projections of future emissions scenarios. Following a CO_2 pulse emission of 100GtC in present-day climate conditions, ESMs (and Earth System Models of Intermediate Complexity – EMICs) display a rapid-draw down of CO_2 with the concentration anomaly reduced by approximately 40% from peak after 20 years and by 60% after 100 years, followed by a much slower decay of concentrations leaving approximately 25% of peak concentration anomaly remaining after 1000 years (Joos et al., 2013). The speed and shape the this pulse decay is dependent on both the background climate state and the size of the pulse, but a substantial fraction of the emission is simulated to remain in the atmosphere after 1000 years in all cases. The effect of this longevity of fossil carbon in the atmosphere, combined with the gradual "recalcitrant" thermal adjustment of the climate system (Held et al., 2010), is to induce a global mean surface temperature ~~(GMST)~~ response to a pulse emission of CO_2 characterised by a rapid warming, over approximately a decade, to a plateau value of ~~GMST~~ global mean surface temperature anomaly (Joos et al., 2013). Warming does not noticeably decrease from this value over the following several hundred years, indicating that, short of artificial CO_2 removal (CDR) or active solar geoengineering, CO_2-induced warming is essentially permanent on human-relevant timescales.

~~As computations of the social cost of carbon require the discounted summation of future climate change-induced economic damages associated with an additional pulse emission of CO_2 above a baseline scenario, the~~ The correct representation of the temporal evolution of the warming response to ~~the~~ a pulse emission is required ~~from~~ for computationally-simple climate-carbon-cycle models. ~~As~~ Aside from the simple climate-carbon-cycle models ~~are not explicitly~~ analysed in Joos et al. (2013), many simple models, including some used in integrated assessment models (IAMs - e.g see Nordhaus (2010)), have not explicitly been evaluated in terms of their pulse-response behaviour, ~~it is~~ and it remains unclear how well ~~this robustly simulated physics is~~ the physical dependences of the pulse-response are represented in such models. The social cost of carbon is conventionally calculated by applying a pulse emission of a specified magnitude in near to present-day conditions as a perturbation on top of a specified future emissions scenario (NAS, 2016). As calculating the social cost of carbon is a key element of many cost-benefit analyses of climate change policy, simple climate-carbon-cycle models used in IAMs should aim to reproduce the pulse-response dependencies on pulse size and background state that have been highlighted in ESMs and EMICs (Joos et al., 2013; Herrington and Zickfeld, 2014).

A second important feature of ~~more complex climate-carbon-cycle models~~ ESMs is the increase in airborne fraction (the ~~percentage~~ fraction of emitted CO_2 that remains in the atmosphere after a ~~period of time) over~~ specified period) with time in scenarios involving substantial ~~levels of~~ emissions or warming ~~(Friedlingstein et al., 2006; Millar et al., 2016)~~(Friedlingstein et al., 2006). An emergent feature of the CMIP5 full-complexity ESMs appears to be that this increase in airborne fraction approximately cancels the logarithmic relationship between CO_2 concentrations and radiative forcing, yielding an approximately linear relationship between cumulative CO_2 emissions and CO_2-induced warming (Matthews et al., 2009; Gillett et al., 2013). This relationship has given rise to the concept of an all-time cumulative 'carbon budget' to restrict warming to a certain level (Ro-

gelj et al., 2016), which has quickly become an important tool in evaluating the required energy-system transitions that are needed to limit warming to below particular thresholds (Gignac and Matthews, 2015; Van Vuuren et al., 2016), as well as the climate implications of the existing capital stock (Davis and Socolow, 2014; Pfeiffer et al., 2016). As simple climate-carbon-cycle models are often used to compute particular carbon budgets in integrated assessment scenarios (e.g. the MAGICC model as used in Meinshausen et al. (2009)), the ability to reproduce the approximate linearity of the relationship between warming and cumulative emissions is a desirable property.

Representing climate response uncertainty is also a crucial factor in the integrated assessment of climate policies. Despite significant advances in climate system understanding, non-negligible uncertainties remain in the response of the coupled climate-carbon-cycle system to emissions of CO_2 (Gillett et al., 2013) implying that climate policies have to be constructed and assessed in the light of this continued uncertainty (Millar et al., 2016). Integrated assessment activities require a representation of the physical climate system that can transparently and simply sample physically-consistent modes of climate response uncertainty, partly in order to capture and assess the possibility of extreme and highly costly responses within the Earth system (often called "fat-tailed" outcomes) (Weitzman, 2011).

In this paper we show that although the impulse-response functions ~~that are~~ provided for the calculation of multi-gas equivalence metrics in IPCC-AR5 (Myhre et al., 2013) ~~,~~ provide a simple and easy to use ~~coupled~~ climate-carbon-cycle model, ~~are~~ this model is insufficient to fully capture ~~these~~ the emergent responses of the coupled climate-carbon-cycle system. Such a ~~state-insensitive~~ state-independent impulse-response model cannot simultaneously reproduce the relationship between emissions, concentrations and temperatures seen over the historical period and the projected response over the 21^{st} century to both ~~high-emission~~ high-emissions and mitigation scenarios as simulated by ESMs ~~.~~ and EMICs. Indeed, such a model formalism would inherently fail to capture the dependence of the evolution of the airborne fraction following a pulse emission on both the background state of the climate and pulse size, as simulated by ESMs in Joos et al. (2013). We therefore propose a simple extension of the standard IPCC-AR5 impulse-response model, coupling the carbon-cycle to the thermal response and to cumulative carbon uptake by terrestrial and marine sinks in order to reproduce the behaviour of the ESMs under a variety of idealised experiments and future emissions scenarios.

Section 2 describes the formalism of the models that we contrast throughout this paper ~~. We then describe,~~ and describes the methodological details of the experiments that we use to analyse the responses of these models. We show and discuss the results of these model validation experiments in section 3, beginning, in section 3.1, with why a state-dependence modification to the IPCC-AR5 carbon-cycle impulse-response function is required, motivating the modified ~~model~~ 'Finite Amplitude Impulse-Response Model' (FAIR) described in section 2. Section 3.2 then evaluates ~~these models' ability in replicating the~~ the ability of FAIR and the unmodified IPCC-AR5 impulse-response models to replicate the dependencies of the response to a pulse-emission on background conditions and pulse size shown in ESMs and EMICs. Section 3.3 evaluates the models' behaviour under a set of idealised experiments in which CO_2 concentrations are increased by a fixed percentage each year starting from pre-industrial values. Section 3.4 discusses uncertainty in ~~the modified simple model~~ FAIR and how probabilistic assessments of climate response to CO_2 emissions could be made using the model. Section 4 provides a concluding summary and discussion.

2 Model ~~descriptions~~description and methods

2.1 The IPCC AR5 Impulse-Response (AR5-IR) model

The IPCC-AR5 proposed an idealised simple ~~climate~~ climate-carbon-cycle model for metric calculations, incorporating a "2-box" or "2-time-constant" model of the temperature response to radiative forcing with a "4-time-constant" impulse-response model of the CO_2 concentration response to emissions (Myhre et al., 2013). This model represents the evolution of atmospheric CO_2 by partitioning emissions of anthropogenic CO_2 between four different reservoirs (all of which are empty in pre-industrial equilibrium) of atmospheric carbon anomaly that each decay with a fixed time constant. ~~The impulse-response function for a unit emission at time $t = 0$ is therefore give as,~~ Four carbon pools are determined to be sufficient to empirically represent the response of atmospheric CO_2 concentration anomalies following a pulse emission of 100GtC, above a specified background concentration of 389ppm, over the 1000 years following the pulse (Joos et al., 2013). These carbon pools do not directly correspond to individual physical processes and instead represent the combined effect of several carbon-cycle mechanisms, however, processes that are guiding analogues to the timescale of the pool decays are summarised in table 1. The evolution of the carbon concentration anomaly in each pool, R_i, is given as,

$$\frac{\mathrm{d}R_i}{\mathrm{d}t} = a_i E - \frac{R_i}{\tau_i} \quad ; \quad i = 1,41-4 \tag{1}$$

where E ~~are~~ is the annual CO_2 emissions, in units of ppm/year (1 ppm = 2.12GtC), a_i is the fraction of carbon emissions entering each reservoir and τ_i the decay time constant for that pool. Atmospheric CO_2 concentrations are given by $C = C_0 + \sum_i R_i$, and radiative forcing by:

$$F = \frac{F_{2\mathrm{X}}}{\ln(2)} \ln\left(\frac{C}{C_0}\right) + F_{\mathrm{ext}} \quad , \tag{2}$$

where C_0 is the pre-industrial CO_2 concentration, $F_{2\mathrm{X}}$ the forcing due to CO_2 doubling ($F_{2\mathrm{X}}$=3.74Wm^{-2}), and F_{ext} the non-CO_2 forcing. ~~GMST anomalies~~ Global mean surface temperature anomalies (T) are computed thus:

$$\frac{\mathrm{d}T_j}{\mathrm{d}t} = \frac{c_j F - T_j}{d_j} \frac{q_j F - T_j}{d_j} \quad ; \quad T = \sum_j T_j \quad ; \quad j = 1,22. \tag{3}$$

~~with coefficients~~ Coefficients a_i, ~~d_j~~ and τ_i are as given in AR5 Chapter 8, tables 8.SM.9 and 8.SM.10 (Myhre et al., 2013)~~. c_j are set to give an~~ , except for τ_0 which is here given a finite value and not set to infinity (all results presented in this paper are insensitive to this choice). The two thermal response timescales, d_j, have been ordered to run from longest to slowest, as with the carbon-cycle response timescales, and are chosen to match the multi-model means of Geoffroy et al. (2013). By considering the analytic solutions of equation 3 under an instantaneous doubling of CO_2 concentrations and a 1%/yr increase in CO_2 concentrations, q_j can be related to the Equilibrium Climate Sensitivity (ECS~~) =2.75K~~ [1]) and Transient Climate Response

[1] The global mean warming resulting from an instantaneous doubling of pre-industrial CO_2 concentrations after allowing the climate system to reach a new equilibrium state.

(TCR) [2]) via the expressions,

$$\mathrm{ECS} = F_{2\times}(q_1 + q_2), \tag{4}$$

and

$$\mathrm{TCR} = F_{2\times}\left(q_1\left(1 - \frac{d_1}{70}\left(1 - \exp\left[-\frac{70}{d_1}\right]\right)\right) + q_2\left(1 - \frac{d_2}{70}\left(1 - \exp\left[-\frac{70}{d_2}\right]\right)\right)\right), \tag{5}$$

Equations 4 and 5 can be inverted to give expressions for q_j in terms of ECS and TCR assuming response timescales (d_j) as given in table 1 (Millar et al., 2015). We choose default values for q_j corresponding to TCR=1.6K ~~(corresponding to $c_1 = 0.46$ and $c_2 = 0.27$ (Millar et al., 2015))~~and ECS=2.75K ($q_1 = 0.33\mathrm{KW}^{-1}\mathrm{m}^2$ and $q_2 = 0.41\mathrm{KW}^{-1}\mathrm{m}^2$), indicative of a typical mid-range climate response to radiative forcing in ESMs (Flato et al., 2013). ~~The four carbon pools, each with a fixed decay time constant, are determined to be sufficient to empirically represent the response of atmospheric CO_2 concentration anomalies following a pulse emission of 100GtC, above a specified background concentration of 389ppm, over the 1000 years following the pulse (Joos et al., 2013). As the fraction of carbon emissions entering each reservoir (a_i) and the decay time constant (τ_i) are determined empirically, they do not in themselves correspond to individual physical processes and instead represent the combined effect of several carbon-cycle mechanisms. However, the distinct range of decay timescales indicates specific physical processes that are strongly associated with the evolution of each carbon reservoir. These are summarised in table 1.~~

We use two versions of the AR5-IR model in this paper, one calibrated to the present-day (AR5-IR) and one calibrated to the pre-industrial (PI-IR) climate response to a pulse emission ~~(PI-IR)~~ respectively. The AR5-IR model is used for the calculation of absolute Global Temperature Potentials (aGTPs) in IPCC-AR5 and has carbon-cycle coefficients that best represent the ESM simulated evolution of a 100GtC pulse emission under approximately present-day conditions. The PI-IR model uses an alternative set of coefficients that are selected to represent the evolution of a 100GtC pulse emission in pre-industrial conditions~~for~~. These parameters are derived from a fit to the multi-model mean of the ensemble of ESMs and EMICs ~~in~~ from Joos et al. (2013) (see table 1 for parameter values).

2.2 A ~~"~~Finite Amplitude Impulse Response ~~"~~ (FAIR) model

In the AR5-IR ~~model~~ and PI-IR models the carbon-cycle ~~constants are not~~ response to a pulse emission is not explicitly affected by rising temperature or CO_2 accumulation and hence only represent the specific response to a particular perturbation scenario. In more comprehensive models, ocean uptake efficiency declines with accumulated CO_2 in ocean sinks (Revelle and Suess, 1957) and uptake of carbon into both terrestrial and marine sinks are reduced by warming (Friedlingstein et al., 2006).

In an attempt to capture some of these dynamics within the simple impulse-response model structure, we ~~here~~ attempt a minimal modification of the AR5-IR model to allow it to mimic the behaviour of ~~more complex models~~ ESMs/EMICs in response to finite-amplitude CO_2 injections, which we call a Finite Amplitude Impulse-Response (FAIR) model. To introduce a state-dependent carbon uptake as simply as possible, we apply a single scaling factor, α, to all four of the time-constants in

[2] The global mean warming at the time of doubled CO_2 concentrations following a 1%/yr increase from pre-industrial values

Parameter	Value - AR5-IR	Value - PI-IR	Value - FAIR	~~Processes~~ Guiding analogues
a_0	0.2173	~~0.1266~~ 0.1545	0.2173	Geological re-absorption
a_1	0.2240	~~0.2607~~ 0.1924	0.2240	Deep ocean invasion / equilibration
a_2	0.2824	~~0.2909~~ 0.2424	0.2824	Biospheric uptake / ocean thermocline invasion
a_3	0.2763	~~0.3218~~ 0.4108	0.2763	Rapid biospheric uptake / ocean mixed-layer invasion
τ_0 (yr)	1x10^6	1x10^6	1x10^6	Geological re-absorption
τ_1 (yr)	394.4	~~302.8~~ 276.7	394.4	Deep ocean invasion/equilibration
τ_2 (yr)	36.54	~~31.61~~ 30.75	36.54	Biospheric uptake / ocean thermocline invasion
τ_3 (yr)	4.304	~~4.240~~ 4.459	4.304	Rapid biospheric uptake / ocean mixed-layer invasion
~~e_1~~ g_1 (KW^{-1}m^2)	~~0.46~~ 0.33	~~0.46~~ 0.33	~~0.46~~ 0.33	Thermal ~~adjustment of upper~~ equilibration of deep ocean
~~e_2~~ g_2 (KW^{-1}m^2)	~~0.27~~ 0.41	~~0.27~~ 0.41	~~0.27~~ 0.41	Thermal ~~equilibration of deep~~ adjustment of upper ocean
d_1 (yr)	~~8.4~~ 239.0	~~8.4~~ 239.0	~~8.4~~ 239.0	Thermal ~~adjustment of upper~~ equilibration of deep ocean
d_2 (yr)	~~409.5~~ 4.1	~~409.5~~ 4.1	~~409.5~~ 4.1	Thermal ~~equilibration of deep~~ adjustment of upper ocean
r_0 (yr)	-	-	~~35~~ 32.40	Pre-industrial iIRF$_{100}$
r_C (yr/GtC)	-	-	~~0.02~~ 0.019	Increase in iIRF$_{100}$ with cumulative carbon uptake
r_T (yr/K)	-	-	~~4.5~~ 4.165	Increase in iIRF$_{100}$ with warming

Table 1. Default parameter values for the simple impulse-response climate-carbon-cycle models used in this paper. ~~Note that, for consistency with (Myhre et al., 2013), the ordering of indices is fast-slow for the thermal response and slow-fast for the carbon cycle.~~

the carbon-cycle of the AR5-IR model, such that the CO_2 concentrations in the 4 "carbon reservoirs" are updated thus:

$$\frac{\mathrm{d}R_i}{\mathrm{d}t} = a_i E - \frac{R_i}{\alpha \tau_i} \quad ; \quad i = 1,4 1-4 \tag{6}$$

To identify a suitable state-dependence, we focus on parameterising variations in the 100-year integrated impulse response function, iIRF$_{100}$. A focus on the integrated impulse response (average airborne fraction over a period of time, multiplied by the length of time period), as opposed to the airborne fraction at a particular point in time, ~~it~~ is more closely related to the impact of emissions on the global energy budget, and also to other metrics such as Global Warming Potential (GWP)~~(Joos et al., 2013)~~. With other coefficients fixed, iIRF$_{100}$ is a monotonic (but non-linear) function of α:

$$\mathrm{iIRF}_{100} = \sum_i \alpha a_i \tau_i \left[1 - \exp\left(\frac{-100}{\alpha \tau_i}\right)\right]. \tag{7}$$

~~Following other simplified carbon-cycle models (Meinshausen et al., 2011a; Glotter et al., 2014), we~~ Equation 7 is derived using the approximation that α is independent of time, and hence equation 7 is only exactly equivalent the iIRF$_{100}$ for infinitesimal perturbations against constant background conditions (in which the approximation of time-independent α becomes exact). We assume iIRF$_{100}$ is a function of accumulated perturbation carbon stock in the land and ocean (equivalent to the amount of emitted carbon that no longer resides in the atmosphere), $C_{\mathrm{acc}} = \sum_t E - (C - C_0)$, and of ~~GMST~~ global mean temperature anomaly from pre-industrial conditions, T. A simple linear relationship appears to give an adequate approxima-

tion to the behaviour of ESMs and EMICs (as will be shown subsequently in section 3):

$$\text{iIRF}_{100} = r_0 + r_C C_{\text{acc}} + r_T T. \quad (8)$$

At each time-step we first compute the required iIRF_{100} using C_{acc} and T from the previous time-step (equation 8). We then numerically solve equation 7 for the compatible value for α, which is then in turn used to update the carbon pool concentrations (equation 6). The total radiative forcing is then computed with equation 2, before changes in global mean temperature are computed with equation 3.

Values of r_0=~~35~~ 32.4 years, r_C=~~0.02~~ 0.019 years/GtC, ~~recalling that 2.12 GtC = 1ppm,~~ and r_T=~~4.5~~ 4.1 years/K, with ECS=2.75K and TCR=1.6K, ~~give a numerically-computed iIRF_{100} of 53 years for a 100 GtC pulse released against a background CO_2 concentration of 389ppm following a historical build-up, consistent with the central estimate of Joos et al. (2013). These parameters also~~ are here used as model default parameters[3]. We choose these parameters to approximately replicate the relationship between warming-driven outgassing of carbon in the bulk of CMIP5 ESMs (see section 3.3), whilst also diagnosing near-observed values of present-day CO_2 emissions to achieve present-day concentrations. The values ~~of r_0, r_T and r_C given here~~ given here as default parameters are intended to be taken only as approximate ~~best-estimate~~ CMIP5-representative values that capture important carbon-cycle dynamics in ESMs. ~~The exact values of these parameters could~~ These values have not been explicitly optimised to any particular goal and can be tuned (along with the other ~~parameters in the model) to best-reproduce the aspect of~~ model parameters) to reproduce specific aspects of individual ESM/EMIC behaviour ~~of interest~~ (e.g. see ~~Figure~~ figure 4). Best-estimate values for the FAIR parameters will depend on exactly what feature of ESM behaviour is the desired target for the optimisation.

~~We compute iIRF_{100} at each time-step using C_{acc} and T from the previous time-step and equation 8, convert to a α using equation 7 and apply to the carbon-cycle equations (equation 6). This means the iIRF_{100} is only exactly reproduced under constant background conditions with infinitesimal perturbations.~~ Values of iIRF_{100} larger than 100 years correspond to a net carbon source to the atmosphere in response to a perturbation, and, as perturbations to the carbon stock in the atmosphere would grow indefinitely, makes the model unstable. In this regime there is no solution for α, so we set iIRF_{100} to a maximum value of ~~95~~ 96.6 years, corresponding~~with these parameters~~ , with the parameters as given in table 1, to α=~~65.4.~~ 100. This physically corresponds to a near-absence of carbon sinks in the Earth system following a very large injection, with very slow rates of decay of atmospheric concentrations. This limit is only reached after 2250 in RCP8.5 (Riahi et al., 2011) of the scenarios considered in this paper and is unimportant for the results presented in section 3.

[3] The carbon-cycle decay timescale scaling factor α is not restricted to be > 1. With default parameters given in table 1 the value of α in the pre-industrial state is 0.11.

3 ~~Results~~

2.1 Experimental set-up

In this section we ~~initially set out the need for the FAIR model by showing that state-independent~~ describe the features of several experimental protocols that have been used to examine coupled climate-carbon-cycle feedbacks in ESMs and EMICs. These experiments, conducted with the AR5-IR, PI-IR and FAIR models, form the core of our analysis of these models in section 3.

2.1.1 Pulse-response experiments

Joos et al. (2013) documented the response of an ensemble of ESMs and EMICs to pulses of various sizes and under various background conditions (black lines in figure 3). In the PD100 experiment (100GtC pulse in approximately present-day background conditions), background emissions are diagnosed that stabilise CO_2 concentrations at 389ppm (after rising as historically observed). In a second experiment, a 100GtC pulse is added to these diagnosed background emissions in the year that CO_2 concentrations reach 389ppm and the resulting CO_2 concentration and temperature evolutions are compared to the case without the pulse emission to isolate the response to the pulse emission alone. Experiments are also conducted for a pulse of 100GtC and 5000GtC in pre-industrial background conditions (PI100 and PI5000 respectively) for a smaller sub-set of models.

We simulate these experiments with the impulse-response ~~model cannot simultaneously reproduce the observed~~ climate-carbon-cycle models by following the experimental protocol exactly as described in Joos et al. (2013). Emissions are derived consistent with the background concentration profile using an inversion of the carbon-cycle ~~response over the historical period and the future~~ equations for the AR5-IR and PI-IR models (equation 1), and the FAIR model (equation 6). A declining but non-zero low level of diagnosed emissions are required to stabilise atmospheric concentrations at the 389 ppm level for all of the models considered.

As well as investigating the response of the default FAIR parameters in these pulse-response experiments, we also investigate how parameter perturbations could allow FAIR to span the range of responses observed in the PD100 and PI100 experiments for the individual models of the Joos et al. (2013) ensemble. We fit the FAIR parameters to the individual model responses in a two-step process. First, the carbon-cycle ~~evolution as projected by ESMs under possible future emissions scenarios (section 3.1). We subsequently evaluate the ability~~ parameters (a_i, r_0, r_T and r_C) of the FAIR model are optimised to minimise the total combined residual sum of squares of the FAIR fit to the Joos et al. (2013) multi-model mean airborne fraction across both the PD100 and PI100 experiments. As a constraint on this fit, we fix the ratio between the r_T and r_C parameters at the value of this for the default parameters given in table 1. This is both to reduce the number of free parameters in the fitting process (the model is underconstrained as pulse-response experiments don't distinguish between temperature-induced and CO_2 uptake-induced carbon-cycle feedbacks), and because the representation of the temperature-induced carbon-cycle feedbacks in the 'radiatively-coupled' prescribed concentration increase experiment (see section 2.1.1) is more sensitive to parameter perturbations that change this ratio than to perturbations that don't alter it (not shown).

After fitting the multi-model mean as described above, we then fit the responses for individual models by minimising the combined PD100 and PI100 residual sum of squares whilst allowing only the r_0, r_T and r_C parameters to vary from the model parameters found in the multi-model mean fit (again the ratio between r_T and r_C is fixed at the value for this ratio for the default parameters). The timeseries of change in global mean surface temperature due to the pulse emission are taken as simulated by the individual models when conducting both stages of these fits.

We also consider the response of the FAIR ~~model to capture the responses shown by ESMs~~ and ~~EMICs under a range of idealised experiments (sections 3.2 and 3.3), before discussing climate response uncertainty in the FAIR model and describing a strategy to sample climate response uncertainty within the model structure (section 3.4)~~ AR5-IR models under the idealised pulse experiments of Herrington and Zickfeld (2014). Herrington and Zickfeld (2014) conducted several experiments with the UVic Earth System Model of intermediate complexity (Weaver et al., 2001). We here emulate the PULSE experiments of Herrington and Zickfeld (2014) by integrating the FAIR and AR5-IR models with historical fossil fuel and land-use CO_2 emissions together with estimates of the historical non-CO_2 radiative forcing, both backed-out from historical concentrations using the MAGICC model (Meinshausen et al., 2011a). Pulse emissions of various sizes were then applied over a two-year period from 2008 in order to restrict total all time cumulative emissions to specified totals (see Herrington and Zickfeld (2014) for details). Non-CO_2 forcings are held constant at 2008 levels after following RCP8.5 trajectories for 2005-2008.

2.2 ~~The necessity for a state-dependent impulse-response model~~

~~A key requirement for simple climate-carbon-cycle models is to reproduce the historical period and the~~

2.1.1 Exponential CO_2 increase experiments

To explore the response to sustained emissions, rather than an emission pulse, we consider the experiments of Gregory et al. (2009) and Arora et al. (2013), in which ESMs are subjected to specified rates of increase in CO_2 concentrations. Concentrations were increased from pre-industrial values at 0.5%yr^{-1}, 1%yr^{-1} and 2%yr^{-1} respectively and consistent emissions were diagnosed for different configurations of the ESMs: a "biogeochemically-coupled" experiment, where the carbon-cycle is only allowed to respond to the direct effect of increasing CO_2 concentrations and not to the resultant warming; a "radiatively-coupled" experiment in which the climate system is allowed to respond to the radiative forcing of CO_2 but the carbon-cycle is only allowed to respond to the simulated warming and not to increasing CO_2; and a "fully-coupled" experiment in which the carbon-cycle is allowed to respond to both warming and increased CO_2. Such idealised scenarios can be highly informative with regard to the physical drivers of carbon-cycle feedbacks under increasing emissions.

Within the FAIR framework we recreate the "biogeochemically-coupled" experiment by setting r_T =0, and approximate the "radiatively-coupled" experiment by evaluating the difference between the "fully-coupled" and "biogeochemically-coupled" experiments (net out-gassing of carbon, the simulated response to the "radiatively-coupled" experiment in the ESMs, cannot be directly simulated in impulse-response models as a pulse emission of carbon always decays over time). Although Gregory et al. (2009) found that the relationship between the experiments was not simply a linear summation at high CO_2 concentrations, this serves as an adequate approximation for our purposes, since our objective is the correct representation of aggregate feedbacks from different

effects in the FAIR model as opposed to a more complex linear and non-linear partitioning. In all experiments concentrations are increased at the prescribed rates until they reach four times their pre-industrial values.

2.1.2 Uncertainty sampling with FAIR

Uncertainty in the thermal response to radiative forcing typically tends to be the dominant factor in the uncertainty in the response of the global climate system to CO_2 emissions (Gillett et al., 2013). ECS and TCR co-vary in global climate models (Knutti et al., 2005; Millar et al., 2015), with TCR typically considered the more policy-relevant parameter and the parameter better constrained by climate observations to date (Frame et al., 2006). Hence varying ECS alone in a probabilistic assessment risks introducing an implicit distribution for TCR that is inconsistent with available observations. Millar et al. (2015) observed that, within the coupled models of the CMIP5 ensemble, TCR and the ratio TCR/ECS (referred to as the Realised Warming Fraction or RWF) are approximately independent. IPCC-AR5 provided formally assessed uncertainty ranges for TCR and ECS (Collins et al., 2013) but not for their ratio. RWFs for the CMIP5 models lie within the range 0.45-0.7, while observationally-constrained estimates typically lie in the upper half of this range (Millar et al., 2015).

We assess the impact on the uncertainty in the FAIR parameters on the representation of the response to a 100GtC pulse emission of CO_2 in 2020 (against a background RCP2.6 concentrations (van Vuuren et al., 2011)) via a large ensemble (300 members) of draws from distributions representative of assessed uncertainty in these parameters. As IPCC-AR5 likely (>66% probability) ranges for a physical climate parameter attempt to capture structural uncertainties that might be present in all studies, IPCC-AR5 likely intervals are generally comparable to the 90% confidence intervals in the underlying studies as opposed to central 66% of the distribution. IPCC-AR5 gives no assessment of the shape of the distribution associated with structural uncertainty as, by definition, this encompasses "unknown unknowns" that are not included in any model or study available. For quantitative modelling purposes, likely ranges are best interpreted as 5-95 percentiles of input distributions for IPCC-AR5 assessed parameters, provided a similar "structural degradation" is applied to interpret the 5-95 percentiles of output quantities as corresponding only to a likely range, propagating the possibility of structural uncertainty in the assessed parameter through the study.

We here assume a bounded (between 0 and 1) Gaussian distribution for RWF (with 5-95 percentiles of 0.45-0.75) and a log-normal distribution for TCR (with 5-95 percentiles of 1.0-2.5K), reproducing the positive skewness (fat high tail) of many estimated distributions for this parameter. A log-normal distribution has some theoretical justification as an appropriate shape for the distribution of a so-called "scale parameter" (one in which uncertainty increases with parameter size) which is arguably the case for TCR (Pueyo, 2012). Convolving these distributions gives a corresponding ECS 5-95 percentile range of 1.6-4.5K, in good agreement with the IPCC-AR5 assessed likely range (1.5-4.5K).

The short thermal response timescale, d_2, is an important determinant of the Initial Pulse-adjustment Time (IPT - the initial e-folding adjustment time of the temperature response to a pulse emission of CO_2 (NAS, 2016)). We sample d_2 using a log-normal distribution (as d_2 is a positive definite parameter) with 5-95% probability interval of 1.6-8.4 years, corresponding to the minimum of the CMIP5 range given in Geoffroy et al. (2013) as the 5th percentile and the HadCM3 value of 8.4 years as the 95th percentile. We consider uncertainties in the carbon cycle by perfectly positively-correlated sampling of r_0, r_T and

r_C with Gaussian distributions for which the 5-95% probability intervals equal to +/- 13% (present-day ~~state of~~ $iIRF_{100}$ +/- 7 years) of their default value. The 300 random draws from each of the above distributions are then used as model parameters for the ~~climate system successfully. Compatibility with~~ integration of the ~~present-day climate state can be important for accurately assessing the scale of future mitigation ambition required to achieve specific policy targets (Rogelj et al., 2011). Atmospheric~~ 2020 100GtC pulse-response scenario described above, leading to a 300 member ensemble of climate outcomes indicative of the combined uncertainty in the FAIR input parameters.

3 Results and discussion

3.1 The necessity for a state-dependent impulse-response model

When the AR5-IR model is integrated under estimated historical emissions from the Global Carbon Project (GCP) (Le Quéré et al., 2015) st from an assumed quasi-equilibrium in 1850, atmospheric CO_2 concentrations increase faster than observed ~~when computed from estimated historical emissions (Le Quéré et al., 2015) with the AR5-IR model~~ (Figure 1a). This ~~leads to~~ is indicative of the under-efficiency of the AR5-IR carbon sinks when continuously integrated over the observed period, resulting in a bias of over 30ppm in 2011 ~~concentrations, due to the slower than observed decay of~~ CO_2 concentrations. Similarly this under-efficiency of carbon sinks requires lower than observed emissions to simulate the observed timeseries of atmospheric CO_2 ~~from the atmosphere over the historical period. The AR5-IR displays a too-large instantaneous airborne fraction over the entire historical period and is less consistent with the observations than the FAIR model (Figure~~ concentrations (figure 1b). Whilst sinks are maintained at their pre-industrial efficiency throughout (by definition) for the PI-IR model, it is only at around 1980 when the FAIR airborne fraction rises above the PI-IR airborne fraction (figure 1c). ~~The~~ This arises due to a combination of a lower pre-industrial $iIRF_{100}$ in the FAIR model for a 100GtC pulse (see section 3.2) as $\alpha < 1$ in the FAIR pre-industrial state, annual emissions that are much less than 100GtC (reducing the $iIRF_{100}$ for these emissions relative to the 100GtC pulse in FAIR but not in PI-IR~~model maintains a lower instantaneous airborne fraction than the~~), and the PI-IR model (and the AR5-IR model~~throughout the historical period, and matches the observed record much better, however neither state-independent impulse-response model matches observations as well as the state-dependent FAIR model. Large~~) not capturing temporary reductions in the airborne fraction associated with volcanic-forced cooling (figure 1d) mediated through the temperature-induced feedback on the carbon-cycle represented in FAIR. Figure 1c shows large amplitude variations in the observed instantaneous airborne fraction ~~can be seen in the observational record~~ that are likely to be driven in large part by unforced variability in the Earth-system and as such ~~would not be expected~~ we would not expect these oscillations to be reproduced by any of ~~these simple~~ the simple climate-carbon-cycle models. More complex carbon-cycle models are required to understand the drivers of these variations and any implications that they have for future carbon-cycle responses. ~~A similar relationship between the models is seen for emissions derived from each model consistent with prescribed observed CO_2 concentrations (Figure 1b), where required emissions are too low relative to observed values over much of the historical period for both the~~ Observed anomalies of global-mean temperature are reproduced well in the FAIR and PI-IR models (figure 1d),

but present day warming is too large in the AR5-IR ~~and PI-IR models~~model, driven by the substantially higher-than-observed present day CO_2 concentrations.

Another key test of simple coupled climate-carbon-cycle models is the ability to replicate the response of ESMs to possible scenarios of future emissions. Commonly-used future scenarios are generally defined in terms of concentration pathways (Van Vuuren et al., 2011) and therefore do not have a model-independent set of emissions associated with them. In this paper we drive all three simple impulse-response climate-carbon-cycle models by a single set of emissions for each future scenario that are ~~derived~~ diagnosed from the MAGICC model ~~(Meinshausen et al., 2011b)~~ (Meinshausen et al., 2011a) in order to allow a comparison of ~~both concentrations and temperatures~~ simulated concentrations between simple models driven by identical inputs. MAGICC has been shown to be a good emulator of the CMIP5 ensemble and therefore offers a comparison by proxy to the projection of CMIP5 ESMs (Meinshausen et al., 2011a). Whilst the PI-IR model might do a better job than the AR5-IR model of reproducing historical concentrations, under high future emissions scenarios such as RCP8.5~~(Riahi et al., 2011)~~, it underestimates end of century concentrations, relative to MAGICC, to an even greater extent than the AR5-IR model (Figure 2a) and concentrations fall from peak even quicker than MAGICC under the high mitigation RCP2.6 scenario (Figure 2b). ~~It is clear~~ The lack of saturation of carbon sinks in the AR5-IR model prevents the simulated concentrations keeping pace with MAGICC in the second half of the 21st century (under RCP8.5) despite having higher concentrations than MAGICC over the historical period and until approximately 2070. AR5-IR concentrations peak significantly higher than MAGICC under RCP2.6 and also decline faster after the concentration peak than simulated in MAGICC and FAIR. The deviation of both of these two models from MAGICC clearly indicates that any state-insensitive impulse-response model is therefore unsuitable, unless modified, for ~~calculations of, for example, the social cost of carbon against realistic baseline trajectories or~~ long integrations with historical and projected emissions.

The FAIR model compares well to MAGICC ~~, particularly for the ambitious mitigation scenario. There is some divergence~~ in both RCP8.5 and RCP2.6, scenarios that span the range of plausible future emissions trajectories. Concentrations simulated by FAIR are marginally higher than MAGICC after 2100 in ~~the high emission scenario~~RCP8.5, but the behaviour of MAGICC (or indeed any other model) under these more extreme forcing scenarios has not been verified. Additionally, concentrations in FAIR peak at a slightly lower value than MAGICC in the RCP2.6 scenario. Whilst comparing the performance of one simple model to another is not as rigorous a test of model performance as comparing directly to the behaviour of ESMs, it is encouraging that the FAIR model shows a close correspondence with a well-known and well-used simple model that has been used extensively to emulate the response of ESMs (Rogelj et al., 2012).

3.2 Response to pulse emission experiments

~~The social cost of carbon is conventionally calculated by applying a pulse emission of a specified magnitude of carbon in near to present-day conditions as a perturbation on top of a certain future emission scenario (NAS, 2016). As calculating the social cost of carbon is a key element of cost-benefit analysis of climate change policy in IAMs, simple climate-carbon-cycle models used in IAMs should aim to reproduce the dependencies of the response to the perturbation on pulse size and background state that has been highlighted in ESMs and EMICs (Joos et al., 2013; Herrington and Zickfeld, 2014).~~

~~Joos et al. (2013) documented the response of an ensemble of ESMs and EMICs to pulses of various sizes and against various different background conditions (black lines in Figure 3). In the PD100 experiment (~~Figure 3 shows the response to a pulse-emission of CO_2 of differing magnitudes and against different backgrounds. For a 100GtC pulse ~~in~~ of carbon set against an approximately present-day ~~background conditions - upper two panels), future emissions are derived that stabilise concentrations at (~~389ppm~~and held constant thereafter. A declining but sustained low level of diagnosed emissions are required to stabilise atmospheric concentrations at a constant level (Figure ~~) concentration background (PD100 - figure 3a). ~~In a second experiment, a 100GtC pulse is added to these calculated emissions in the year that concentrations exceed 389ppm and the resulting concentration and temperature anomalies are compared to the case without the pulse emission to isolate the coupled response to the pulse emission alone (Figure 3b). After 100 years the pulse in the concentration anomaly in the fully coupled FAIR model has decayed to 0.46 ~~, the FAIR simulated concentration anomaly associated with the pulse decays to 0.43 of its initial value after 100 years, slightly greater than the multi-model average of the ESM responses ~~of (~~0.41~~, but, the ~~). The FAIR $iIRF_{100}$ of ~~53 years is consistent with ~~50.5 years lies within the ESM multi-model ~~mean of 52.4 years (Joos et al., 2013)~~spread (see table 2). Excluding temperature feedbacks (the "biogeochemically-coupled~~' version ~~" version of FAIR - setting $r_T = 0$) on the carbon-cycle in FAIR increases the decay of the ~~temperature ~~concentration response to the pulse over the century following the pulse emission~~which reduces ~~, reducing the $iIRF_{100}$ ~~airborne fraction by 11~~by 10%. The ~~"fully-coupled" FAIR model shows temperature ~~AR5-IR, PI-IR and the FAIR model all show temperature anomalies due to the pulse initially adjusting rapidly followed by near-constant temperature over the remainder of the century ~~. ~~as displayed by the Joos et al. (2013) multi-model mean (with internal variability superimposed on-top of this signal). Peak temperature anomaly is achieved after 12 years, consistent with the value of 10 years found by Ricke and Caldeira (2014).

Figure 3~~c and 3d also ~~b and 3c show the response to a 100GtC and a 5000GtC pulse respectively, applied in pre-industrial conditions (named PI100 and PI5000 respectively). ~~Similarly to the response shown by ESMs, the ~~The 100GtC pre-industrial pulse decays faster than the present-day case (as in the ESMs and EMICs), due to reduced saturation of the land and ocean carbon sinks ~~. With these parameters, the FAIR ~~in the background state. FAIR simulates an $iIRF_{100}$ ~~is approximately 30~~approximately 32% lower in the pre-industrial case ~~compared ~~(34.3 years) relative to the present day, ~~consistent with corresponding ratio in the Joos et al. (2013) ensemble , with its value of 36 years within the ~~lying just within the ensemble spread of PI100 $iIRF_{100}$ simulated by the models of Joos et al. (2013) (34-47 years~~range of the ESMs~~). The magnitude of the FAIR simulated temperature response is similar in both the PD100 and PI100 cases due to the increased radiative efficiency of a pulse of CO_2 at lower background concentrations counteracting the faster decay of carbon out of the atmosphere. ~~The 89~~FAIR simulates a 100% increase of $iIRF_{100}$ in the 5000GtC pre-industrial pulse relative to the 100GtC pre-industrial ~~, whilst smaller than ~~pulse, consistent with the approximate doubling observed in the ESMs~~, shows that the FAIR model ~~. This clearly demonstrates that FAIR can capture the dependence of the pulse-response on pulse size as well as background conditions, whilst the AR5-IR model displays identical pulse response independent of pulse size or background conditions. The very rapid drawdown of CO_2 simulated by FAIR in the initial timestep following the 5000GtC pulse is in-part a function of the annual timestep in the model (carbon sinks remain at their pre-industrial efficiencies over the entirety of the first year despite accumulating a substantial amount of carbon over that period) and could be alleviated by using a smaller timestep.

~~A difference between the FAIR model and the ESMs is that restricting~~ Restricting temperature-induced feedbacks on the carbon-cycle does not result in a substantial reduction in the $iIRF_{100}$ for the pre-industrial 100GtC pulse experiment (the "fully-coupled" and "biogeochemically-coupled" experiments lie on top of each other in figure 3~~c~~b), whereas a 13% reduction in $iIRF_{100}$ is observed for the ~~ESMs (Joos et al., 2013) (not shown). It is only for~~ Bern3D-LPJ model examined in Joos et al. (2013). This is due to the FAIR formalism, where evolution in global mean temperature and cumulative carbon uptake increase $iIRF_{100}$ from a fixed pre-industrial value (r_0), therefore preventing substantial differences between "fully-coupled" and "biogeochemically-coupled" PI100 experiment in which $iIRF_{100} \sim r_0$. For the 5000GtC pre-industrial pulse experiment ~~that~~ we see a reduction in the $iIRF_{100}$ (17%) associated with suppression of the temperature-induced feedbacks on the carbon cycle in FAIR~~.~~, consistent with the approximate 15% reduction in $iIRF_{100}$ for Bern3D-LPJ. As the impact of the temperature-induced carbon-cycle feedbacks have only been assessed in a single EMIC, it is unclear how consistent or not FAIR is with the indicative range of behaviour that would be simulated by the ESM and EMIC ensemble.

~~Significant~~ In figure 4 we show emulations of the individual models in the Joos et al. (2013) ensemble using a single set of parameters for both the PD100 (figure 4a) and PI100 (figure 4b) simulations. In the fitting process we choose to fix the ratio of the r_T and r_C parameters to the default value (see section 2.1.1) as the fully coupled PD100 and PI100 experiments do not distinguish between temperature-induced and CO_2-induced feedbacks. Whilst significant diversity is seen in the range of responses to the PD100 and PI100 experiments across different ESMs/EMICs~~(grey shading in Figures 3b and 3c). Whilst this diversity is ultimately~~ , attributable to a range of differences in carbon-cycle process representations within ~~the~~ these models, variations in just a sub-set of the FAIR parameters are sufficient to span the ranges of responses in both the PD100 and PI100 experiments ~~, as well as the ratio between the two responses. Figure 4 shows this by fitting individual model responses in a two-step process. First, the carbon-cycle parameters of the FAIR model are optimised to minimise the combined residual sum of squares of the FAIR fit to the Joos et al. (2013) multi-model mean airborne fraction in the PD100 and PI100 experiments (whilst maintaining the same ratio between the r_T and r_C parameters as the default parameters given in section 2 and assuming fixed τ_i at their table 1 values). Then, as a second step, the response for individual models are fitted by minimising the combined PD100 and PI100 residual sum of squares whilst allowing only the r_0~~(fits are worse for models in which inter-annual variability is simulated), ~~r_T and r_C parameters to vary from the model parameters found in the first stage, whilst again maintaining the same~~ as well achieving the correct ratio between the ~~r_T and r_C parameters as the default and therefore reducing the effective degrees of freedom of the fit to just two. The timeseries of change in GMST are taken as given by the individual models. While a much better fit could be obtained by adjusting all the parameters of the FAIR model, this subset appears sufficient to successfully capture much of the response to both the~~ PD100 and PI100 ~~experiments for individual models , as well as their range of behaviour (Figure 4). The FAIR model offers a simple framework to emulate the range of ESMresponses whilst at the same time maintaining the dependency on background condition and pulse size for the specific model in question~~responses for models across the ESM/EMIC ensemble.

~~As an final test of the FAIR model's sensitivity to pulse size, we also consider the response of the FAIR and~~

Experiment - $iIRF_{100}$ (yr)	AR5-IR ~~models under the idealised pulse experiments of Herrington and Zickfeld (2014). Herrington a~~
PD100	
PD100 (no climate-carbon-cycle feedbacks)	
PI100	
PI100 (no climate-carbon-cycle feedbacks)	
PI5000	
PI5000 (no climate-carbon-cycle feedbacks)	

Table 2. $iIRF_{100}$ values in the experiments and models shown in figure 3. The star indicates that the values for suppressed climate-carbon-cycle feedbacks are calculated assuming the 13% (PI100) and 15% (PI5000) reductions in $iIRF_{100}$ from the fully coupled experiment observed from the Bern3D-LPJ model in Joos et al. (2013), the only model which conducted the experiments.

3.2.1 Temporal dependence of CO_2-induced warming on pulse size

Ricke and Caldeira (2014) used a version of the AR5-IR model to find that the maximum warming from a pulse emissions of CO_2 occurs approximately a decade after emission, but ~~as shown here (Figure 5) and as highlighted by Zickfeld and Herrington (2015), not accounting for feedbacks on the carbon-cycle fails to capture the plateau of CO_2-induced warming over the century following emission . For all pulse sizes (denoted with different linestyles)contrasting the fully coupled FAIR (thick blue) and~~ didn't account for important sensitivities of the timing of peak warming on the size of the emission pulse (Zickfeld and Herrington, 2015). Figure 5 shows that for the present-day pulse emissions experiments of Herrington and Zickfeld (2014) (see section 2.1.1), despite concentrations rapidly peak and then subsequently declining in all cases (figure 5b), the ~~AR5-IR (red)models shows that including carbon-cycle feedbacks is essential to prevent a substantial decay in the temperature anomaly over the first 100 years following the pulse emission. At higher pulse sizes, the temperature response~~ total warming in the FAIR model fails to plateau as quickly ~~as at lower pulses, where~~ at higher pulse sizes compared to at lower pulse sizes (figure 5a), as the balance between carbon-cycle cooling and long-timescale thermal warming takes centuries to reach ~~balance (Figure 3 of Herrington and Zickfeld (2014)).~~

3.3 ~~Response to idealised concentration increase experiments~~

~~To explore the response to sustained emissions, rather than an emission pulse, we consider the experiments of Gregory et al. (2009) and Arora et al. (2013), in which ESMs are subjected to specified rates of increase in CO_2 concentrations. Concentrations were increased from pre-industrial values at 0.5%yr^{-1}, 1%yr^{-1} and 2%yr^{-1} respectively and consistent emissions were derived for different configurations of the ESMs: a "biogeochemically-coupled" experiment~~equilibrium. This behaviour was demonstrated for the UVic EMIC in figure 3 of Herrington and Zickfeld (2014), and is not captured by the state-independent AR5-IR model, where the ~~carbon-cycle is only allowed to respond to the direct effect of increasing CO_2 concentrations and not to the resultant warming ; a "radiatively-coupled" experiment in which the climate system is allowed to respond to the radiative forcing of CO_2 but the carbon-cycle is only allowed to respond to the simulated warming and not to increasing CO_2; and a "fully-coupled"~~

experiment in which the carbon-cycle is allowed to respond to both warming and CO_2 concentrations (light pastel coloured lines in Figure 6) for the 1%/yr concentration increase scenario. Such idealised scenarios can be highly informative with regard to the physical drivers of carbon-cycle feedbacks under increased emissions . Successfully emulating the approximate balance between warming-induced and biogeochemically-induced contributions to carbon-cycle feedbacks could be important for integrated assessment of solar radiation management scenarios and mitigation scenarios in which the balance of contributions to warming from CO_2 and non-CO_2 sources changes significantly in shape of the futurewarming and concentration response is independent of pulse-size. As the magnitude of future cumulative emissions are uncertain and could exceed multiple TtC under minimal climate policy scenarios (e.g. RCP8.5) correctly capturing the dependence of peak warming on injection size in simple climate-carbon-cycle models is important for correctly assessing the multi-millennial impacts of climate policy (Clark et al., 2016).

Within the FAIR framework we recreate the "biogeochemically-coupled" experiment by setting r_T =0, and approximate the "radiatively-coupled" experiment by evaluating the difference between the "

3.3 Response to idealised concentration increase experiments

The set of "fully-coupled" and "", "biogeochemically-coupled" experiments . Although Gregory et al. (2009) found that the relationship between the experiments was not simply a linear summation at high CO_2 concentrations, this serves as an adequate approximation for our purposes here, since our objective is the correct representation of aggregate feedbacks rather than a breakdown into specific contributions.

Similarly to the ESMs from Arora et al. (2013), the coupling " and "radiatively-coupled" experiments from Gregory et al. (2009) (see section 2.1.1) can help to isolate the contributions from temperature-induced and direct carbon-induced effects on overall carbon-cycle feedbacks. Coupling between temperature changes and the carbon-cycle in the FAIR model acts to suppress carbon uptake, shown by the difference between the "fully-coupled" and "biogeochemically-coupled" (thick and thin lines in Figure 6a, blue lines in figure 6a) under a 1%/yr CO_2 concentration increase scenario, consistent with the behaviour shown by the ESMs from Arora et al. (2013) (thick and thin pastel coloured lines). This is a mechanism that is absent (by construction) in the AR5-IR model. The coupling with cumulative carbon uptake in the FAIR model also increases airborne fraction in the later stages of the experiment relative to earlier stages (Figure 1c), as illustrated by the approximately linear increase in C_{acc} in the "biogeochemically-coupled" experiment, also consistent with ESM responses. A constant airborne fraction necessarily gives an approximately quadratic increase in C_{acc} in this experiment, as illustrated by the AR5-IR model. Figure Figure 6b shows C_{acc} as a function of atmospheric CO_2 concentration for several rates of prescribed concentrations increase: again, the FAIR model captures the concave-downward form of this diagnostic shown by the ESM ensemble for all rates of concentration increase, in contrast to the AR5-IR model.

Whilst oceanic Oceanic carbon-cycle feedbacks are almost exclusively driven by biogeochemical effects (Glotter et al., 2014), for simple . However, simple global climate-carbon-cycle models to be of use in representing the entire climate system, they need to need to also capture dependencies of the land carbon cycle uptake on warming. Aside from 3 ESMs that display global-mean carbon-cycles relatively insensitive to warming, Figure figure 6c shows a coherent relationship between tempera-

ture increases and the size of the carbon outgassing back to the atmosphere ~~(Arora et al., 2013). The impact of GMST increase on cumulative uptake, or the difference between the biogeochemically coupled and fully coupled experiments shown in Figure 6a, as a function of warming, indicating that values of r_T close to 4.5yr~~in the 'radiatively-coupled' FAIR experiment similar, in both shape and magnitude, to that displayed in the ESMs under the 1%/~~K allow the FAIR model to reproduce this relationship well~~yr concentration increase experiment of Arora et al. (2013). 1%yr^{-1}, 0.5%yr^{-1} and 2%yr^{-1} experiments in FAIR all lie along the same line in ~~panel (c)~~figure 6c, indicating minimal scenario dependence of this effect in FAIR, in contrast to the two ESMs analysed in Gregory et al. (2009).

The initial decrease in cumulative airborne fraction (the ~~time-integrated instantaneous airborne fraction~~ fraction of all past emissions remaining in the atmosphere) followed by a subsequent increase (~~Figure~~ figure 6d) ~~displayed by the FAIR model~~ simulated by FAIR is a feature of the response of many ESMs under a 1%/yr increasing CO_2 scenario. In contrast, the IPCC-AR5 model shows a steady decrease in the cumulative airborne fraction with higher concentrations due to the state-invariant rates at which a pulse of carbon is removed from the atmosphere. The initial decrease in cumulative airborne fraction followed by subsequent increase can be understood in terms of the saturation of carbon sinks. If atmospheric anomalies of carbon decay with fixed timescales, τ_i (as in the AR5-IR model case), then the instantaneous airborne fraction remains constant in time, which necessarily means that cumulative airborne fraction must decline over time (as emissions from previous years decay further, so the cumulative fraction of the emitted carbon continually decays from the instantaneous airborne fraction). However, if carbon sinks become saturated, the instantaneous airborne fraction would be expected to increase with time (this is represented in the FAIR model by increases to the decay timescales through the parameterised increase in ~~iRF100). As~~ iRF_{100}). Therefore, as more recent emissions (which increase monotonically ~~under the 1%/yr scenario~~with time in the prescribed concentration increase experiments) have a higher instantaneous airborne fraction, the initial decrease in cumulative airborne fraction stops and ~~then~~ subsequently begins to increase as this accelerating saturation becomes the dominant effect.

3.4 ~~Uncertainty and probabilistic~~ Probabilistic parameter sampling within the FAIR model

~~Uncertainty is a crucial factor in the integrated assessment of climate policies. Despite significant advances in climate system understanding, non-negligible uncertainties remain in the responses of the coupled climate-carbon-cycle system to emissions of CO_2 (Gillett et al., 2013). Uncertainty in aspects of the climate response to CO_2 remains broad and climate policies have to be constructed and assessed in the light of this continued uncertainty (Millar et al., 2016). Integrated assessment activities require a representation of the physical climate system that can transparently and simply sample physically-consistent modes of climate response uncertainty.~~

The impulse-response formulation of the physical climate response to radiative forcing used by both the AR5-IR and FAIR models (equation 3) offers a convenient structure for simply sampling plausible ranges of TCR and ECS~~, as a~~. Figures 7a and 7b show the isolated impact of thermal response uncertainty under idealised prescribed concentration scenarios, namely a 1%/yr CO_2 concentration increase (figure 7a) and an instantaneous quadrouple of CO_2 concentrations from pre-industrial values (figure 7b). A unique combination of TCR and ECS (for fixed response time-scales d_j) are associated with a unique combination of the model parameters ~~c_j (see Millar et al. (2015) for details). Panels~~ q_j (equations 4 and 5).

The blue shading in panels a) and b) ~~in~~ of figure 7 show ~~how the~~ the simulated temperature response for the likely range of TCR and ECS as assessed by IPCC-AR5 (TCR: 1.0-2.5K and ECS: 1.5-4.5K)~~can spanned for assessing the climate response to any radiative forcing scenario~~. Expressed in terms of q_i, the top end of the IPCC-AR5 thermal uncertainty range (TCR=2.5K, ECS=4.5K) corresponds to $q_1 = 0.57\text{KW}^{-1}\text{m}^2$ and $q_2 = 0.63\text{KW}^{-1}\text{m}^2$, with the lower end (TCR=1.0K, ECS=1.5K) corresponding to $q_1 = 0.14\text{KW}^{-1}\text{m}^2$ and $q_2 = 0.26\text{KW}^{-1}\text{m}^2$. The thermal response model given by equation 3 fully spans the range of responses seen in the CMIP5 ensemble with these ranges of parametric uncertainty for these fixed concentration integrations. AR5-IR results are not shown in figure 7a and 7b as the experiments shown are conducted with prescribed CO_2 concentrations and therefore demonstrate thermal response uncertainty only - which is identical in the FAIR and AR5-IR/PI-IR models.

A robust feature of the carbon-cycle response in all ESMs is an increase in the cumulative airborne fraction over time associated with a saturation of carbon sinks (upward curving black lines in ~~Figure 7cimply that a rising fraction of cumulative emissions remain resident in the atmosphere~~figure 7c). Unlike the AR5-IR model, which displays a slowly declining cumulative airborne fraction over time due to the state-independence of its response function, ~~coherent~~ correlated and time invariant perturbations of +/-13% ~~(approximately equivalent to a present-day iIRF$_{100}$ change of +/- 7 years) to~~ to the r_0, r_T and r_C parameters (combined with perturbations to ~~e_1 and e_2~~ q_1 and q_2 consistent with the IPCC-AR5 likely ranges) in the FAIR model ~~all show increasing cumulative airborne fraction over time (blue shading in Figure 7c) and approximately span~~ approximately spans the range of responses seen in the CMIP5 models under a $1\%\text{yr}^{-1}$ concentration increase scenario (blue shading in figure 7c) and show increasing cumulative airborne fraction over time as a common feature.

Crucially, the FAIR model also captures the ~~straight-line relationship~~ approximately linear relationship (to first order) between cumulative carbon emissions and human-induced warming (~~Figure~~ figure 7d) that was highlighted in the IPCC 5th Assessment, and is becoming an integral part of climate change policy analysis (Millar et al., 2016). When integrated ~~, the FAIR model,~~ with parameter settings given in section 2, the FAIR model has a Transient Response to Cumulative Emissions (TCRE[4]) =~~1.5~~1.3K/TtC (thick blue line in ~~Figure~~ figure 7d). As there is some downward curvature apparent in the FAIR relationship between warming and cumulative emissions across the range of cumulative emissions shown in figure 7d (as also displayed by the CMIP5 ESMs) the numerical value of the TCRE is only exactly valid for the first 1000GtC of an emissions injection. Perturbations to the model parameters as described above (and identical to Figure 7c) allow the IPCC-AR5 likely TCRE range of 0.8-2.5K/TtC to be spanned (Figure 7d). In contrast, the AR5-IR model, with a constant airborne fraction, shows a ~~clear~~ more-pronounced concave-downward shape in ~~a~~ the plot of realised warming against cumulative carbon emissions, ~~because~~ as the decline of the cumulative airborne fraction is unable to compensate (as it does in more complex models) for the logarithmic relationship between CO_2 concentration and radiative forcing (Millar et al., 2016). The ~~FAIR model also displays some~~ curvature in the relationship between warming and cumulative emissions in FAIR displays most prominent curvature at high cumulative emissions, consistent with the behaviour of ESMs (Leduc et al., 2015).

[4]TCRE is defined as the annual mean global surface temperature change per unit of cumulative CO_2 emissions, usually 1000GtC, in a scenario with continuing emissions (Collins et al., 2013).

Integrated assessment of climate change often requires probabilistic projections of the climate response to CO_2 emissions, partly in order to capture and assess the possibility of extreme, and highly costly, sensitivities within the Earth system (often called "fat-tailed" outcomes) (Weitzman, 2011). Uncertainty in the global climate response to emissions of CO_2 is associated with several factors, which are each considered in turn here.

Uncertainty in the thermal response to radiative forcing typically tends to dominate Figure 8 shows the combined effect of uncertainty in the response of the global climate system to CO_2 emissions (Gillett et al., 2013). ECS and TCR co-vary in global climate models (Knutti et al., 2005; Millar et al., 2015), with TCR typically considered the more policy-relevant parameter and the parameter better constrained by climate observations to date (Frame et al., 2006; Gillett et al., 2013). Hence varying ECS alone in a probabilistic assessment risks introducing an implicit distribution for TCR that is inconsistent with available observations. Millar et al. (2015) observed that, within the coupled models of the CMIP5 ensemble, TCR and the ratio TCR/ECS (referred to as the Realised Warming Fraction or RWF) are approximately independent. IPCC-AR5 provided formally assessed uncertainty ranges for TCR and ECS (Collins et al., 2013) but not for their ratio. RWFs for the CMIP5 models lie within the range 0.45-0.7, while observationally-constrained estimates typically lie in the upper half of this range (Millar et al., 2015).

As IPCC-AR5 likely (>66% probability) ranges for a physical climate parameter attempt to capture structural uncertainties that might be present in all studies, therefore, IPCC-AR5 likely intervals are generally comparable to the 90% confidence intervals in the underlying studies. IPCC-AR5 gives no assessment of the shape of the distribution associated with structural uncertainty as, by definition, this encompasses "unknown unknowns" that are not included in any model or study available. For quantitative modelling purposes, likely ranges are best interpreted as 5-95 percentiles of input distributions for IPCC-AR5 assessed parameters, provided a similar "structural degradation" is applied to interpret the 5-95 percentiles of output quantities as corresponding only to a likely range, propagating the possibility of structural uncertainty in the assessed parameter through the study. We here assume a bounded (between 0 FAIR model parameters, both thermal parameters (the joint distribution of TCR and 1) Gaussian distribution for RWF and a log-normal distribution for TCR, reproducing the positive skewness (fat high tail) of many estimated distributions for this parameter. A log-normal distribution has some theoretical justification for a so-called "scale parameter", or one in which uncertainty increases with parameter size, which is arguably the case for TCR (Pueyo, 2012). Convolving a bounded Gaussian RWF distribution (with 5-95 percentiles of 0.45-0.75) with a log-normal TCR distribution (with 5-95 percentiles of 1.0-2.5K), gives a corresponding ECS 5-95 percentile range of 1.6-4.5K, in good agreement with the IPCC-AR5 assessed likely range (1.5-4.5K). A sample of 300 ECS and TCR values drawn from these distributions are ECS is shown in figure 8a.

Another key uncertainty is the short thermal response timescale, d_1, an important determinant of the Initial Pulse-adjustment Time (IPT) , the initial e-folding adjustment time of the temperature response to a pulse emission of CO_2 (NAS, 2016). This can be approximated for the FAIR model as IPT=$d_1(1 - a_3)$. Throughout this paper we have used the IPCC-AR5 default value for d_1 of 8.4 years, but this is longer than indicated by most climate models (Geoffroy et al., 2013). We therefore sample the short thermal response timescale using a Gaussian distribution with a median value of 4 years and a 5-95% probability interval

of 2-8 years. This corresponds to an approximate median estimate of 2.8 years with 5-95 percentile range of 1.4-5.6 years for the IPT.

We consider uncertainties in the carbon cycle by sampling r_0, r_T) and r_C with Gaussian distributions of 5-95% probability intervals equal to +/- 13% (present-day $iIRF_{100}$ +/- 7 years) of their default value. Combined with the thermal response uncertainty sampling, the carbon-cycle feedback parameters. Representative uncertainty in FAIR parameters (see section 2.1.2) propagates into an emergent 5-95% range (based on 300 draws from the input parameter distributions) for TCRE (figure 8eb) of 1.0-2.5K/TtC is when integrated under a 1%/yr CO_2 concentration increase scenario (figure 8b), broadly consistent with the IPCC-AR5 *likely* range (0.8-2.5K/TtC).

Sampling these parameters independently, as described above, produces a range of responses to a 100 GtC pulse emissions The temperature anomaly associated with a 100GtC pulse in 2020 against the background of the an RCP2.6 scenario (figure 8d). background is shown in figure 8c. Across the ensemble a range of responses in both magnitude and shape are observed. However, we consistently observe a rapid warming on the order of a decade followed by an approximate warming plateau (at differing values) that persists for a century or more. Such behaviour is broadly consistent, in all cases, with the range of pulse-response behaviour observed across the ensemble of ESMs in Joos et al. (2013).

4 Conclusions

In this paper we have presented a simple Finite Amplitude Impulse Response (FAIR) climate-carbon-cycle model, which adjusts the carbon-cycle impulse-response function based on feedbacks from the warming of the climate and cumulative CO_2 uptake , through a parameterisation of the 100-year integrated impulse-response function, $iIRF_{100}$. This metric provides a potential We use this metric of carbon-cycle response as a parallel to those used to assess the thermal response to radiative forcing, namely the Transient Climate Response (TCR) and the Equilibrium Climate Sensitivity (ECS). Although a useful composite metric for the coupled climate-carbon-cycle system exists, the Transient Climate Response to Cumulative Emissions (TCRE), future studies of carbon cycle behaviour could report on ranges of $iIRF_{100}$, and importantly for carbon cycle feedbacks, the evolution of this metric over time under specific emissions scenarios, in order to isolate the changing response of the carbon cycle.

We have shown that including both explicit CO_2 uptake- and temperature- induced feedbacks are essential to capture ESM behaviour. Neglecting temperature-induced feedbacks on the carbon-cycle would prevent a simple climate-carbon-cycle model being able to capture the important dependences of carbon uptake on warming displayed in 'radiatively-coupled' ESM experiments, largely driven by responses of the land carbon-cycle. Similarly, neglecting CO_2 uptake-feedbacks would fail to incorporate well-understood physical mechanisms governing the response of ocean carbonate chemistry to anthropogenic CO_2 emissions.

Important dependences of the carbon-cycle response to pulse size, background conditions and the suppression of temperature-induced feedbacks are generally well captured by the FAIR model. As present-day pulse responses are an essential part of calculations of the social cost of carbon (Marten, 2011), the The inclusion of climate-carbon-cycle feedbacks in the FAIR model

offers an improvement on several simple and transparent climate-carbon-cycle models that have been proposed for policy analysis which either incorporate no feedbacks on the carbon-cycle or do not fully capture the operation of these feedbacks in ESMs. Successfully emulating the approximate balance between warming-induced and biogeochemically-induced contributions to carbon-cycle feedbacks could be important for integrated assessment of solar radiation management scenarios and mitigation scenarios in which the balance of contributions to warming from CO_2 and non-CO_2 sources changes significantly in the future.

Throughout this paper we have used $iIRF_{100}$ as a central metric of the climate system response. This represents an inherent value choice about the timescales of the coupled climate-carbon-cycle system prioritised for correct representation in a simple climate-carbon-cycle model. A time horizon of 100 years captures important aspects of the climate response to a pulse emission of CO_2 relevant over typical economic discounting timescales, whereas a longer time horizon could be used to prioritise the millennial timescale response. For studies that are primarily focused on the climate response over multi-millennial timescales it may be most appropriate to retune the values of the FAIR model to capture the dependencies shown by ESMs/EMICs over this time period (e.g. Zickfeld et al. (2013)). However, over these time periods Earth system feedbacks not simulated in ESMs and EMICs may become important, questioning the validity of emulating particular aspects of very long-term ESM behaviour. Although a useful composite metric for the coupled climate-carbon-cycle system already exists, the Transient Climate Response to Cumulative Emissions (TCRE), future studies of carbon cycle behaviour could usefully report on ranges of $iIRF_{100}$, and importantly for carbon cycle feedbacks, the evolution of this metric over time under specific emissions scenarios, in order to isolate the changing response of the carbon cycle and to enable emulators such as FAIR to span the ranges and capture the dependencies of $iIRF_{100}$ that are observed in state-of-the-art models.

We believe that the FAIR model could be a useful tool for offering a simple and transparent framework for assessing the implications of CO_2 emissions for climate policy analyses. It offers a structure that both replicates the essential physical mechanisms of the climate ~~system~~system's response to cumulative emissions, whilst at the same time it can easily be modified to sample representative climate response uncertainty in either the thermal climate response component, the unperturbed carbon-cycle or the coupled climate-carbon-cycle response to anthropogenic CO_2 emissions. Tuning of parameters within ~~the FAIR framework~~ FAIR allows the range of ESM behaviour to be emulated whilst maintaining the physically-understood dependency of the carbon-cycle pulse-response on background conditions and pulse size exhibited by a particular ESM. This model structure could thus be adapted to be an effective emulator of CMIP6 ESM responses under a variety of scenarios.

Author contributions. RJM, ZRN and MRA developed the FAIR model formulation. PF and MRA identified the need for the feedback term in the AR5-IR model while RJM developed the final formulation. MRA designed the tests and RJM made the figures, except Figure 4 which was made by ZRN. RJM wrote the first draft of the manuscript and all authors contributed to the editing and revisions of the paper.

References

Allen, M. R., Frame, D. J., Huntingford, C., Jones, C. D., Lowe, J. A., Meinshausen, M., and Meinshausen, N.: Warming caused by cumulative carbon emissions towards the trillionth tonne, Nature, 458, 1163–1166, 2009.

Arora, V. K., Boer, G. J., Friedlingstein, P., Eby, M., Jones, C. D., Christian, J. R., Bonan, G., Bopp, L., Brovkin, V., Cadule, P., et al.: Carbon–concentration and carbon–climate feedbacks in CMIP5 Earth system models, Journal of Climate, 26, 5289–5314, 2013.

Clark, P. U., Shakun, J. D., Marcott, S. A., Mix, A. C., Eby, M., Kulp, S., Levermann, A., Milne, G. A., Pfister, P. L., Santer, B. D., et al.: Consequences of twenty-first-century policy for multi-millennial climate and sea-level change, Nature Climate Change, 2016.

Collins, M. et al.: Long-term Climate Change: Projections, Commitments and Irreversibility, Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)], 2013.

Davis, S. J. and Socolow, R. H.: Commitment accounting of CO_2 emissions, Environmental Research Letters, 9, 084 018, 2014.

Flato, G. et al.: Evaluation of Climate Models, Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)], 2013.

Frame, D., Stone, D., Stott, P., and Allen, M.: Alternatives to stabilization scenarios, Geophysical Research Letters, 33, 2006.

Friedlingstein, P., Cox, P., Betts, R., Bopp, L., Von Bloh, W., Brovkin, V., Cadule, P., Doney, S., Eby, M., Fung, I., et al.: Climate-carbon cycle feedback analysis: Results from the C4MIP model intercomparison, Journal of Climate, 19, 3337–3353, 2006.

Geoffroy, O., Saint-Martin, D., Olivié, D. J., Voldoire, A., Bellon, G., and Tytéca, S.: Transient climate response in a two-layer energy-balance model. Part I: Analytical solution and parameter calibration using CMIP5 AOGCM experiments, Journal of Climate, 26, 1841–1857, 2013.

Gignac, R. and Matthews, H. D.: Allocating a 2°C cumulative carbon budget to countries, Environmental Research Letters, 10, 075 004, 2015.

Gillett, N., Arora, V., Matthews, D., and Allen, M.: Constraining the Ratio of Global Warming to Cumulative CO_2 Emissions Using CMIP5 Simulations, Journal of Climate, 26, 6844–6858, 2013.

Glotter, M. J., Pierrehumbert, R. T., Elliott, J. W., Matteson, N. J., and Moyer, E. J.: A simple carbon cycle representation for economic and policy analyses, Climatic Change, 126, 319–335, 2014.

Gregory, J. M., Jones, C., Cadule, P., and Friedlingstein, P.: Quantifying carbon cycle feedbacks, Journal of Climate, 22, 5232–5250, 2009.

Held, I. M., Winton, M., Takahashi, K., Delworth, T., Zeng, F., and Vallis, G. K.: Probing the fast and slow components of global warming by returning abruptly to preindustrial forcing, Journal of Climate, 23, 2418–2427, 2010.

Herrington, T. and Zickfeld, K.: Path independence of climate and carbon cycle response over a broad range of cumulative carbon emissions, Earth System Dynamics, 5, 409, 2014.

Hof, A. F., Hope, C. W., Lowe, J., Mastrandrea, M. D., Meinshausen, M., and van Vuuren, D. P.: The benefits of climate change mitigation in integrated assessment models: the role of the carbon cycle and climate component, Climatic Change, 113, 897–917, 2012.

Joos, F., Roth, R., Fuglestvedt, J., Peters, G., Enting, I., Bloh, W. v., Brovkin, V., Burke, E., Eby, M., Edwards, N., et al.: Carbon dioxide and climate impulse response functions for the computation of greenhouse gas metrics: a multi-model analysis, Atmospheric Chemistry and Physics, 13, 2793–2825, 2013.

Knutti, R., Joos, F., Müller, S., Plattner, G.-K., and Stocker, T.: Probabilistic climate change projections for CO2 stabilization profiles, Geophysical Research Letters, 32, 2005.

Le Quéré, C., Moriarty, R., Andrew, R., Canadell, J., Sitch, S., Korsbakken, J., Friedlingstein, P., Peters, G., Andres, R., Boden, T., et al.: Global Carbon Budget 2015, Earth System Science Data, 7, 349–396, 2015.

Leduc, M., Matthews, H. D., and de Elía, R.: Quantifying the limits of a linear temperature response to cumulative CO_2 emissions, Journal of Climate, 28, 9955–9968, 2015.

Marten, A. L.: Transient temperature response modeling in IAMs: the effects of over simplification on the SCC, Economics: The Open-Access, Open-Assessment E-Journal, 5, 2011.

Matthews, H. D., Gillett, N. P., Stott, P. A., and Zickfeld, K.: The proportionality of global warming to cumulative carbon emissions, Nature, 459, 829–832, 2009.

Meinshausen, M., Meinshausen, N., Hare, W., Raper, S. C., Frieler, K., Knutti, R., Frame, D. J., and Allen, M. R.: Greenhouse-gas emission targets for limiting global warming to 2°C, Nature, 458, 1158–1162, 2009.

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

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

Millar, R., Allen, M., Rogelj, J., and Friedlingstein, P.: The cumulative carbon budget and its implications, Oxford Review of Economic Policy, 32, 323–342, 2016.

Millar, R. J., Otto, A., Forster, P. M., Lowe, J. A., Ingram, W. J., and Allen, M. R.: Model structure in observational constraints on transient climate response, Climatic Change, 131, 199–211, 2015.

Morice, C. P., Kennedy, J. J., Rayner, N. A., and Jones, P. D.: Quantifying uncertainties in global and regional temperature change using an ensemble of observational estimates: The HadCRUT4 data set, Journal of Geophysical Research: Atmospheres, 117, 2012.

Myhre, G. et al.: Anthropogenic and Natural Radiative Forcing: Supplementary Material, Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)], 2013.

NAS: Assessment of Approaches to Updating the Social Cost of Carbon: Phase 1 Report, Tech. rep., National Academies of Sciences, Engineering, and Medicine., 2016.

Nordhaus, W. D.: Economic aspects of global warming in a post-Copenhagen environment, Proceedings of the National Academy of Sciences, 107, 11 721–11 726, 2010.

Pfeiffer, A., Millar, R., Hepburn, C., and Beinhocker, E.: The '2°C capital stock' for electricity generation: Committed cumulative carbon emissions from the electricity generation sector and the transition to a green economy, Applied Energy, 2016.

Pueyo, S.: Solution to the paradox of climate sensitivity, Climatic Change, 113, 163–179, 2012.

Revelle, R. and Suess, H. E.: Carbon dioxide exchange between atmosphere and ocean and the question of an increase of atmospheric CO_2 during the past decades, Tellus, 9, 18–27, 1957.

Riahi, K., Rao, S., Krey, V., Cho, C., Chirkov, V., Fischer, G., Kindermann, G., Nakicenovic, N., and Rafaj, P.: RCP 8.5—A scenario of comparatively high greenhouse gas emissions, Climatic Change, 109, 33–57, 2011.

Ricke, K. L. and Caldeira, K.: Maximum warming occurs about one decade after a carbon dioxide emission, Environmental Research Letters, 9, 124 002, 2014.

Rogelj, J., Hare, W., Chen, C., and Meinshausen, M.: Discrepancies in historical emissions point to a wider 2020 gap between 2°C benchmarks and aggregated national mitigation pledges, Environmental Research Letters, 6, 024 002, 2011.

Rogelj, J., Meinshausen, M., and Knutti, R.: Global warming under old and new scenarios using IPCC climate sensitivity range estimates, Nature Climate Change, 2, 248–253, 2012.

Rogelj, J., Schaeffer, M., Friedlingstein, P., Gillett, N., Van Vuuren, D., Riahi, K., Allen, M., and Knutti, R.: Differences between carbon budget estimates unravelled, Nature Climate Change, 6, 245–252, 2016.

Taylor, K. E., Stouffer, R. J., and Meehl, G. A.: An overview of CMIP5 and the experiment design, Bulletin of the American Meteorological Society, 93, 485–498, 2012.

Van Vuuren, D. P., Edmonds, J., Kainuma, M., Riahi, K., Thomson, A., Hibbard, K., Hurtt, G. C., Kram, T., Krey, V., Lamarque, J.-F., et al.: The representative concentration pathways: an overview, Climatic Change, 109, 5–31, 2011.

van Vuuren, D. P., Stehfest, E., Elzen, M. G., Kram, T., Vliet, J., Deetman, S., Isaac, M., Goldewijk, K. K., Hof, A., Beltran, A. M., et al.: RCP2.6: exploring the possibility to keep global mean temperature increase below 2°C, Climatic Change, 109, 95–116, 2011.

Van Vuuren, D. P., Van Soest, H., Riahi, K., Clarke, L., Krey, V., Kriegler, E., Rogelj, J., Schaeffer, M., and Tavoni, M.: Carbon budgets and energy transition pathways, Environmental Research Letters, 11, 075 002, 2016.

Weaver, A. J., Eby, M., Wiebe, E. C., Bitz, C. M., Duffy, P. B., Ewen, T. L., Fanning, A. F., Holland, M. M., MacFadyen, A., Matthews, H. D., et al.: The UVic Earth System Climate Model: Model description, climatology, and applications to past, present and future climates, Atmosphere-Ocean, 39, 361–428, 2001.

Weitzman, M.: Fat-tailed uncertainty in the economics of catastrophic climate change, Review of Environmental Economics and Policy, 5, 275–292, 2011.

Zickfeld, K. and Herrington, T.: The time lag between a carbon dioxide emission and maximum warming increases with the size of the emission, Environmental Research Letters, 10, 031 001, 2015.

Zickfeld, K., Eby, M., Alexander, K., Weaver, A. J., Crespin, E., Fichefet, T., Goosse, H., Philippon-Berthier, G., Edwards, N. R., Holden, P. B., et al.: Long-term climate change commitment and reversibility: An EMIC intercomparison, Journal of Climate, 26, 5782–5809, 2013.

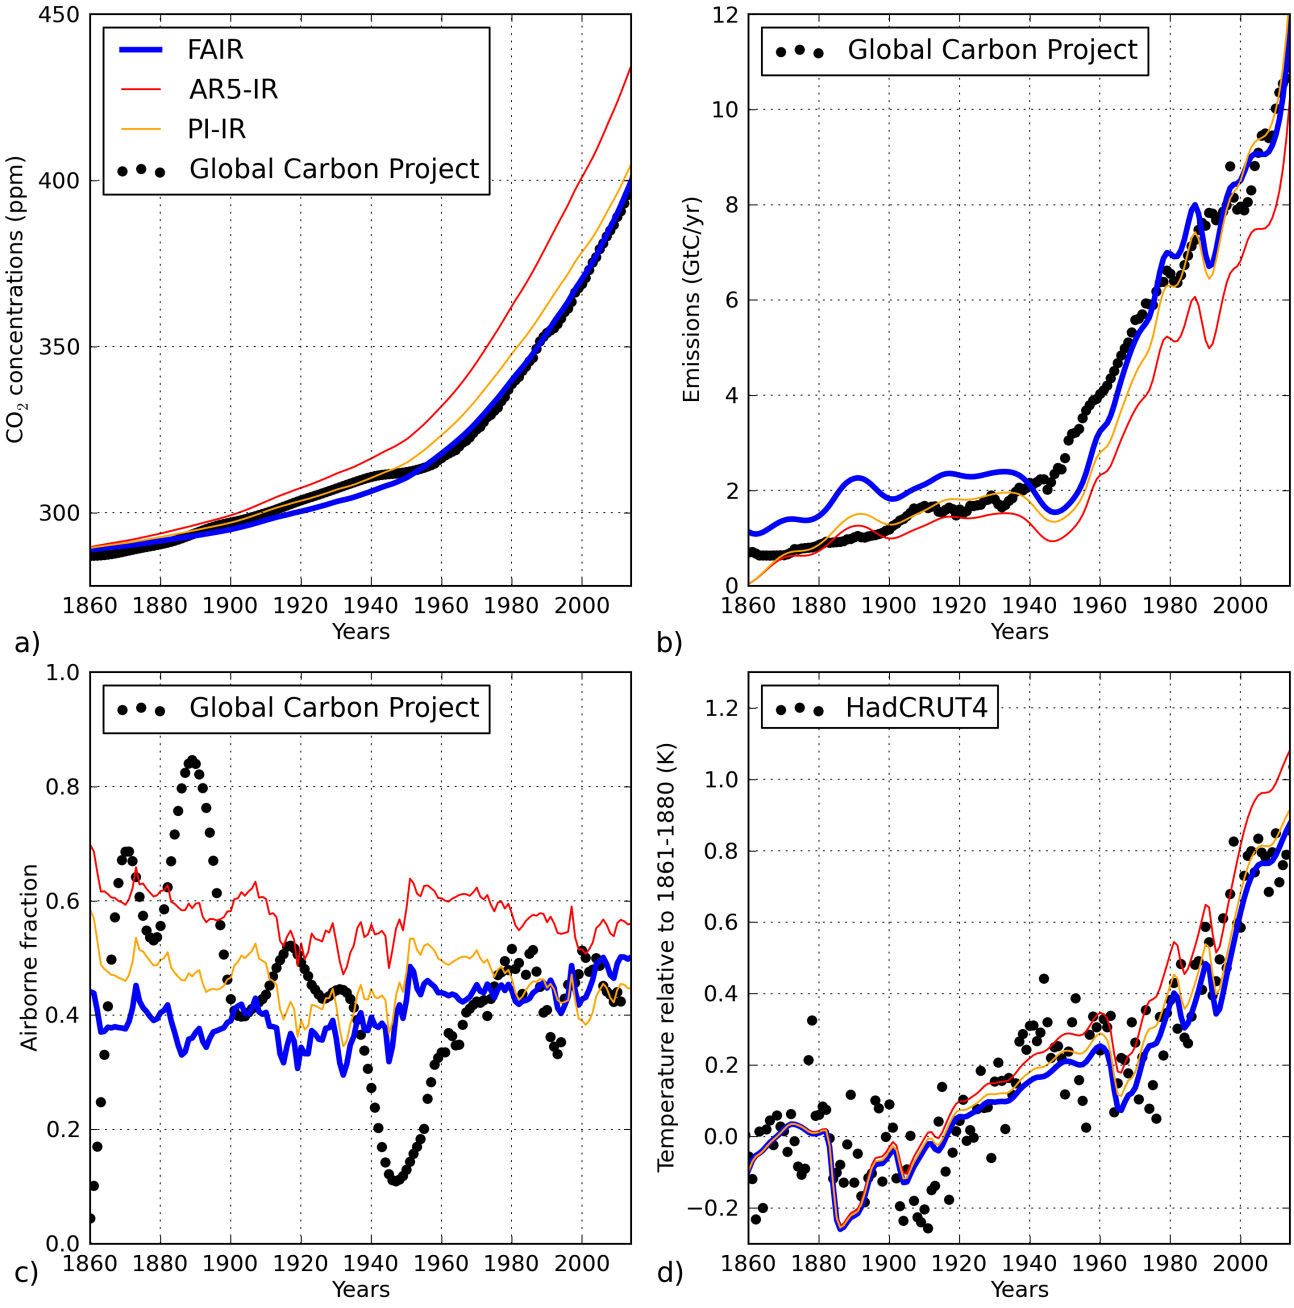

Figure 1. Historical validation of the FAIR (blue), AR5-IR (red) and PI-IR (orange) models. Panel a) shows ~~the~~ simulated CO_2 ~~concentration response~~ concentrations when ~~integrated under~~ driven by historical emissions (and historical non-CO_2 ~~radiative forcing for the RCP scenarios~~) as estimated by MAGICC. Panel b) shows ~~the derived~~ diagnosed CO_2 emissions consistent with historical concentrations. Panel c) shows the evolution of annual airborne fraction (smoothed with a 7-year running mean for the observations) ~~in the models when driven by historical emissions (as in panel a)). Panel~~ , and d) ~~shows~~ the warming anomaly, in the models when driven by historical emissions. Historical observations are shown as black dots in all panels. ~~Panels~~ , a), b) and c) ~~all~~ show ~~data from~~ Le Quéré et al. (2015) and ~~panel d) shows the~~ HadCRUT4 (Morice et al., 2012) ~~dataset~~ is shown in d). All simulations are commenced from assumed quasi-equilibrium carbon-cycle states in 1850.

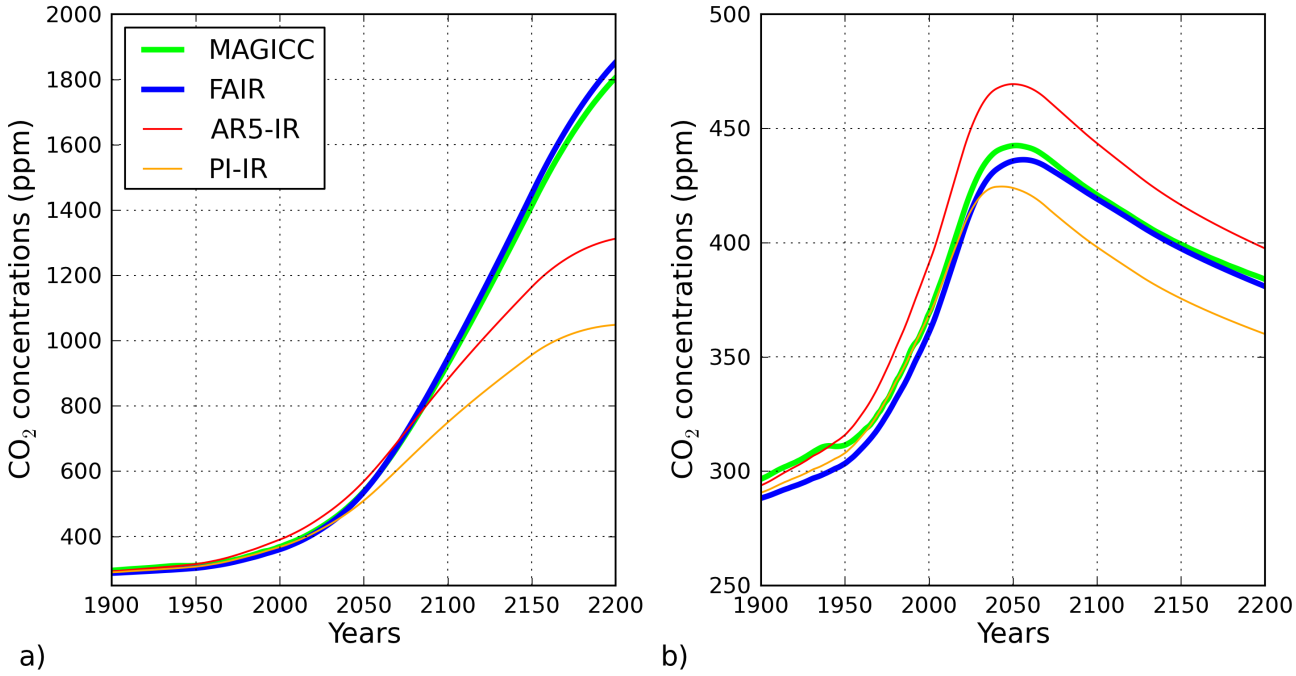

Figure 2. ~~Panels a) and b) shows the~~ CO_2 concentrations under RCP8.5 (a) and RCP2.6 ~~respectively for the~~ (b). FAIR (blue), AR5-IR (red), PI-IR (orange) and MAGICC (green) ~~models. Panel c) shows the temperature response under~~ are shown in both ~~RCP2.6 and RCP8.5~~panels. ~~Panel d) shows the evolution of total warming (full) and CO_2-induced warming (dashed) as a function of cumulative carbon emissions.~~

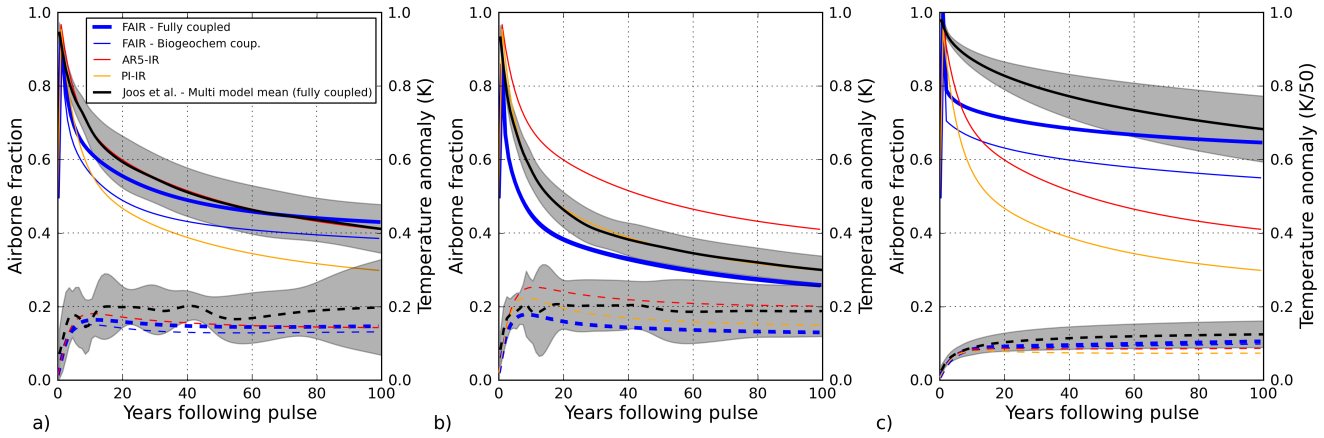

Figure 3. Response to pulse emission experiments of Joos et al. (2013). Panel a) shows the ~~"baseline" emissions (left-hand axis, solid) and warming (right-hand axis, dashed) when concentrations are stabilised at 389ppm for the FAIR (blue) and AR5-IR (red) models. Panel b) shows the~~ response to a 100GtC imposed on present-day (389ppm) background conditions (PD100 experiment)~~. Panel c~~, panel b) ~~shows~~ the response to a 100GtC pulse in pre-industrial conditions (PI100 experiment) ~~. Panel d~~and panel c) ~~shows~~ the response to a 5000GtC pulse in pre-industrial conditions (PI5000 experiment), with the warming normalised by the increase in pulse size between panels ~~c~~b) and ~~d~~c). ~~The black~~ Airborne fraction (left hand axis) is represented by solid lines in all panels ~~b~~and warming (right hand axis) by dashed lines. FAIR is shown as thick blue lines, ~~c)~~ AR5-IR as red and ~~d)~~ PI-IR as orange. The black lines in all panels shows the Joos et al. (2013) multi-model mean for airborne fraction (solid) and warming (dashed), with the ~~black~~ grey shading indicating one standard deviation uncertainty across the ensemble. Thin blue lines denote the biogeochemically-coupled version of FAIR.

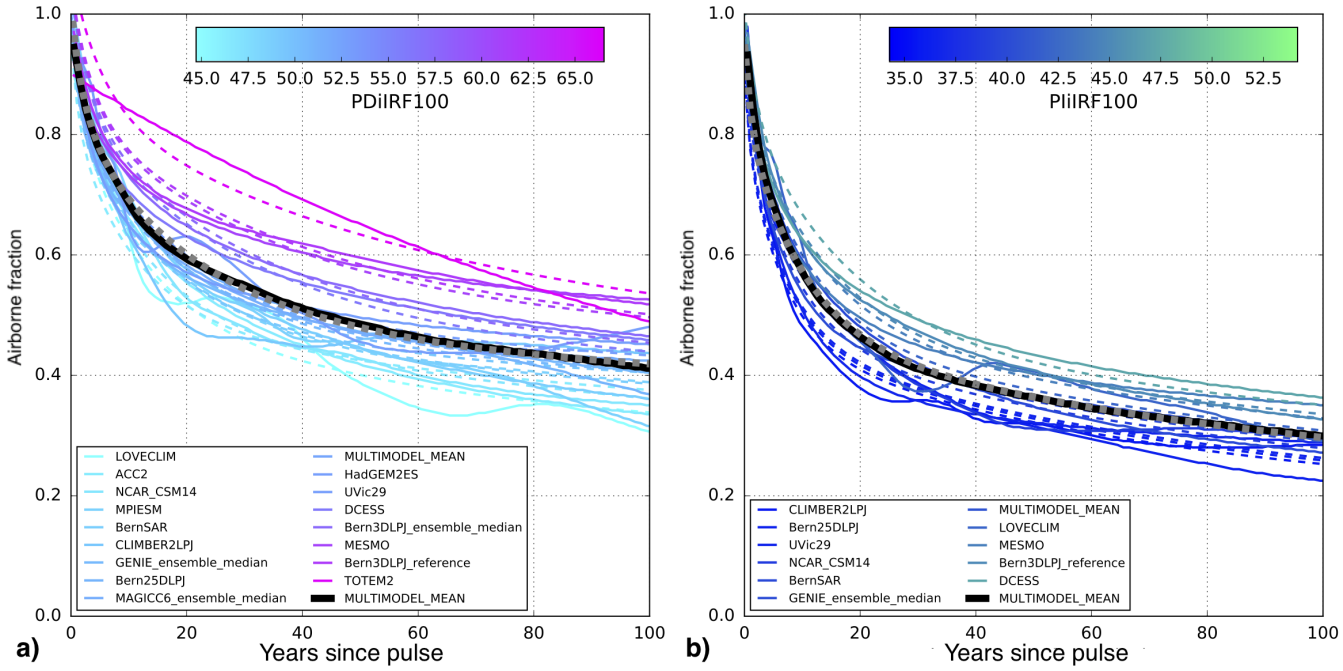

Figure 4. Fitting individual models from Joos et al. (2013) with FAIR. Panel a) shows the remaining airborne fraction for the PD100 experiment and panel b) for those models that additionally completed the PI100 experiment. Solid lines show the original model response coloured by the $iIRF_{100}$ values. Emulations with FAIR are shown by the same coloured dashed lines. The multi-model mean is shown by a solid black line with the FAIR fit denoted by a dashed grey line.

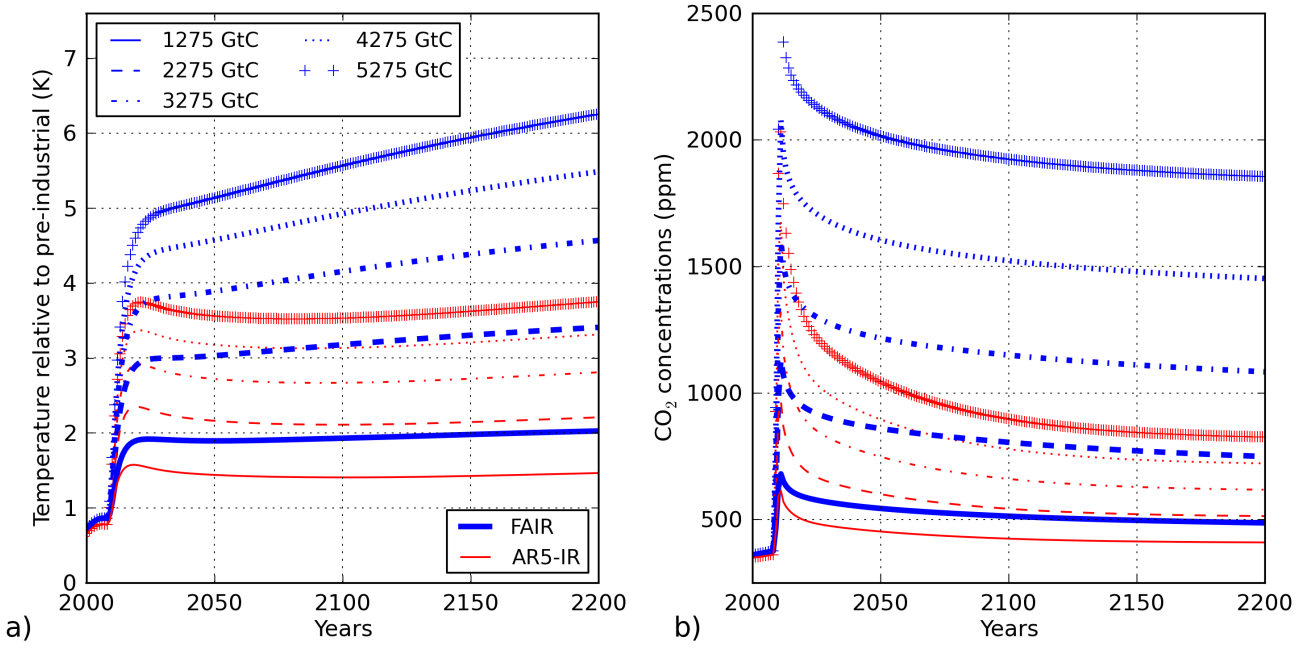

Figure 5. Dependency of temperature response on pulse size. Panel a) shows the global mean surface temperature ~~(GMST)~~ response to the pulse experiments of Herrington and Zickfeld (2014). ~~Pulse emissions are applied over a 2-year period from 2008, with differing total cumulative carbon emissions denoted by different line styles.~~ Responses are shown for the FAIR (blue) and AR5-IR (red) models. Panel b) shows the corresponding concentration response.

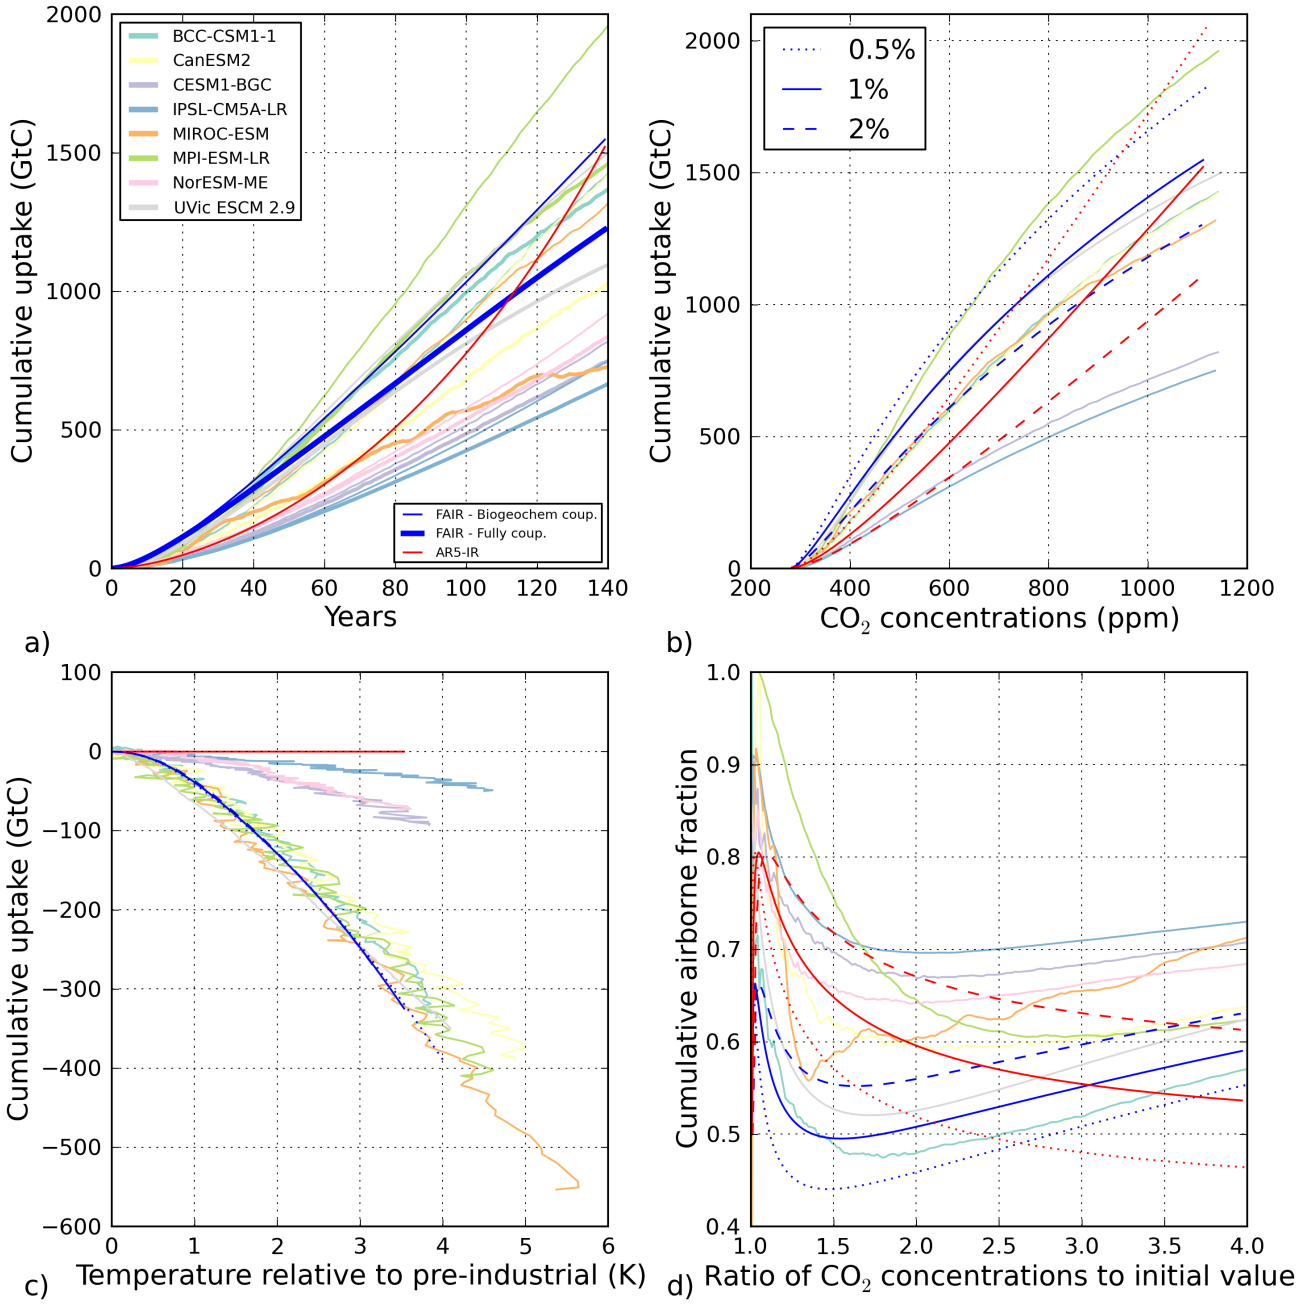

Figure 6. Response to idealised concentration increase experiments from Gregory et al. (2009) for the FAIR (blue) and AR5-IR (red) models. Light pastel colours show the ESMs from Joos et al. (2013) for the 1%/yr concentration increase scenario only. Panel a) shows the cumulative ~~total~~ ocean and land carbon ~~uptake~~ uptakee over time in the "fully coupled" 1%yr^{-1} concentration increase scenario. Panel b) shows the evolution of cumulative ~~total~~ ocean and land carbon ~~uptake~~ uptakee as a function of atmospheric concentration in the "biogeochemically coupled" experiment for 1%yr^{-1} (solid), 2%yr^{-1} (dashed) and 0.5%yr^{-1} (dotted) experiments. Panel c) shows the cumulative uptake as a function of temperature in the "radiatively coupled" experiment. Panel d) shows the evolution of the cumulative airborne fraction as a function of the proportional concentration increase for the "fully coupled" experiments.

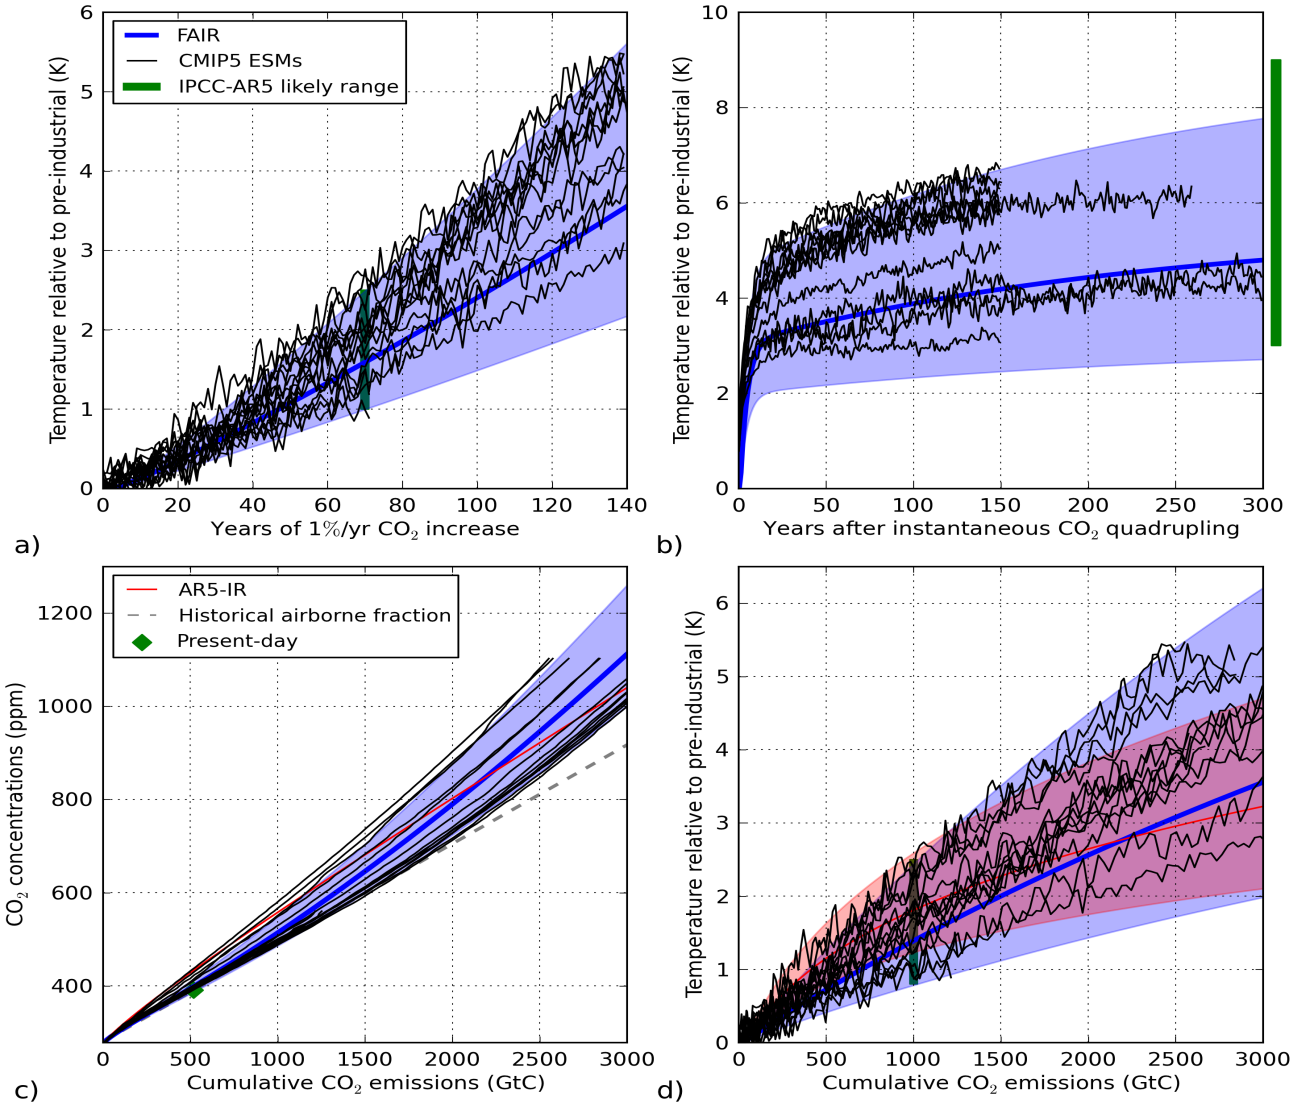

Figure 7. Climate response uncertainties in the FAIR (blue), AR5-IR (red) and CMIP5 (black) models. Panel a) shows the temperature responses to a 1%/yr concentration increase scenario. The ~~purple~~ green bar indicates the IPCC-AR5 TCR likely range. The blue shading in panels a) and b) shows the response of FAIR under IPCC-AR5 upper and lower likely TCR and ECS ranges. Panel b) shows the responses to an instantaneous quadrupling of atmospheric CO_2 which is held fixed subsequently. The ~~purple~~ green bar indicates the assessed equilibrium warming compatible with the IPCC-AR5 ECS likely range. Panel c) shows concentrations as a function of cumulative emissions in the 1%/yr scenario. ~~Upward curving lines indicate an increase cumulative airborne fraction.~~ The plumes in panels c) and d) show the response for the IPCC-AR5 likely TCR and ECS ranges, with an additional ~~+/- 10%~~ perturbation to the r_0, r_T and r_C parameters for the high/low end the likely ranges respectively in the FAIR model. The dashed grey line indicates a constant cumulative airborne fraction that is consistent with the present-day state ~~o~~ of the climate system (green diamond). Panel d) shows warming as a function of cumulative emissions in the 1%/yr scenario. ~~Straight lines indicate a constant TCRE.~~ The ~~purple~~ green bar shows the IPCC-AR5 likely 0.8-2.5K/TtC assessed range for TCRE.

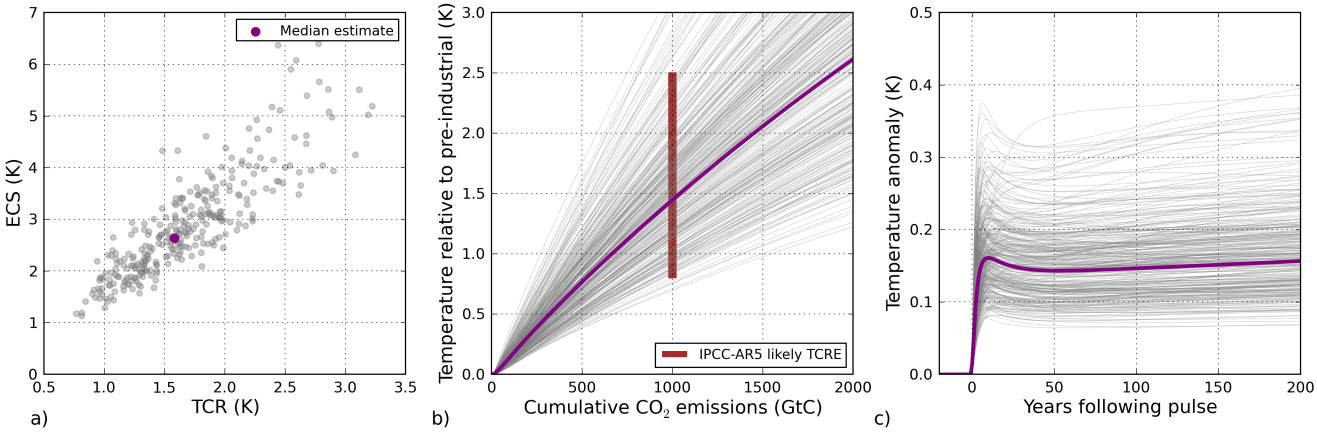

Figure 8. Probabilistic sampling in the FAIR model. Grey lines show 300 random draws from the input parameter distributions, as described in ~~the text~~section 3.4. Panel (a) shows the joint distribution of TCR and ECS. Panel (b), ~~the concentration response under MAGICC-derived RCP2.6 emissions. Panel (c),~~ warming as a function of cumulative emissions in the $1\%\text{yr}^{-1}$ concentration increase experiment. The brown bar in panel ~~e~~b) ~~represent~~ represents the IPCC-AR5 likely TCRE range. Panel (~~d~~c), the warming response to a 100GtC pulse emitted in 2020 on top of the MAGICC-derived RCP2.6 emissions. The purple ~~line/~~dot ~~represents~~ (a) and lines (b and c) represent the median ~~estimate in all panels~~parameters of the distributions.