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

Richard J. Millar[1,2], 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 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 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 CO_2 over the remainder of the century 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 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 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 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 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 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 correct representation of the temporal evolution of the warming response to the pulse emission is required from computationally-simple climate-carbon-cycle models. As simple climate-carbon-cycle models are not explicitly evaluated in terms of their pulse-response behaviour, it is unclear how well this robustly simulated physics is represented in such models.

A second important feature of more complex climate-carbon-cycle models is the increase in airborne fraction (the percentage of emitted CO_2 that remains in the atmosphere after a period of time) over time in scenarios involving substantial levels of emissions or warming (Friedlingstein et al., 2006; Millar et al., 2016). 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 (Rogelj 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 (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. Meinshausen et al. (2009)), the ability to reproduce the approximate linearity of the relationship between warming and cumulative emissions is a desirable property.

In this paper we show that the impulse-response functions that are provided for the calculation of multi-gas equivalence metrics in IPCC-AR5 (Myhre et al., 2013), a simple and easy to use coupled climate-carbon-cycle model, are insufficient to fully capture these emergent responses of the climate-carbon-cycle system. Such a state-insensitive 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 and mitigation scenarios as simulated by ESMs. 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, in section 3.1, why a state-dependence modification to the IPCC-AR5 carbon-cycle impulse-response function is required, motivating the modified model described in section 2. Section 3.2 then evaluates these models' ability in replicating 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 and how probabilistic assessments of climate response to CO_2 emissions could be made. Section 4 provides a concluding summary and discussion.

2 Model descriptions

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

The IPCC-AR5 proposed an idealised simple climate 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 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,

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

where E are annual CO_2 emissions, in units of ppm/year (1 ppm = 2.12GtC). Atmospheric CO_2 concentrations are given by $C = C_0 + \sum_i R_i$, and radiative forcing by:

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

where C_0 is the pre-industrial CO_2 concentration, $F_{2\mathrm{X}}$ the forcing due to CO_2 doubling, and F_{ext} the non-CO_2 forcing. GMST anomalies are computed thus:

$$\frac{\mathrm{d}T_j}{\mathrm{d}t} = \frac{c_j F - T_j}{d_j} \quad ; \quad T = \sum_j T_j \quad ; \quad j = 1,2 \tag{3}$$

with coefficients a_i, d_j and τ_i 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 Equilibrium Climate Sensitivity (ECS) =2.75K and Transient Climate Response (TCR) =1.6K (corresponding to $c_1 = 0.46$ and $c_2 = 0.27$ (Millar et al., 2015)), 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

Parameter	Value - AR5-IR	Value - PI-IR	Value - FAIR	Processes
a_0	0.2173	0.1266	0.2173	Geological re-absorption
a_1	0.2240	0.2607	0.2240	Deep ocean invasion / equilibration
a_2	0.2824	0.2909	0.2824	Biospheric uptake / ocean thermocline invasion
a_3	0.2763	0.3218	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	394.4	Deep ocean invasion/equilibration
τ_2 (yr)	36.54	31.61	36.54	Biospheric uptake / ocean thermocline invasion
τ_3 (yr)	4.304	4.240	4.304	Rapid biospheric uptake / ocean mixed-layer invasion
c_1	0.46	0.46	0.46	Thermal adjustment of upper ocean
c_2	0.27	0.27	0.27	Thermal equilibration of deep ocean
d_1 (yr)	8.4	8.4	8.4	Thermal adjustment of upper ocean
d_2 (yr)	409.5	409.5	409.5	Thermal equilibration of deep ocean
r_0 (yr)	-	-	35	Pre-industrial iIRF$_{100}$
r_C (yr/GtC)	-	-	0.02	Increase in iIRF$_{100}$ with cumulative carbon uptake
r_T (yr/K)	-	-	4.5	Increase in iIRF$_{100}$ with warming

Table 1. Default parameter values for 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.

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 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 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 the multi-model mean of the ensemble of ESMs and EMICs in Joos et al. (2013) (see table 1 for parameter values).

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

In the AR5-IR model the carbon-cycle constants are not 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 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 the carbon-cycle of the AR5-IR model, such that the CO_2 concentrations in the 4 "carbon reservoirs" are updated thus:

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

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), 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:

$$\mathrm{iIRF}_{100} = \sum_i \alpha a_i \tau_i \left[1 - \exp\left(\frac{-100}{\alpha\tau_i}\right)\right]. \tag{5}$$

Following other simplified carbon-cycle models (Meinshausen et al., 2011a; Glotter et al., 2014), 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 anomaly from pre-industrial conditions, T. A simple linear relationship appears to give an adequate approximation to the behaviour of ESMs and EMICs (as will be shown subsequently in section 3):

$$\mathrm{iIRF}_{100} = r_0 + r_C C_{\mathrm{acc}} + r_T T. \tag{6}$$

Values of r_0=35 years, r_C =0.02 years/GtC, recalling that 2.12 GtC = 1ppm, and r_T =4.5 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 approximately replicate the relationship between warming-driven outgassing of carbon in the bulk of CMIP5 ESMs (see section 3.3). The values of r_0, r_T and r_C given here are intended to be taken only as approximate best-estimate values that capture important carbon-cycle dynamics in ESMs. The exact values of these parameters could be tuned (along with the other parameters in the model) to best-reproduce the aspect of ESM/EMIC behaviour of interest (e.g. see Figure 4).

We compute iIRF_{100} at each time-step using C_{acc} and T from the previous time-step and equation 6, convert to a α using equation 5 and apply to the carbon-cycle equations (equation 4). 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 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 years,

corresponding with these parameters to α=65.4. 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.

3 Results

In this section we initially set out the need for the FAIR model by showing that state-independent impulse-response model cannot simultaneously reproduce the observed carbon-cycle response over the historical period and the future carbon-cycle evolution as projected by ESMs under possible future emissions scenarios (section 3.1). We subsequently evaluate the ability 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).

3.1 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 present-day state of the climate system successfully. Compatibility with 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 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 a bias of over 30ppm in 2011 concentrations, due to the slower than observed decay of 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 1c). The PI-IR model maintains a lower instantaneous airborne fraction than 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 amplitude variations in the 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 to be reproduced by any of these simple 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 AR5-IR and PI-IR models.

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 from the MAGICC model (Meinshausen et al., 2011b) in order to allow a comparison of both concentrations and temperatures between simple models. 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 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 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 after 2100 in the high emission scenario, but the behaviour of MAGICC (or indeed any other model) under these more extreme forcing scenarios has not been verified. 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 (100GtC pulse in 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 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 of its initial value, slightly greater than the multi-model average of the ESM responses of 0.41, but, the iIRF_{100} of 53 years is consistent with the ESM multi-model mean of 52.4 years (Joos et al., 2013). Excluding temperature feedbacks (the "biogeochemically-coupled' version - setting $r_T = 0$) on the carbon-cycle increases the decay of the temperature response to the pulse over the century following the pulse emission which reduces the iIRF_{100} airborne fraction by 11%. The "fully-coupled" FAIR model shows temperature initially adjusting rapidly followed by near-constant temperature over the remainder of the century.

Figure 3c and 3d also 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 100GtC pre-industrial pulse decays

faster than the present-day case, due to reduced saturation of the land and ocean carbon sinks. With these parameters, the FAIR $iIRF_{100}$ is approximately 30% lower in the pre-industrial case compared to the present day, consistent with corresponding ratio in the Joos et al. (2013) ensemble, with its value of 36 years within the 34-47 years range of the ESMs. The magnitude of 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% increase of $iIRF_{100}$ in the 5000GtC pre-industrial pulse relative to the 100GtC pre-industrial, whilst smaller than the approximate doubling observed in the ESMs, shows that the FAIR model 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.

A difference between the FAIR model and the ESMs is that 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 3c), whereas a 13% reduction in $iIRF_{100}$ is observed for the ESMs (Joos et al., 2013) (not shown). It is only for the 5000GtC pre-industrial pulse experiment that we see a reduction in the $iIRF_{100}$ associated with suppression of the temperature-induced feedbacks on the carbon cycle in FAIR.

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 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, r_T and r_C parameters to vary from the model parameters found in the first stage, whilst again maintaining the same 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 ESM responses whilst at the same time maintaining the dependency on background condition and pulse size for the specific model in question.

As an final test of the FAIR model's sensitivity to pulse size, we also consider the response of the FAIR and 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 (as derived from historical concentrations using the MAGICC model) together with estimates of the

historical radiative forcing from non-CO_2 factors. 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 (Riahi et al., 2011) trajectories for 2005-2008.

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 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 in the FAIR model fails to plateau as quickly as at lower pulses, where 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 $0.5\%\text{yr}^{-1}$, $1\%\text{yr}^{-1}$ and $2\%\text{yr}^{-1}$ respectively and consistent emissions were derived 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 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 the future.

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. 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 between temperature changes and the carbon-cycle in the FAIR model acts to suppress carbon uptake, shown by the difference between the thick and thin lines in Figure 6a, 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 6b shows C_{acc} as a function of atmospheric CO_2 concentration: again, the FAIR model captures the concave-downward form of this diagnostic, in contrast to the AR5-IR model.

Whilst oceanic carbon-cycle feedbacks are almost exclusively driven by biogeochemical effects (Glotter et al., 2014), for simple climate-carbon-cycle models to be of use in representing the entire climate system, they need to capture dependencies of the land carbon cycle on warming. Aside from 3 ESMs that display global-mean carbon-cycles insensitive to warming, Figure 6c shows a coherent relationship between temperature 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/K allow the FAIR model to reproduce this relationship well. $1\%\text{yr}^{-1}$, $0.5\%\text{yr}^{-1}$ and $2\%\text{yr}^{-1}$ experiments all lie along the same line in panel (c), 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) followed by a subsequent increase (Figure 6d) displayed by the FAIR model 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 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 iIRF100). As more recent emissions (which increase monotonically under the 1%/yr scenario) have a higher instantaneous airborne fraction, the initial decrease in cumulative airborne fraction stops and then begins to increase as this accelerating saturation becomes the dominant effect.

3.4 Uncertainty and 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 offers a convenient structure for simply sampling plausible ranges of TCR and ECS, as 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 a) and b) in figure 7 show how 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.

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 7c imply that a rising fraction of cumulative emissions remain resident in the atmosphere). 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 perturbations of +/-13% (approximately equivalent to a present-day iIRF$_{100}$ change of +/- 7 years) to the r_0 ,r_T and r_C parameters (combined with perturbations to c_1 and c_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 the range of responses seen in the CMIP5 models under a $1\%\text{yr}^{-1}$ concentration increase scenario.

Crucially, the FAIR model also captures the straight-line relationship between cumulative carbon emissions and human-induced warming (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, has a Transient Response to Cumulative Emissions (TCRE) =1.5K/TtC (thick blue line in Figure 7d). 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 concave-downward shape in a plot of realised warming against cumulative carbon emissions, because 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 at high cumulative emissions, consistent with the behaviour of ESMs (Leduc et al., 2015).

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 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 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 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 emergent 5-95% range (based on 300 draws from the input parameter distributions) for TCRE (figure 8c) of 1.0-2.5K/TtC is 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 in 2020 against the background of the RCP2.6 scenario (figure 8d). 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 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. 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 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.

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's response to cumulative emissions, whilst at the same time 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 allows the range of ESM behaviour to be emulated whilst maintaining the physically-understood dependency of 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.

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.

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.

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.

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.

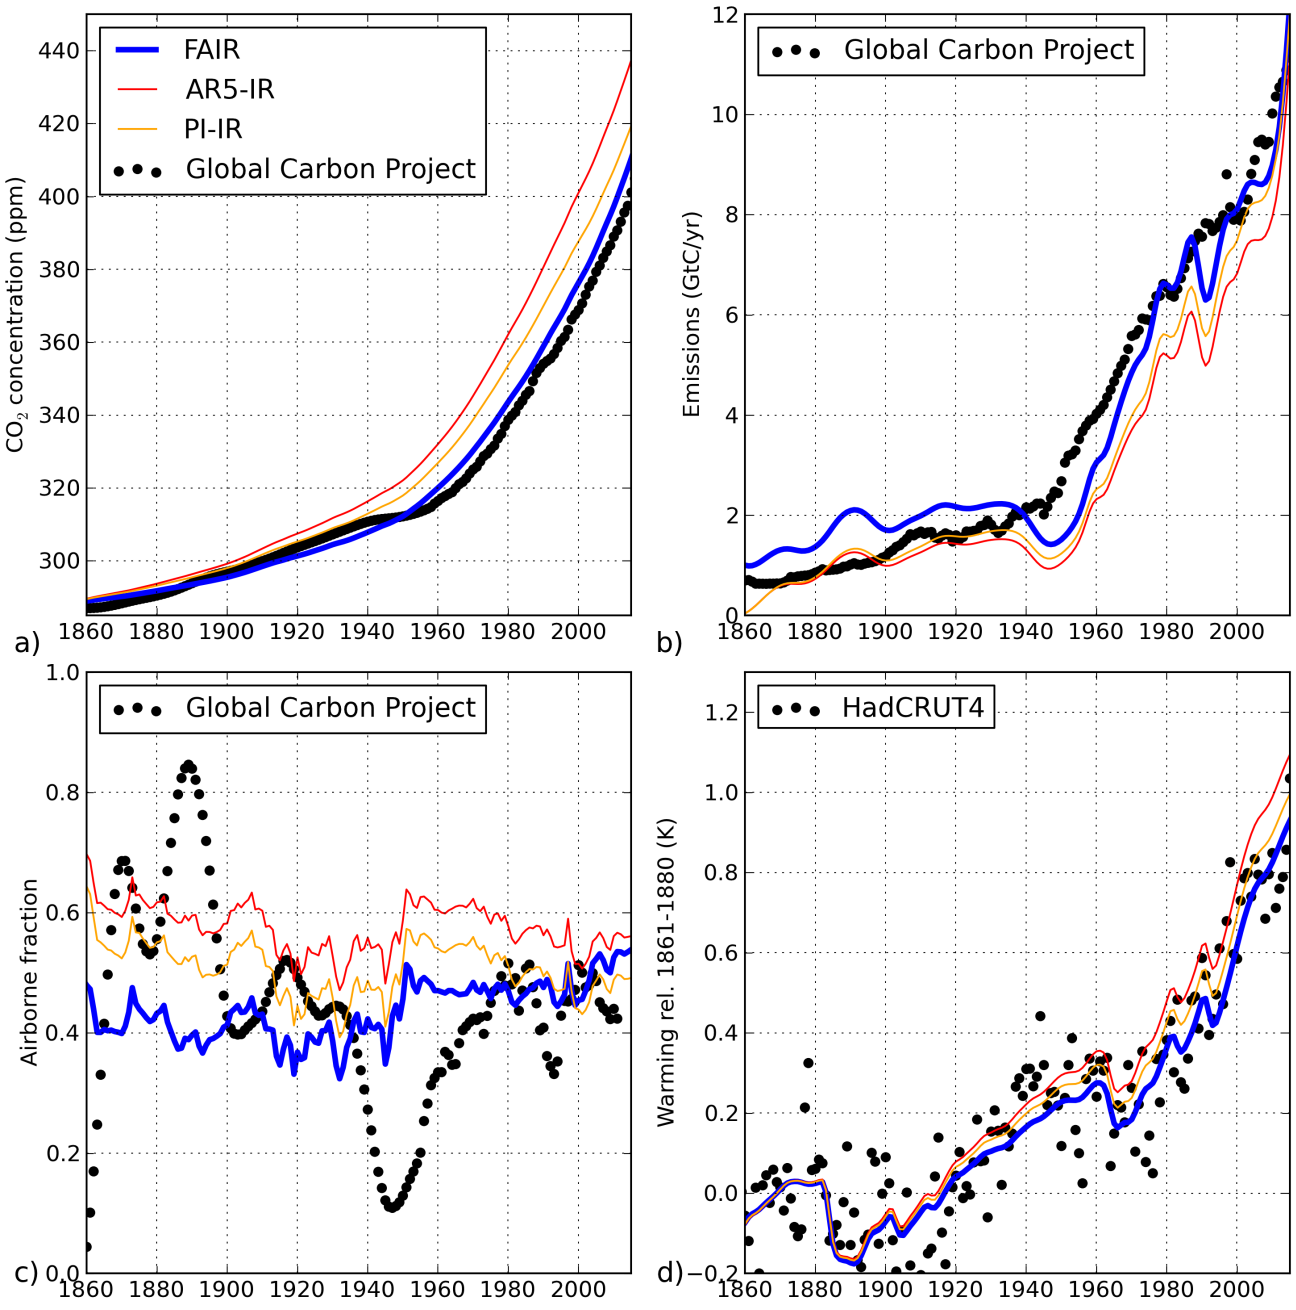

Figure 1. Historical validation of the FAIR (blue), AR5-IR (red) and PI-IR (orange) models. Panel a) shows the CO_2 concentration response when integrated under historical emissions (and historical non-CO_2 radiative forcing for the RCP scenarios). Panel b) shows the derived CO_2 emissions consistent 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 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. All simulations are commenced from assumed quasi-equilibrium carbon-cycle states in 1850.

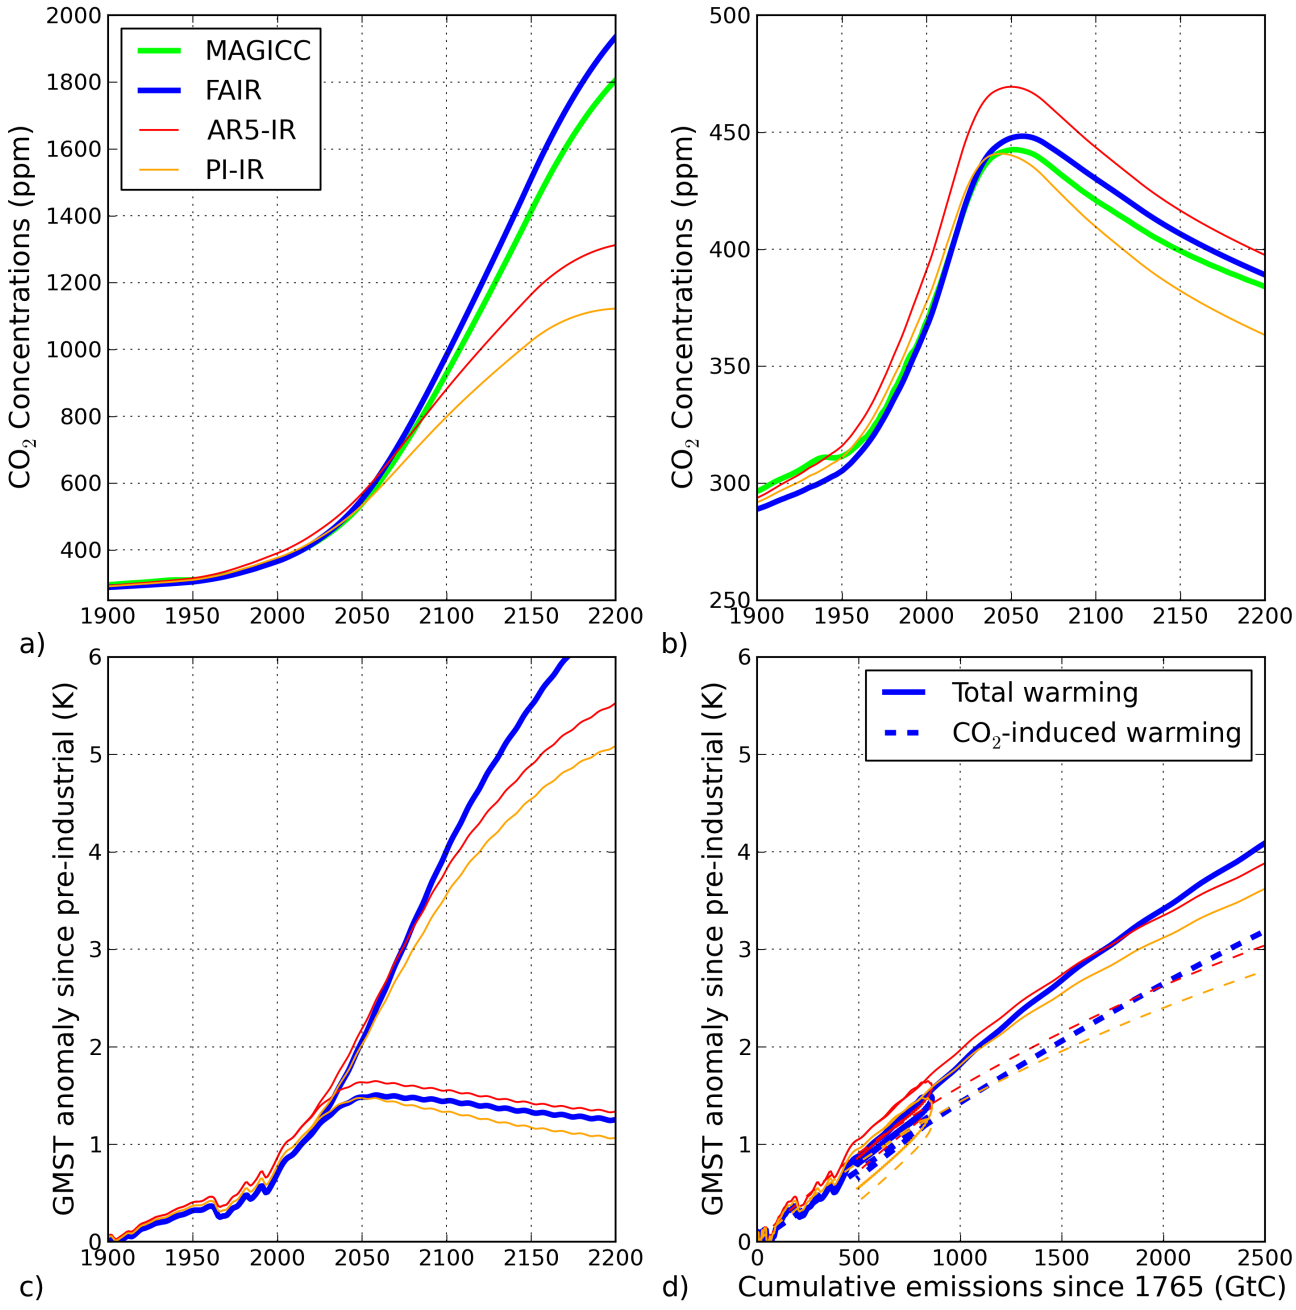

Figure 2. Panels a) and b) shows the CO_2 concentrations under RCP8.5 and RCP2.6 respectively for the FAIR (blue), AR5-IR (red), PI-IR (orange) and MAGICC (green) models. Panel c) shows the temperature response under both RCP2.6 and RCP8.5. Panel d) shows the evolution of total warming (full) and CO_2-induced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

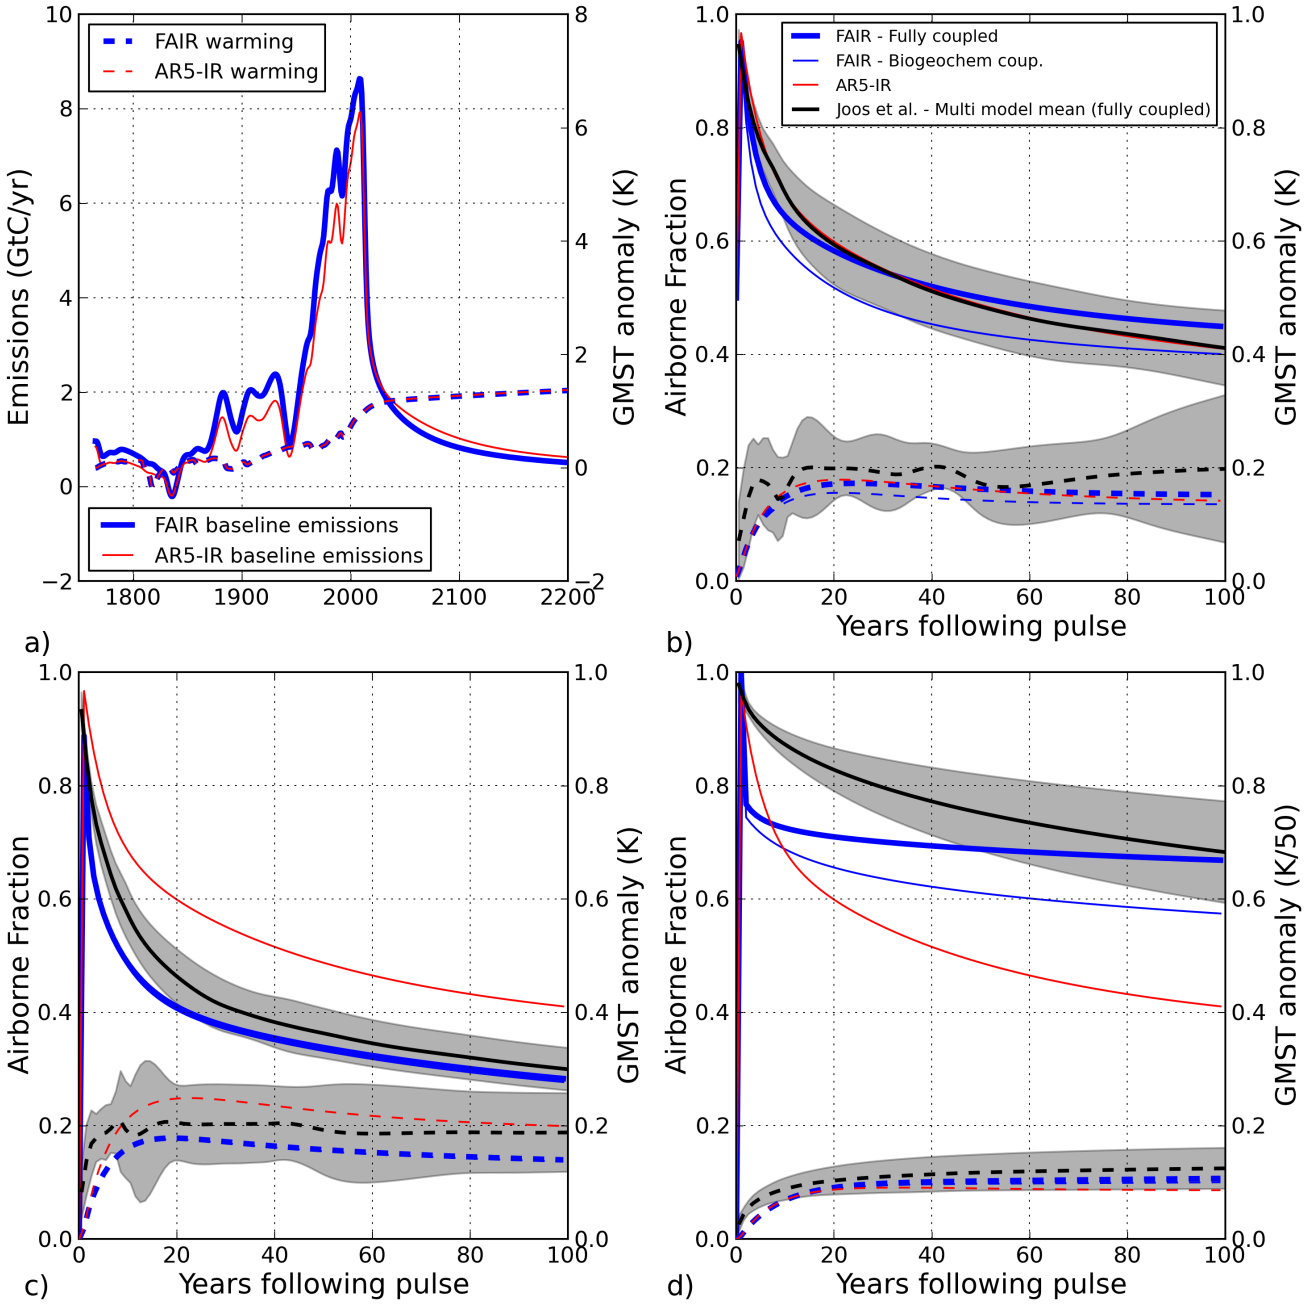

Figure 3. Response to pulse emission experiments of Joos et al. (2013). Panel a) shows the "baseline" emissions (left-hand axis, solid) and warming (right-hand axis, dashed) when concentrations are stabilised at 389ppm for the FAIR (blue) and AR5-IR (red) models. Panel b) shows the response to a 100GtC imposed on present-day (389ppm) background conditions (PD100 experiment). Panel c) shows the response to a 100GtC pulse in pre-industrial conditions (PI100 experiment). Panel d) 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) and d). The black lines in panels b), c) and d) shows the Joos et al. (2013) multi-model mean for airborne fraction (solid) and warming (dashed), with the black shading indicating one standard deviation uncertainty across the ensemble.

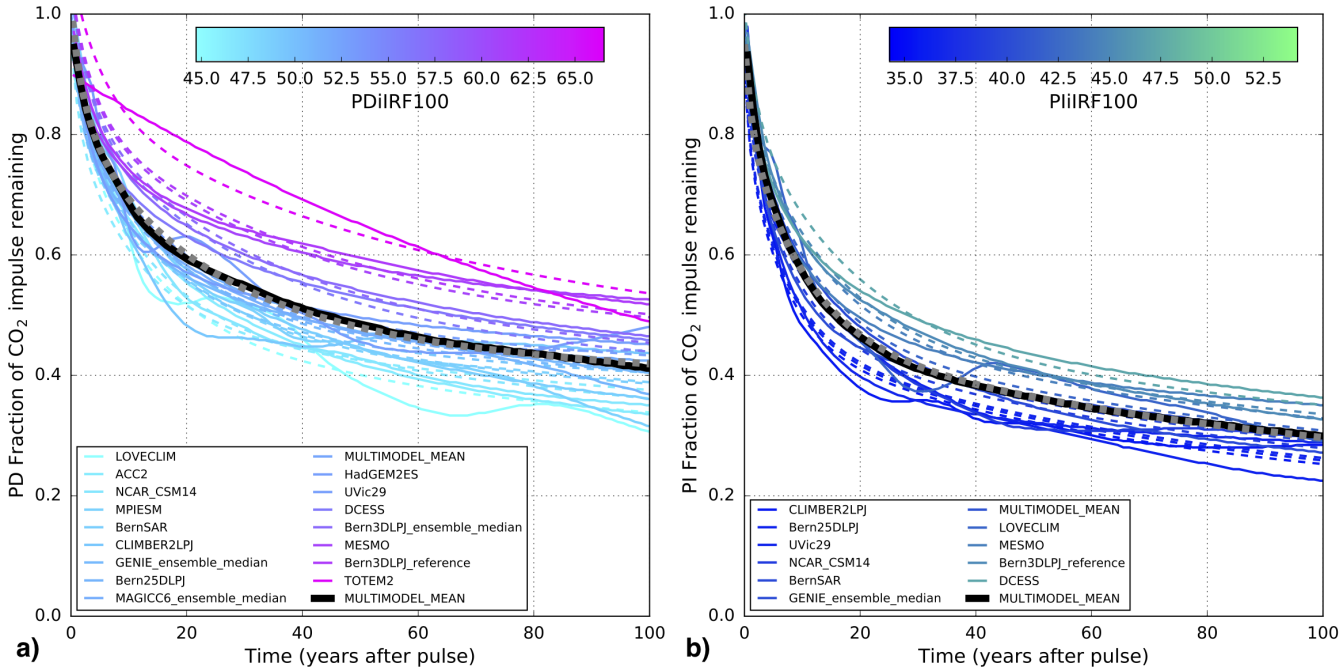

Figure 4. Fitting individual models from Joos et al. (2013) with FAIR. Panel a) shows the remaining airborne fraction for the PD100 experiment and panel b) for those models that additionally completed the PI100 experiment. Solid lines show the original model response coloured by the $iIRF_{100}$ values. 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.

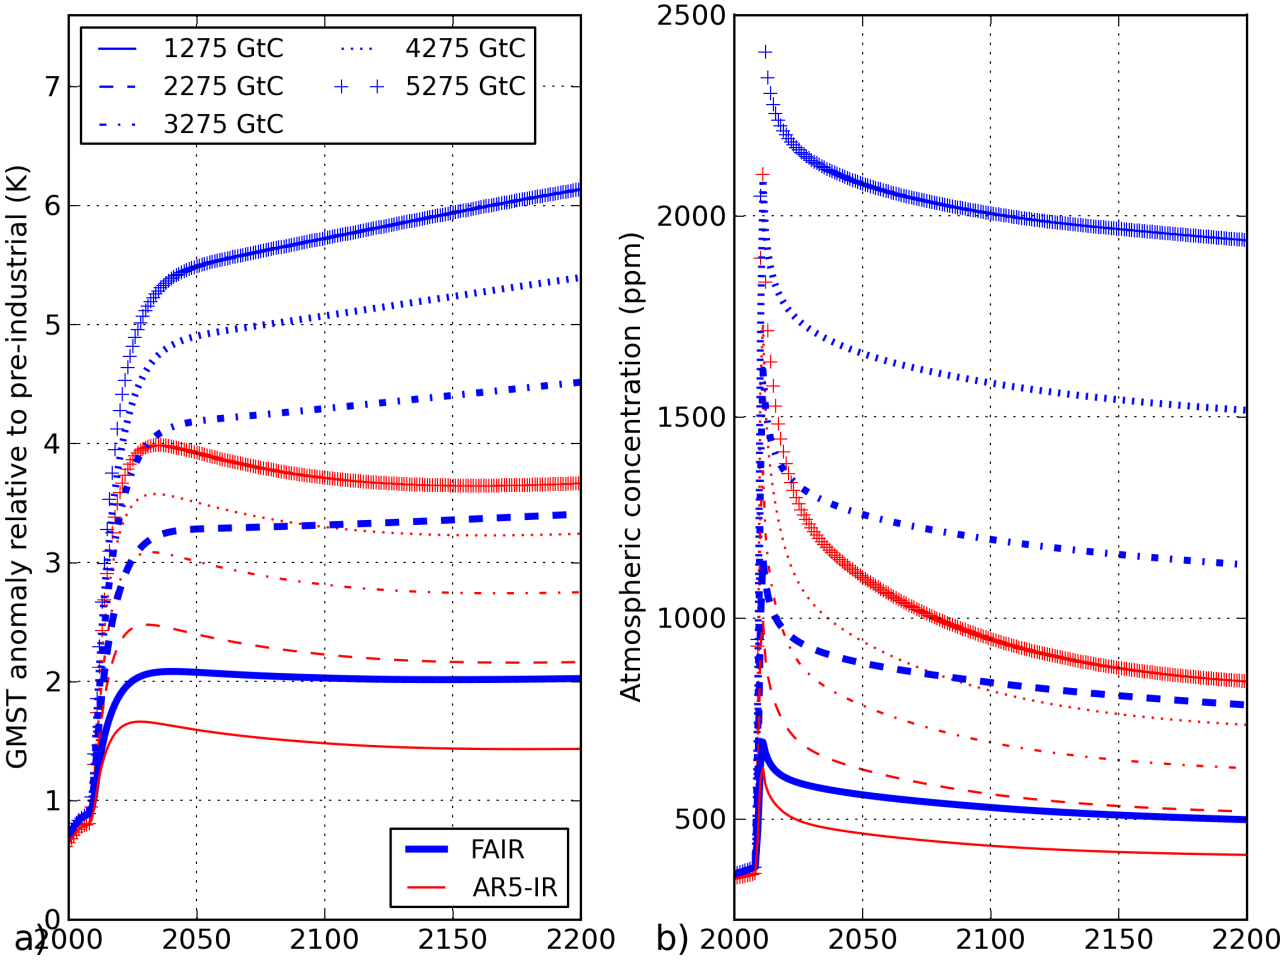

Figure 5. Panel a) shows the global mean surface temperature (GMST) response to the pulse experiments of Herrington and Zickfeld (2014). Pulse emissions are applied over a 2-year period from 2008, with differing total cumulative carbon emissions denoted by different line styles. Responses are shown for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

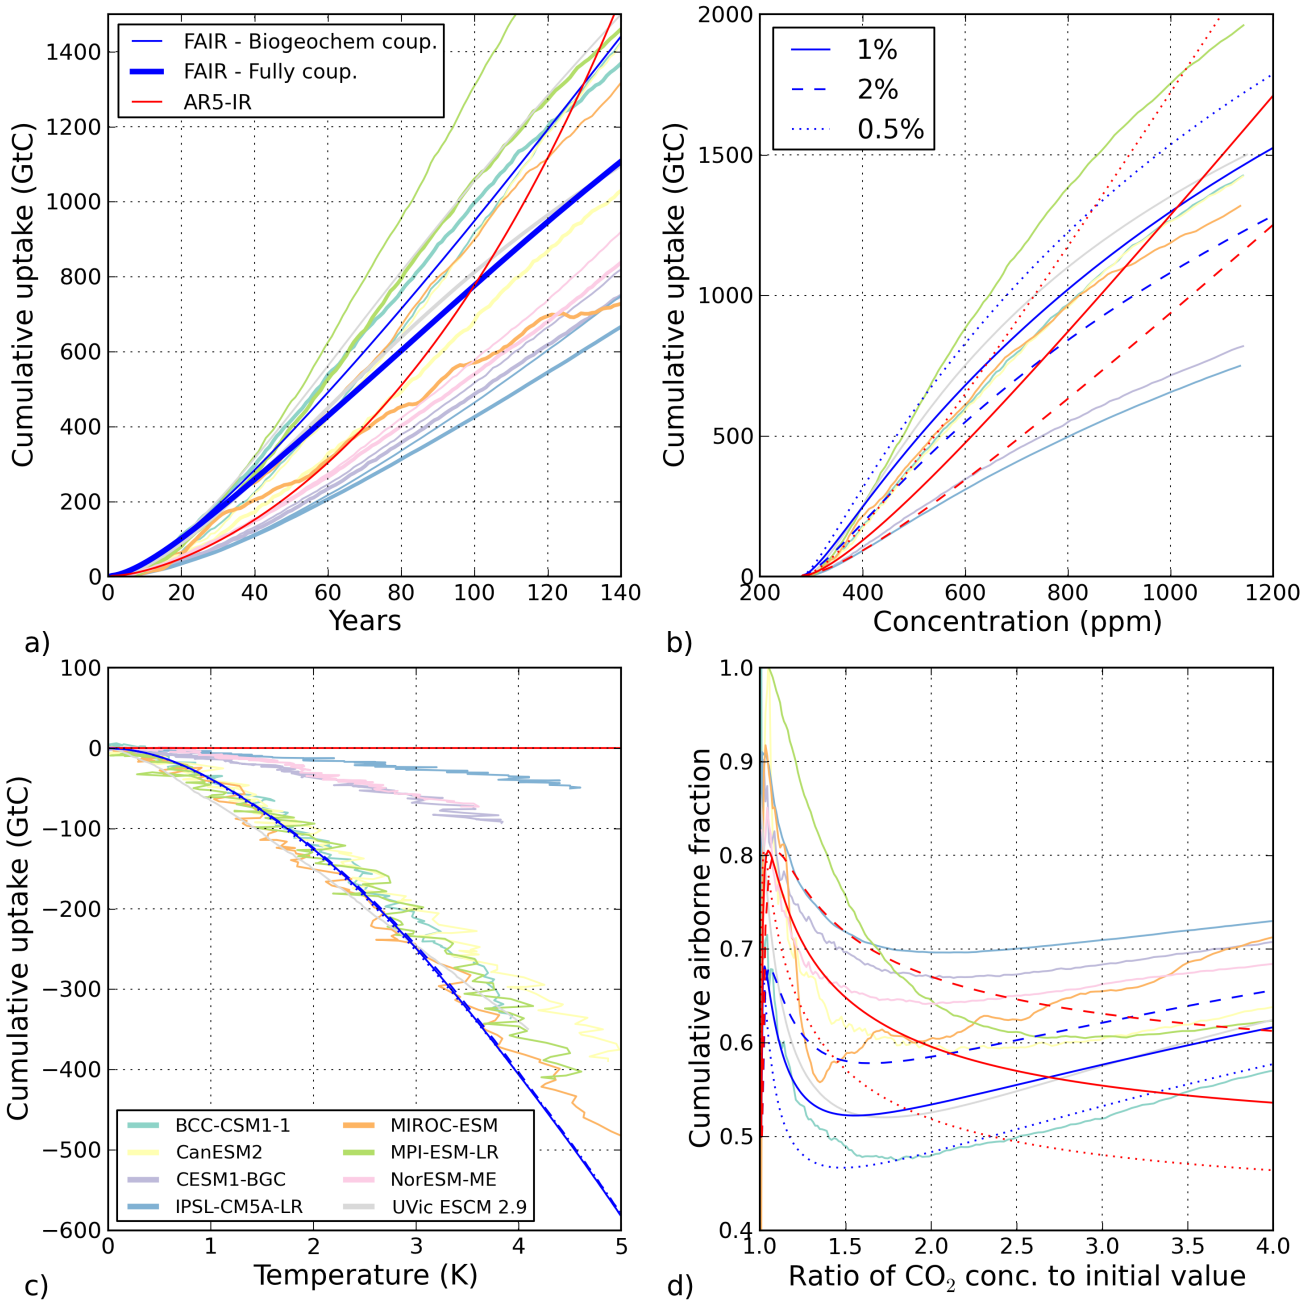

Figure 6. Response to idealised concentration increase experiments from Gregory et al. (2009) for the FAIR (blue) and AR5-IR (red) models. Light pastel colours show the ESMs from Joos et al. (2013) for the 1%/yr concentration increase scenario only. Panel a) shows the cumulative total carbon uptake over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

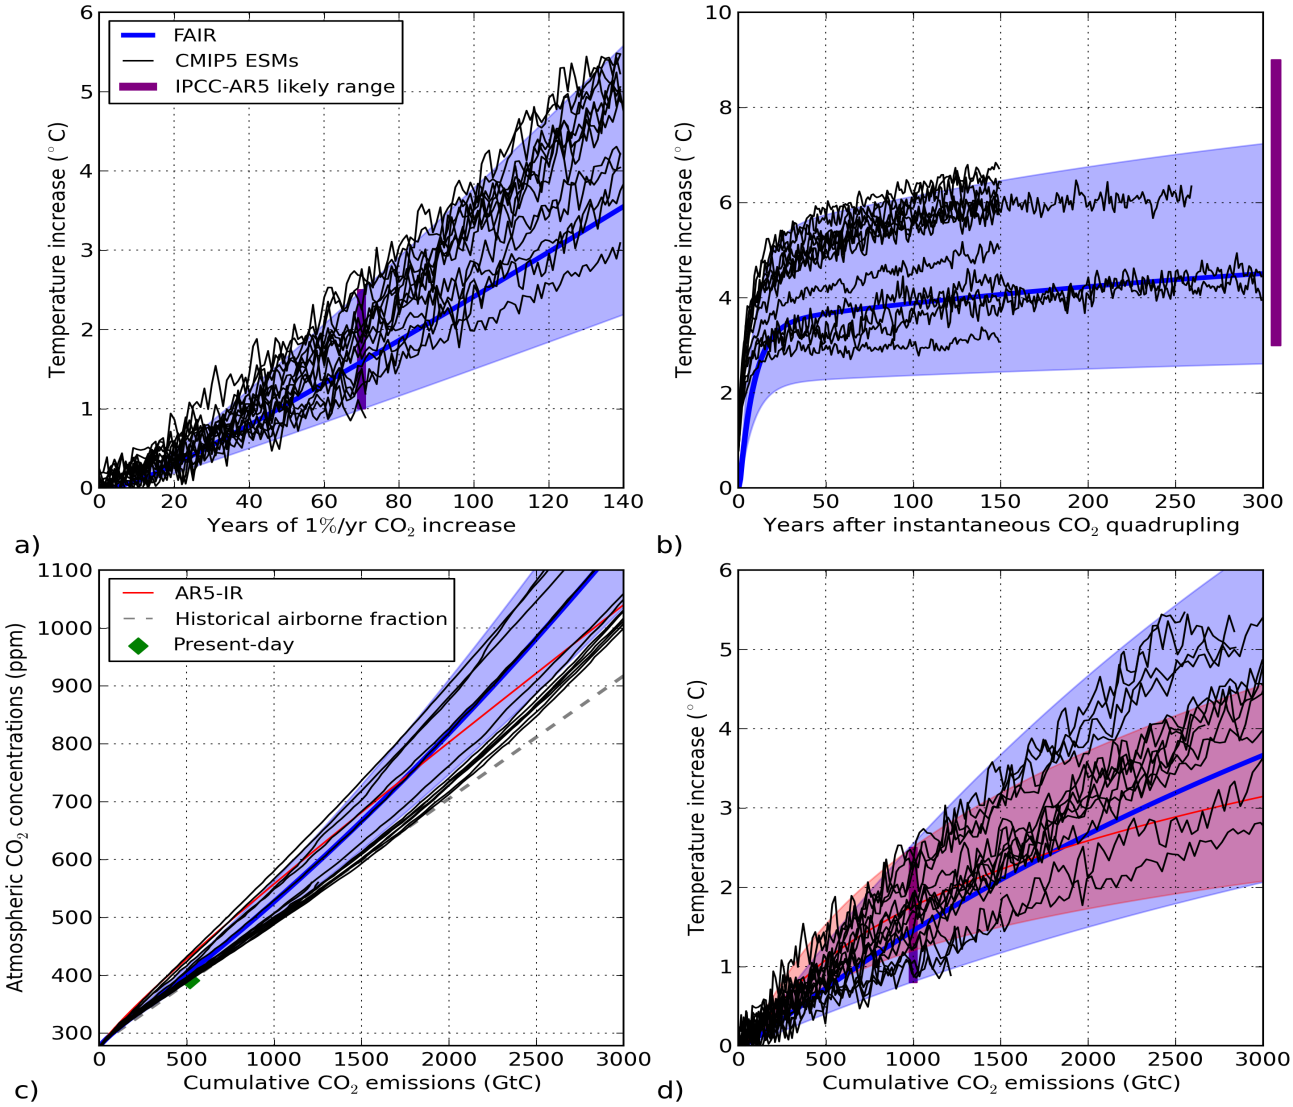

Figure 7. Climate response uncertainties in the FAIR (blue), AR5-IR (red) and CMIP5 (black) models. Panel a) shows the temperature responses to a 1%/yr concentration increase scenario. The purple bar indicates the IPCC-AR5 TCR likely range. The blue shading in panels a) and b) shows the response of FAIR 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 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 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 bar shows the IPCC-AR5 likely 0.8-2.5K/TtC assessed range for TCRE.

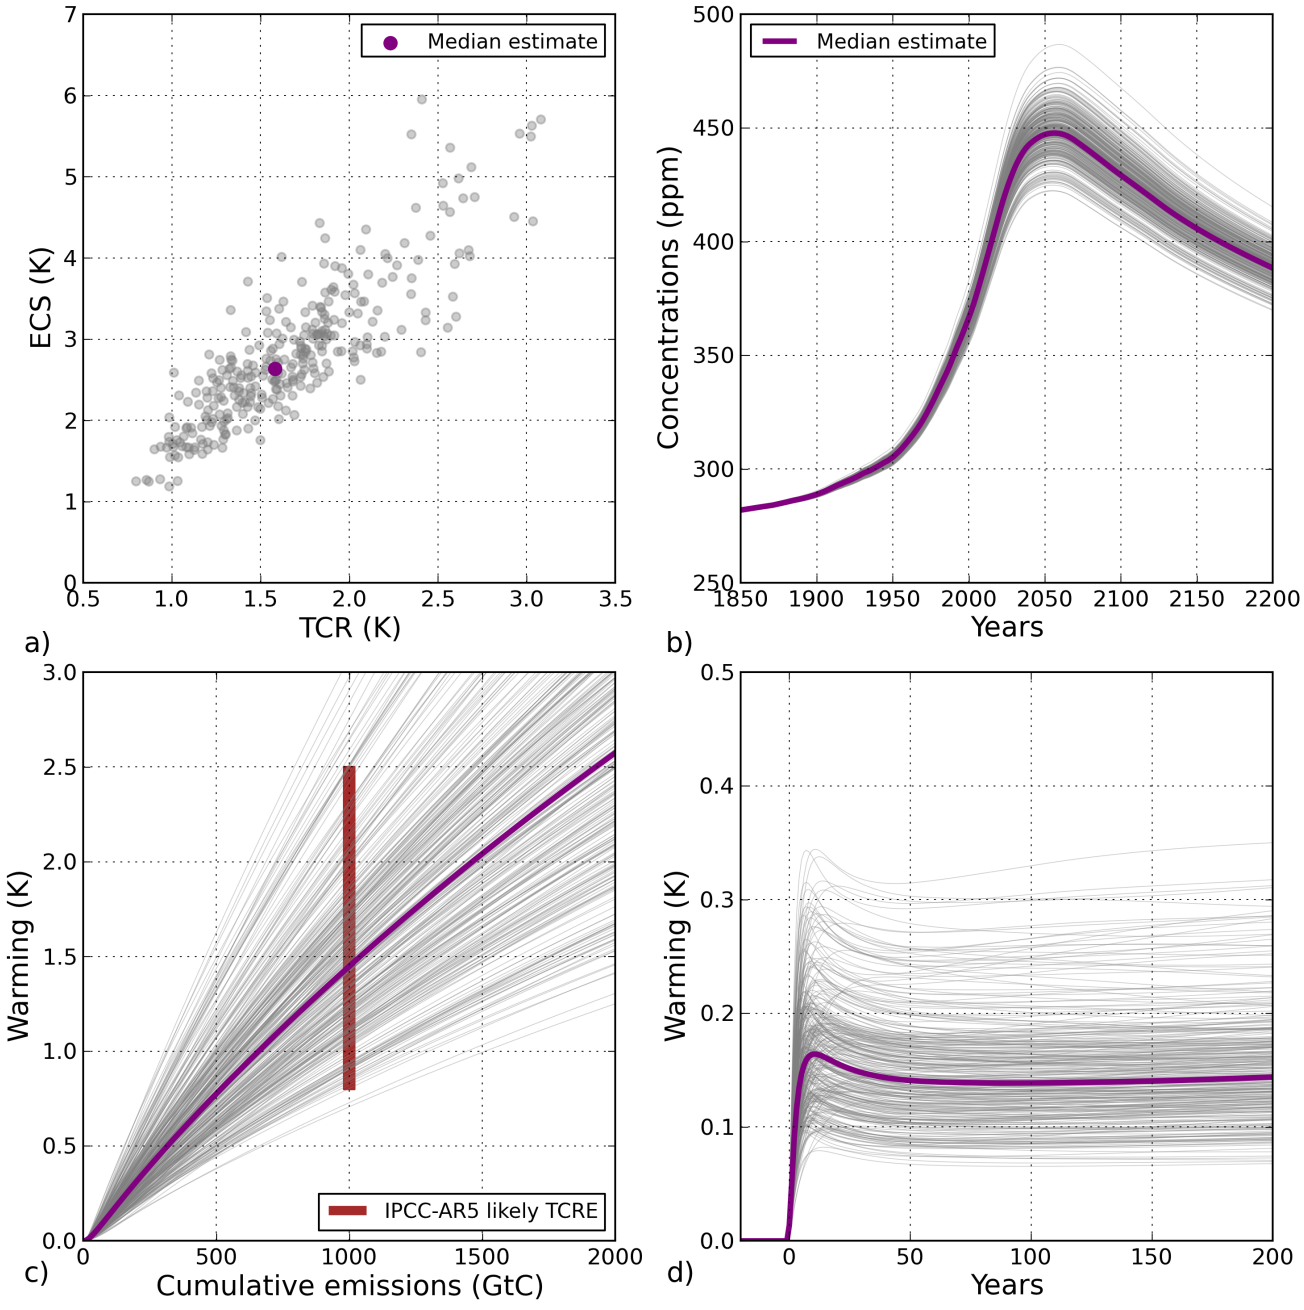

Figure 8. Probabilistic sampling in the FAIR model. Grey lines show 300 random draws from the input parameter distributions, as described in the text. Panel (a) shows the joint distribution of TCR and ECS. Panel (b), the concentration response under MAGICC-derived RCP2.6 emissions. Panel (c), warming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