

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

Richard J. Millar^{1,2,3}, Zebedee. R. Nicholls¹, Pierre Friedlingstein³, and Myles R. Allen^{1,2,4}

¹Department of Physics, University of Oxford, Oxford, UK

²Oxford Martin Net Zero Carbon Investment Initiative, Oxford Martin School, University of Oxford, Oxford, UK

³Department of Mathematics, University of Exeter, Exeter, UK

⁴Environmental 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 (CO₂) 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 CO₂ airborne fraction. 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 CO₂ on a range of timescales and under a range of experimental designs.

1 Introduction

In the long term, future climate changes will largely be determined by future cumulative CO₂ 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). Linking specific CO₂ 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 CO₂ 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, 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 CO₂ concentrations and global mean temperature under both idealised settings and under possible projections of future emissions scenarios. Following a

CO₂ pulse emission of 100GtC in present-day climate conditions, ESMs (and Earth System Models of Intermediate Complexity – EMICs) display a rapid-draw down of CO₂ 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
5 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 response to a pulse emission of CO₂ characterised by a rapid warming, over approximately a decade, to a plateau value of global mean surface temperature anomaly (Joos et al., 2013). Warming does not noticeably decrease from this value
10 over the following several hundred years, indicating that, short of artificial CO₂ removal (CDR) or active solar geoengineering, CO₂-induced warming is essentially permanent on human-relevant timescales.

The correct representation of the temporal evolution of the warming response to a pulse emission is required for computationally-simple climate-carbon-cycle models. Aside from the simple climate-carbon-cycle models 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
15 been evaluated in terms of their pulse-response behaviour, and it remains unclear how well 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
20 background state that have been highlighted in ESMs and EMICs (Joos et al., 2013; Herrington and Zickfeld, 2014).

A second important feature of ESMs is the increase in airborne fraction (the fraction of emitted CO₂ that remains in the atmosphere after a specified period) with time in scenarios involving substantial emissions or warming (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 CO₂ concentrations and radiative forcing, yielding an approximately linear
25 relationship between cumulative CO₂ emissions and CO₂-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 (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
30 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
35 climate-carbon-cycle system to emissions of CO₂ (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).

5 In this paper we show that although the impulse-response functions provided for the calculation of multi-gas equivalence metrics in IPCC-AR5 (Myhre et al., 2013) provide a simple and easy to use climate-carbon-cycle model, this model is insufficient to fully capture the emergent responses of the coupled climate-carbon-cycle system. Such a 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 21st century to both high-emissions and mitigation scenarios
10 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.

15 Section 2 describes the formalism of the models that we contrast throughout this paper 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 ‘Finite Amplitude Impulse-Response Model’ (FAIR) described in section 2. Section 3.2 then evaluates the ability of FAIR and the unmodified IPCC-AR5 impulse-response
20 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 CO₂ concentrations are increased by a fixed percentage each year starting from pre-industrial values. Section 3.4 discusses uncertainty in FAIR and how probabilistic assessments of climate response to CO₂ emissions could be made using the model. Section 4 provides a concluding summary and discussion.

25 **2 Model description and methods**

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

The IPCC-AR5 proposed an idealised simple 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 CO₂ concentration response to emissions (Myhre et al., 2013). This model represents the evolution of atmospheric CO₂
30 by partitioning emissions of anthropogenic CO₂ 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. Four carbon pools are determined to be sufficient to empirically represent the response of atmospheric CO₂ 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,

$$\frac{dR_i}{dt} = a_i E - \frac{R_i}{\tau_i} \quad ; \quad i = 1 - 4 \quad (1)$$

- 5 where E is the annual CO₂ 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 CO₂ concentrations are given by $C = C_0 + \sum_i R_i$, and radiative forcing by:

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

- where C_0 is the pre-industrial CO₂ concentration, F_{2X} the forcing due to CO₂ doubling ($F_{2X}=3.74\text{Wm}^{-2}$), and F_{ext} the non-CO₂ forcing. Global mean surface temperature anomalies (T) are computed thus:

$$\frac{dT_j}{dt} = \frac{q_j F - T_j}{d_j} \quad ; \quad T = \sum_j T_j \quad ; \quad j = 1, 2. \quad (3)$$

- Coefficients a_i , and τ_i are as given in AR5 Chapter 8, tables 8.SM.9 and 8.SM.10 (Myhre et al., 2013), 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 CO₂ concentrations and a 1%/yr increase in CO₂ concentrations, q_j can be related to the Equilibrium Climate Sensitivity (ECS¹) and Transient Climate Response (TCR²) via the expressions,

$$\text{ECS} = F_{2X}(q_1 + q_2), \quad (4)$$

and

$$25 \quad \text{TCR} = F_{2X} \left(q_1 \left(1 - \frac{d_1}{70} \left(1 - \exp\left[-\frac{70}{d_1}\right] \right) \right) + q_2 \left(1 - \frac{d_2}{70} \left(1 - \exp\left[-\frac{70}{d_2}\right] \right) \right) \right), \quad (5)$$

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 and ECS=2.75K ($q_1 = 0.33\text{KW}^{-1}\text{m}^2$ and $q_2 = 0.41\text{KW}^{-1}\text{m}^2$), indicative of a typical mid-range climate response to radiative forcing in ESMs (Flato et al., 2013).

- 25 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 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

¹The global mean warming resulting from an instantaneous doubling of pre-industrial CO₂ concentrations after allowing the climate system to reach a new equilibrium state.

²The global mean warming at the time of doubled CO₂ concentrations following a 1%/yr increase from pre-industrial values

Parameter	Value - AR5-IR	Value - PI-IR	Value - FAIR	Guiding analogues
a_0	0.2173	0.1545	0.2173	Geological re-absorption
a_1	0.2240	0.1924	0.2240	Deep ocean invasion / equilibration
a_2	0.2824	0.2424	0.2824	Biospheric uptake / ocean thermocline invasion
a_3	0.2763	0.4108	0.2763	Rapid biospheric uptake / ocean mixed-layer invasion
τ_0 (yr)	1×10^6	1×10^6	1×10^6	Geological re-absorption
τ_1 (yr)	394.4	276.7	394.4	Deep ocean invasion/equilibration
τ_2 (yr)	36.54	30.75	36.54	Biospheric uptake / ocean thermocline invasion
τ_3 (yr)	4.304	4.459	4.304	Rapid biospheric uptake / ocean mixed-layer invasion
q_1 (KW ⁻¹ m ²)	0.33	0.33	0.33	Thermal equilibration of deep ocean
q_2 (KW ⁻¹ m ²)	0.41	0.41	0.41	Thermal adjustment of upper ocean
d_1 (yr)	239.0	239.0	239.0	Thermal equilibration of deep ocean
d_2 (yr)	4.1	4.1	4.1	Thermal adjustment of upper ocean
r_0 (yr)	-	-	32.40	Pre-industrial iIRF ₁₀₀
r_C (yr/GtC)	-	-	0.019	Increase in iIRF ₁₀₀ with cumulative carbon uptake
r_T (yr/K)	-	-	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.

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. These parameters are derived from a fit to the multi-model mean of the ensemble of ESMs and EMICs from Joos et al. (2013) (see table 1 for parameter values).

5 2.2 A Finite Amplitude Impulse Response (FAIR) model

In the AR5-IR and PI-IR models the carbon-cycle response to a pulse emission is not explicitly affected by rising temperature or CO₂ accumulation and hence only represent the specific response to a particular perturbation scenario. In more comprehensive models, ocean uptake efficiency declines with accumulated CO₂ 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).

10 In an attempt to capture some of these dynamics within the simple impulse-response model structure, we attempt a minimal modification of the AR5-IR model to allow it to mimic the behaviour of ESMs/EMICs in response to finite-amplitude CO₂ 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 CO₂ concentrations in the 4 “carbon reservoirs” are updated thus:

$$15 \quad \frac{dR_i}{dt} = a_i E - \frac{R_i}{\alpha \tau_i} \quad ; \quad i = 1 - 4 \quad (6)$$

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, 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). With other coefficients fixed, $iIRF_{100}$ is a monotonic (but non-linear) function of α :

$$iIRF_{100} = \sum_i \alpha a_i \tau_i \left[1 - \exp\left(\frac{-100}{\alpha \tau_i}\right) \right]. \quad (7)$$

Equation 7 is derived using the approximation that α is independent of time, and hence equation 7 is only exactly equivalent to 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), $C_{acc} = \sum_t E - (C - C_0)$, and of global mean temperature 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):

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

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 $r_0=32.4$ years, $r_C=0.019$ years/GtC, and $r_T=4.1$ years/K, with $ECS=2.75K$ and $TCR=1.6K$, are here used as model default parameters³. 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 given here as default parameters are intended to be taken only as approximate CMIP5-representative values that capture important carbon-cycle dynamics in ESMs. These values have not been explicitly optimised to any particular goal and can be tuned (along with the other model parameters) to reproduce specific aspects of individual ESM/EMIC behaviour (e.g. see 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.

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 96.6 years, corresponding, with the parameters as given in table 1, to $\alpha=100$. 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.

³The 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.

2.3 Experimental set-up

In this section we 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.

5 2.3.1 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 CO₂ 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 CO₂ concentrations reach 389ppm and the resulting CO₂ 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 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 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 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 CO₂ 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 2.3.2) 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 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 CO₂ emissions together with estimates of the historical non-CO₂ radiative forcing, both backed-out from historical concentrations using the MAGICC model (Meinshausen et al., 2011). 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-CO₂ forcings are held constant at 2008 levels after following RCP8.5 trajectories for 2005-2008.

2.3.2 Exponential CO₂ 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 CO₂ 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 CO₂ 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 CO₂ but the carbon-cycle is only allowed to respond to the simulated warming and not to increasing CO₂; and a “fully-coupled” experiment in which the carbon-cycle is allowed to respond to both warming and increased CO₂. 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 CO₂ 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.

2.3.3 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 CO₂ 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 CO₂ in 2020 (against a background RCP2.6 concentrations (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). Convoluting 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 CO₂ (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 iIRF₁₀₀ +/- 7 years) of their default value. The 300 random draws from each of the above distributions are then used as model parameters for the integration of the 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.

30 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) starting from an assumed quasi-equilibrium in 1850, atmospheric CO₂ concentrations increase faster than observed

(Figure 1a). This 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 CO₂ concentrations. Similarly this under-efficiency of carbon sinks requires lower than observed emissions to simulate the observed timeseries of atmospheric CO₂ 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 5 1980 when the FAIR airborne fraction rises above the PI-IR airborne fraction (figure 1c). This arises due to a combination of a lower pre-industrial iIRF₁₀₀ 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 iIRF₁₀₀ for these emissions relative to the 100GtC pulse in FAIR but not in PI-IR), and the PI-IR model (and the AR5-IR model) 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 10 represented in FAIR. Figure 1c shows large amplitude variations in the observed instantaneous airborne fraction that are likely to be driven in large part by unforced variability in the Earth-system and as such we would not expect these oscillations to be reproduced by any of 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. 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 15 in the AR5-IR model, driven by the substantially higher-than-observed present day CO₂ 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 20 are diagnosed from the MAGICC model (Meinshausen et al., 2011) in order to allow a comparison of 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., 2011). 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, it underestimates end of century concentrations, relative to MAGICC, to an even greater extent than 25 the AR5-IR model (Figure 2a) and concentrations fall from peak even quicker than MAGICC under the high mitigation RCP2.6 scenario (Figure 2b). 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 30 these two models from MAGICC clearly indicates that any state-insensitive impulse-response model is therefore unsuitable, unless modified, for long integrations with historical and projected emissions.

The FAIR model compares well to MAGICC 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 RCP8.5, but the behaviour of MAGICC (or indeed any other model) under these more extreme forcing scenarios has not been verified. 35 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

5 Figure 3 shows the response to a pulse-emission of CO₂ of differing magnitudes and against different backgrounds. For a 100GtC pulse of carbon set against an approximately present-day (389ppm) concentration background (PD100 - figure 3a), 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 (0.41). The FAIR iIRF₁₀₀ of 50.5 years lies within the ESM multi-model spread (see table 2). Excluding temperature feedbacks (the “biogeochemically-coupled” version of FAIR - setting
10 $r_T = 0$) on the carbon-cycle in FAIR increases the decay of the concentration response to the pulse over the century following the pulse emission, reducing the iIRF₁₀₀ by 10%. The 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).

15 Figure 3b and 3c show the response to a 100GtC and a 5000GtC pulse respectively, applied in pre-industrial conditions (named PI100 and PI5000 respectively). 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 in the background state. FAIR simulates an iIRF₁₀₀ approximately 32% lower in the pre-industrial case (34.3 years) relative to the present day, lying just within the ensemble spread of PI100 iIRF₁₀₀ simulated by the models of Joos et al. (2013) (34-47 years). The magnitude of the FAIR
20 simulated temperature response is similar in both the PD100 and PI100 cases due to the increased radiative efficiency of a pulse of CO₂ at lower background concentrations counteracting the faster decay of carbon out of the atmosphere. FAIR simulates a 100% increase of iIRF₁₀₀ in the 5000GtC pre-industrial pulse relative to the 100GtC pre-industrial pulse, consistent with the approximate doubling observed in the ESMs. 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
25 independent of pulse size or background conditions. The very rapid drawdown of CO₂ 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.

Restricting temperature-induced feedbacks on the carbon-cycle does not result in a substantial reduction in the iIRF₁₀₀
30 for the pre-industrial 100GtC pulse experiment (the “fully-coupled” and “biogeochemically-coupled” experiments lie on top of each other in figure 3b), whereas a 13% reduction in iIRF₁₀₀ is observed for the 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 iIRF₁₀₀ from a fixed pre-industrial value (r_0), therefore preventing substantial differences between “fully-coupled” and “biogeochemically-coupled” PI100 experiment in which iIRF₁₀₀ $\sim r_0$. For the 5000GtC pre-industrial pulse experiment

Experiment - $iIRF_{100}$ (yr)	AR5-IR	PI-IR	FAIR	Joos et al. (2013)
PD100	52.5	40.8	50.5	52.4 (± 11.3)
PD100 (no climate-carbon-cycle feedbacks)	-	-	45.2	-
PI100	52.5	40.8	34.3	40.6 (± 4.2)
PI100 (no climate-carbon-cycle feedbacks)	-	-	33.8	35.3*
PI5000	52.5	40.8	68.6	76.9 (± 6.2)
PI5000 (no climate-carbon-cycle feedbacks)	-	-	60.0	65.4*

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.

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.

- 5 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 2.3.1) 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, attributable to a range of differences in carbon-cycle process
- 10 representations within 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 (fits are worse for models in which inter-annual variability is simulated), as well achieving the correct ratio between the PD100 and PI100 responses for models across the ESM/EMIC ensemble.

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
- 15 of CO_2 occurs approximately a decade after emission, but 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 2.3.1), despite concentrations rapidly peak and then subsequently declining in all cases (figure 5b), the total warming in the FAIR model fails to plateau as quickly at higher pulse sizes compared to at lower pulse sizes (figure 5a), as the balance between carbon-cycle cooling and long-timescale
- 20 thermal warming takes centuries to reach 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 shape of the warming 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).

3.3 Response to idealised concentration increase experiments

The set of “fully-coupled”, “biogeochemically-coupled” and “radiatively-coupled” experiments from Gregory et al. (2009) (see section 2.3.2) 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 blue lines in figure 6a) under a 1%/yr CO₂ 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. Figure 6b shows C_{acc} as a function of atmospheric CO₂ 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.

Oceanic carbon-cycle feedbacks are almost exclusively driven by biogeochemical effects (Glotter et al., 2014). However, simple global climate-carbon-cycle models need to also capture dependencies of the land carbon uptake on warming. Aside from 3 ESMs that display global-mean carbon-cycles relatively insensitive to warming, figure 6c shows a relationship between temperature increases and the size of the carbon outgassing back to the atmosphere in the ‘radiatively-coupled’ FAIR experiment similar, in both shape and magnitude, to that displayed in the ESMs under the 1%/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 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 fraction of all past emissions remaining in the atmosphere) followed by a subsequent increase (figure 6d) simulated by FAIR is a feature of the response of many ESMs under a 1%/yr increasing CO₂ 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 $iIRF_{100}$). Therefore, as more recent emissions (which increase monotonically with time in the prescribed concentration increase experiments) have a higher instantaneous airborne fraction, the initial decrease in cumulative airborne fraction stops and subsequently begins to increase as this accelerating saturation becomes the dominant effect.

3.4 Probabilistic parameter sampling within the FAIR model

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. Figures 7a and 7b show the isolated impact of thermal response uncertainty under idealised prescribed concentration scenarios, namely a 1%/yr CO₂ concentration increase (figure 7a) and an instantaneous quadrupling of CO₂ 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 q_j (equations 4 and 5). The blue shading in panels a) and b) of figure 7 show 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). Expressed in terms of q_j , 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 CO₂ 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 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, correlated and time invariant perturbations of +/-13% to the r_0 , r_T and r_C parameters (combined with perturbations to q_1 and q_2 consistent with the IPCC-AR5 likely ranges) in the FAIR model approximately spans the range of responses seen in the CMIP5 models under a 1%/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 approximately linear relationship (to first order) 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 with parameter settings given in section 2, the FAIR model has a Transient Response to Cumulative Emissions (TCRE⁴) = 1.3K/TtC (thick blue line in 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 more-pronounced concave-downward shape in the plot of realised warming against cumulative carbon emissions, as the decline of the cumulative airborne fraction is unable to compensate (as it does in more complex models) for the logarithmic relationship between CO₂ concentration and radiative forcing (Millar

⁴TCRE is defined as the annual mean global surface temperature change per unit of cumulative CO₂ emissions, usually 1000GtC, in a scenario with continuing emissions (Collins et al., 2013).

et al., 2016). The 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).

Figure 8 shows the combined effect of uncertainty in the FAIR model parameters, both thermal parameters (the joint distribution of TCR and ECS is shown in figure 8a) and carbon-cycle feedback parameters. Representative uncertainty in FAIR parameters (see section 2.3.3) propagates into an emergent 5-95% range (based on 300 draws from the input parameter distributions) for TCRE (figure 8b) of 1.0-2.5K/TtC when integrated under a 1%/yr CO₂ concentration increase scenario (figure 8b), broadly consistent with the IPCC-AR5 *likely* range (0.8-2.5K/TtC). The temperature anomaly associated with a 100GtC pulse in 2020 against an RCP2.6 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.

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 CO₂ uptake through a parameterisation of the 100-year integrated impulse-response function, $iIRF_{100}$. 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).

We have shown that including both explicit CO₂ 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 CO₂ uptake-feedbacks would fail to incorporate well-understood physical mechanisms governing the response of ocean carbonate chemistry to anthropogenic CO₂ 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. 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 CO₂ and non-CO₂ 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 CO₂ 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 iIRF₁₀₀, 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 iIRF₁₀₀ 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 CO₂ 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 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 CO₂ emissions. Tuning of parameters within 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.:
5 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.,
10 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 CO₂ 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.
- 15 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., Voltaire, 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.
- 20 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 CO₂ 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
25 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,
30 *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., Fuglestedt, 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
35 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 CO₂ emissions, *Journal of Climate*, 28, 9955–9968, 2015.
- 5 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.
- 10 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, 2011.
- 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.
- 15 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.
- 20 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.
- 25 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 CO₂ during the past decades, *Tellus*, 9, 18–27, 1957.
- 30 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.
- 35 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.
- 5 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.
- 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.
- 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,
10 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
15 emission, *Environmental Research Letters*, 10, 031 001, 2015.
- Zickfeld, K., Eby, M., Alexander, K., Weaver, A. J., Cresspin, 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.

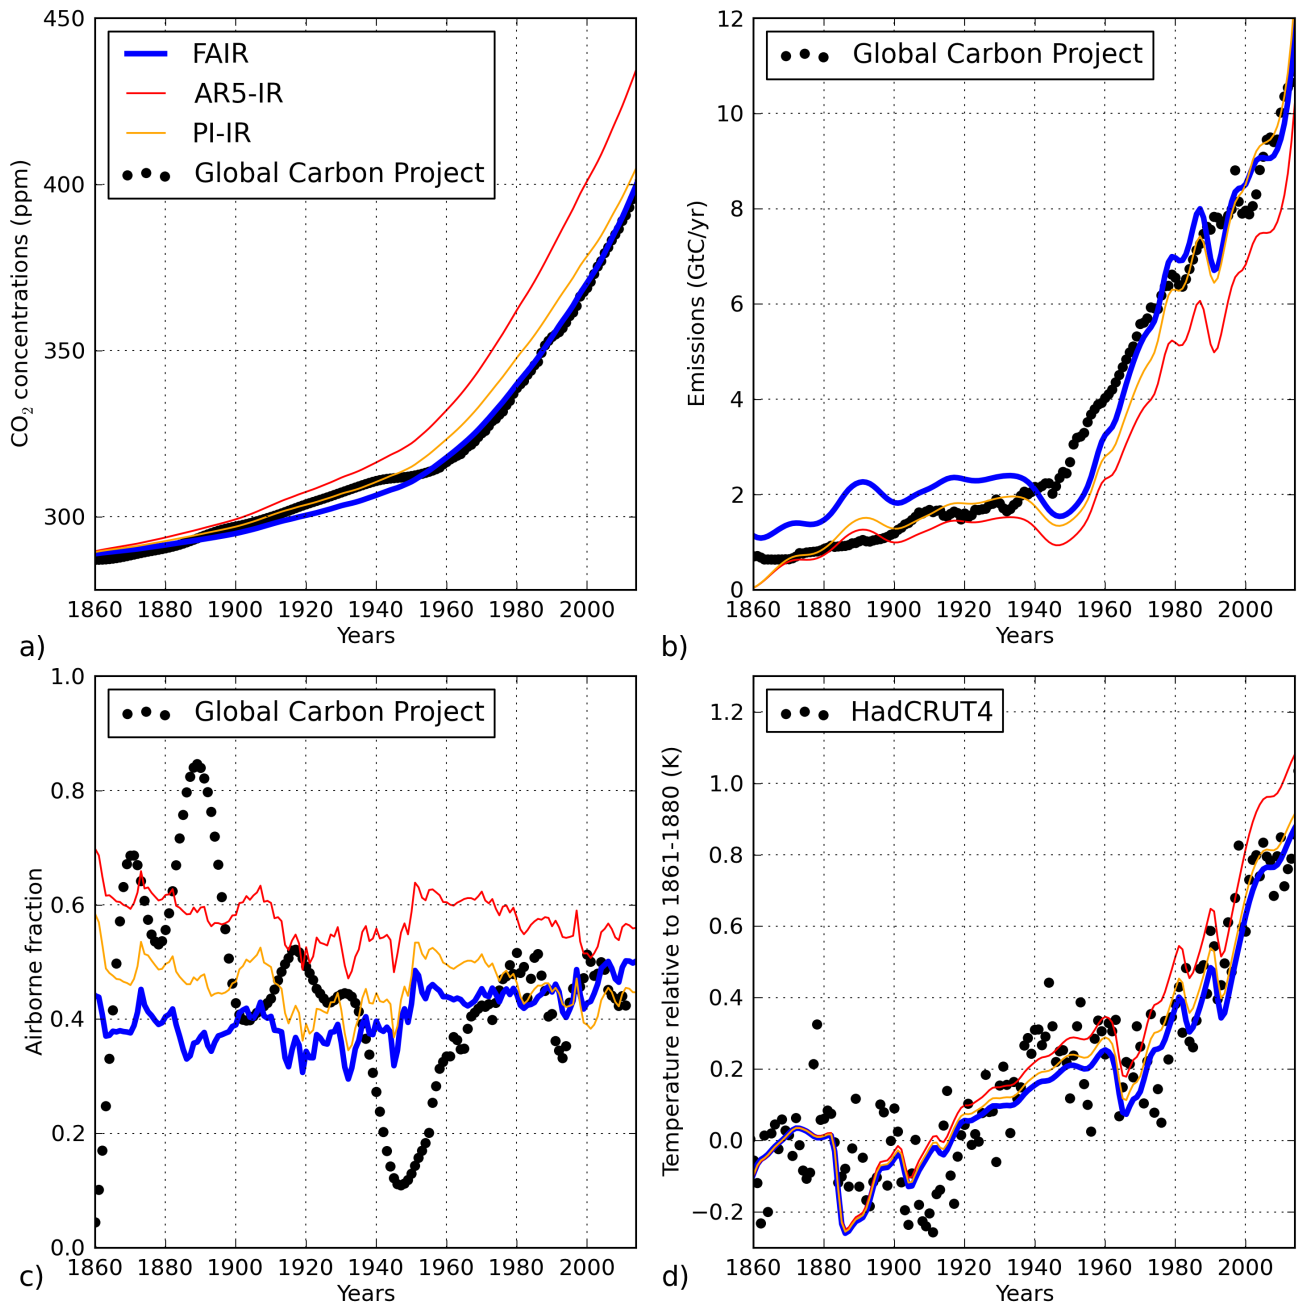


Figure 1. Historical validation of the FAIR (blue), AR5-IR (red) and PI-IR (orange) models. Panel a) shows simulated CO₂ concentrations when driven by historical emissions (and historical non-CO₂) as estimated by MAGICC. Panel b) shows diagnosed CO₂ emissions consistent with historical concentrations. Panel c) shows the evolution of annual airborne fraction (smoothed with a 7-year running mean for the observations), and d) the warming anomaly, in the models when driven by historical emissions. Historical observations are shown as black dots in all panels, a), b) and c) show Le Quéré et al. (2015) and HadCRUT4 (Morice et al., 2012) is shown in d). All simulations are commenced from assumed quasi-equilibrium carbon-cycle states in 1850.

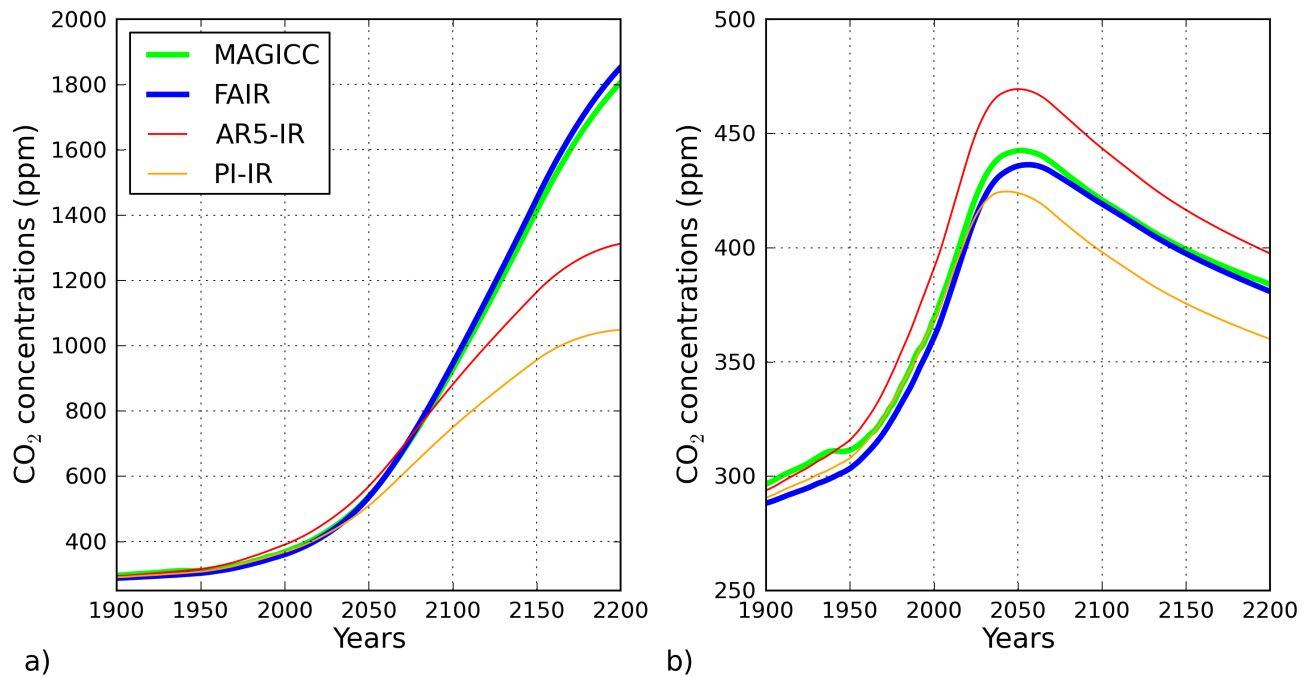


Figure 2. CO₂ concentrations under RCP8.5 (a) and RCP2.6 (b). FAIR (blue), AR5-IR (red), PI-IR (orange) and MAGICCC (green) are shown in both panels.

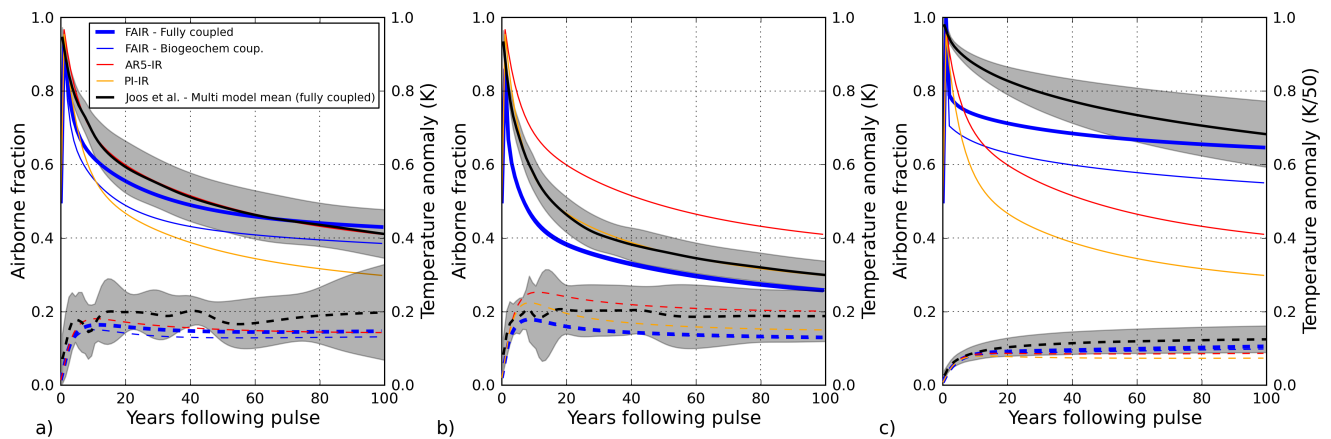


Figure 3. Response to pulse emission experiments of Joos et al. (2013). Panel a) shows the response to a 100GtC imposed on present-day (389ppm) background conditions (PD100 experiment), panel b) the response to a 100GtC pulse in pre-industrial conditions (PI100 experiment) and panel c) the response to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

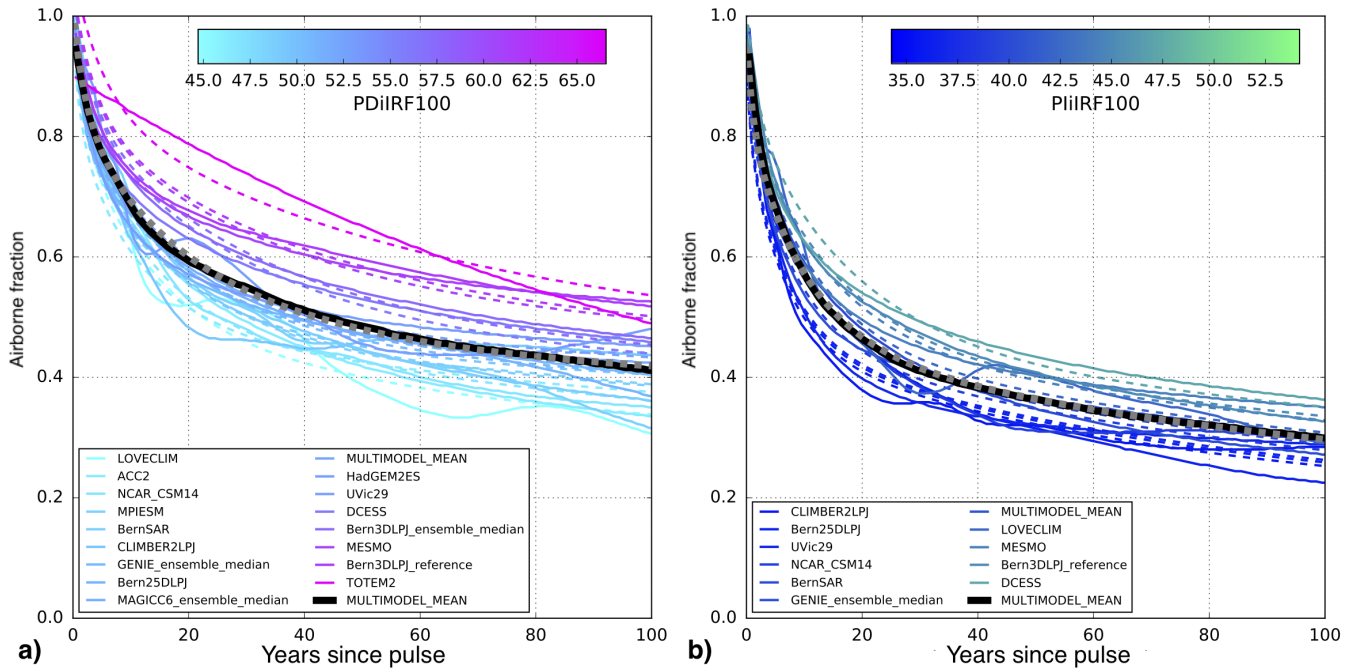


Figure 4. Fitting individual models from Joos et al. (2013) with FAIR. Panel a) shows the remaining airborne fraction for the PD100 experiment and panel b) for those models that additionally completed the PI100 experiment. Solid lines show the original model response coloured by the $iIRF_{100}$ values.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

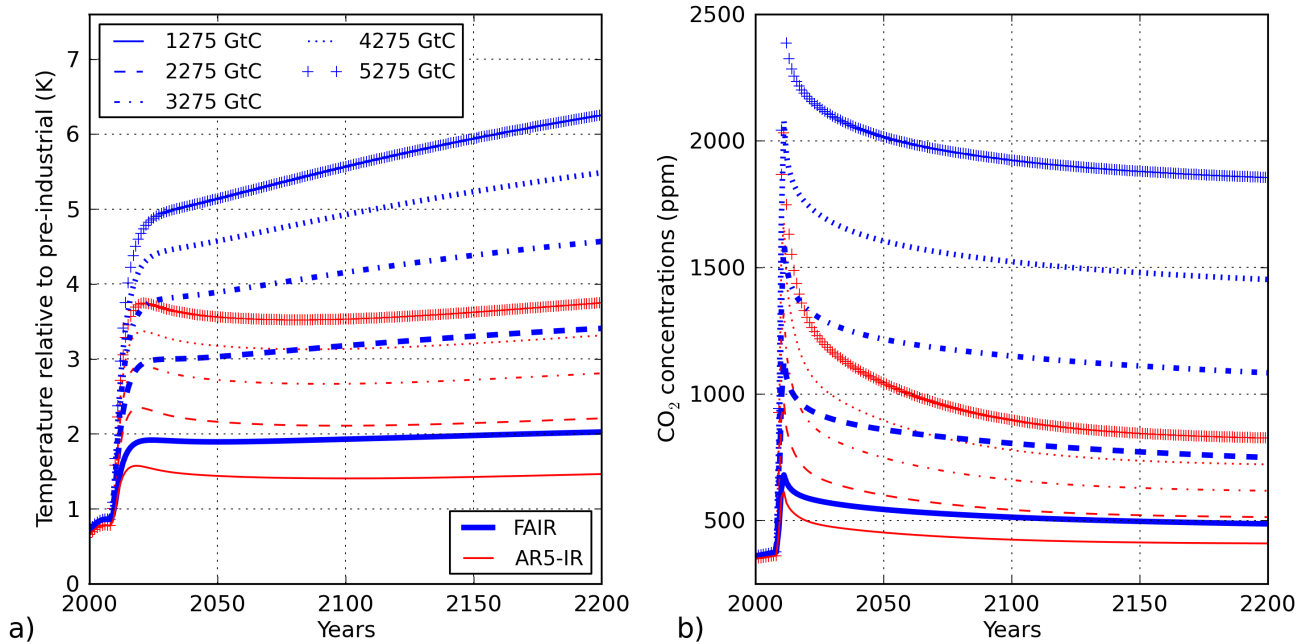


Figure 5. Dependency of temperature response on pulse size. Panel a) shows the global mean surface temperature response to the pulse experiments of Herrington and Zickfeld (2014). Responses are shown for the FAIR (blue) and AR5-IR (red) models. Panel b) shows the corresponding concentration response.

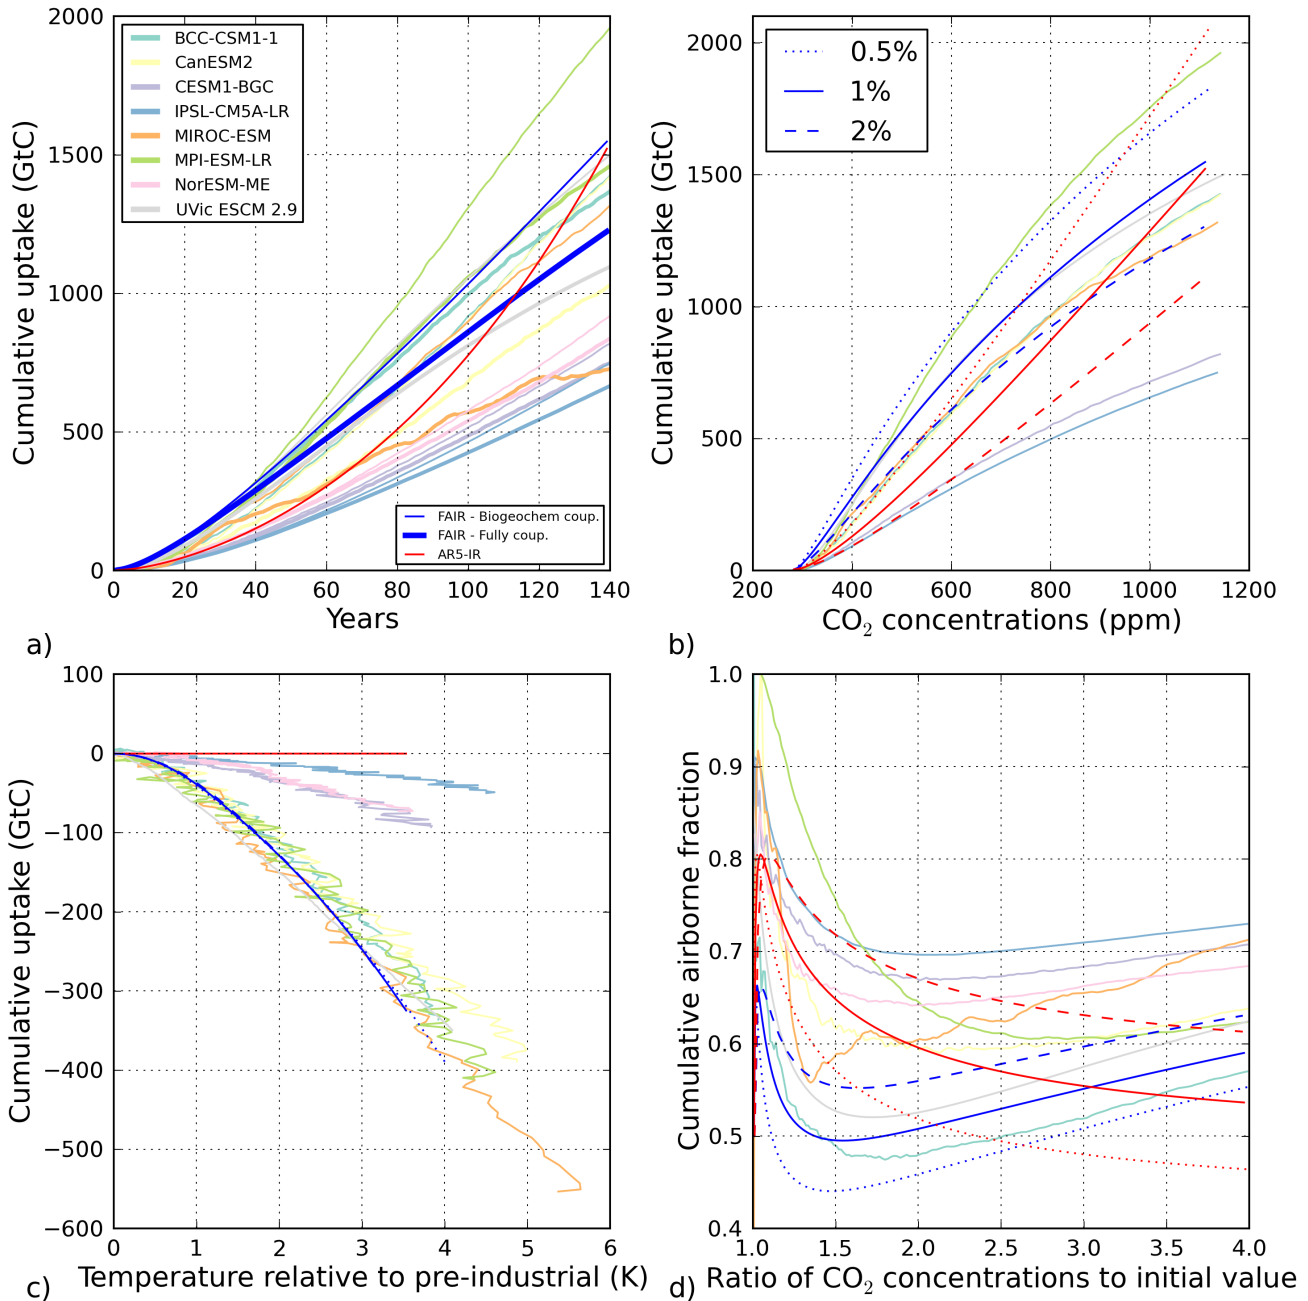


Figure 6. Response to idealised concentration increase experiments from Gregory et al. (2009) for the FAIR (blue) and AR5-IR (red) models. Light pastel colours show the ESMs from Joos et al. (2013) for the 1%/yr concentration increase scenario only. Panel a) shows the cumulative ocean and land carbon uptake over time in the “fully coupled” 1%/yr⁻¹ concentration increase scenario. Panel b) shows the evolution of cumulative ocean and land carbon uptake 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.

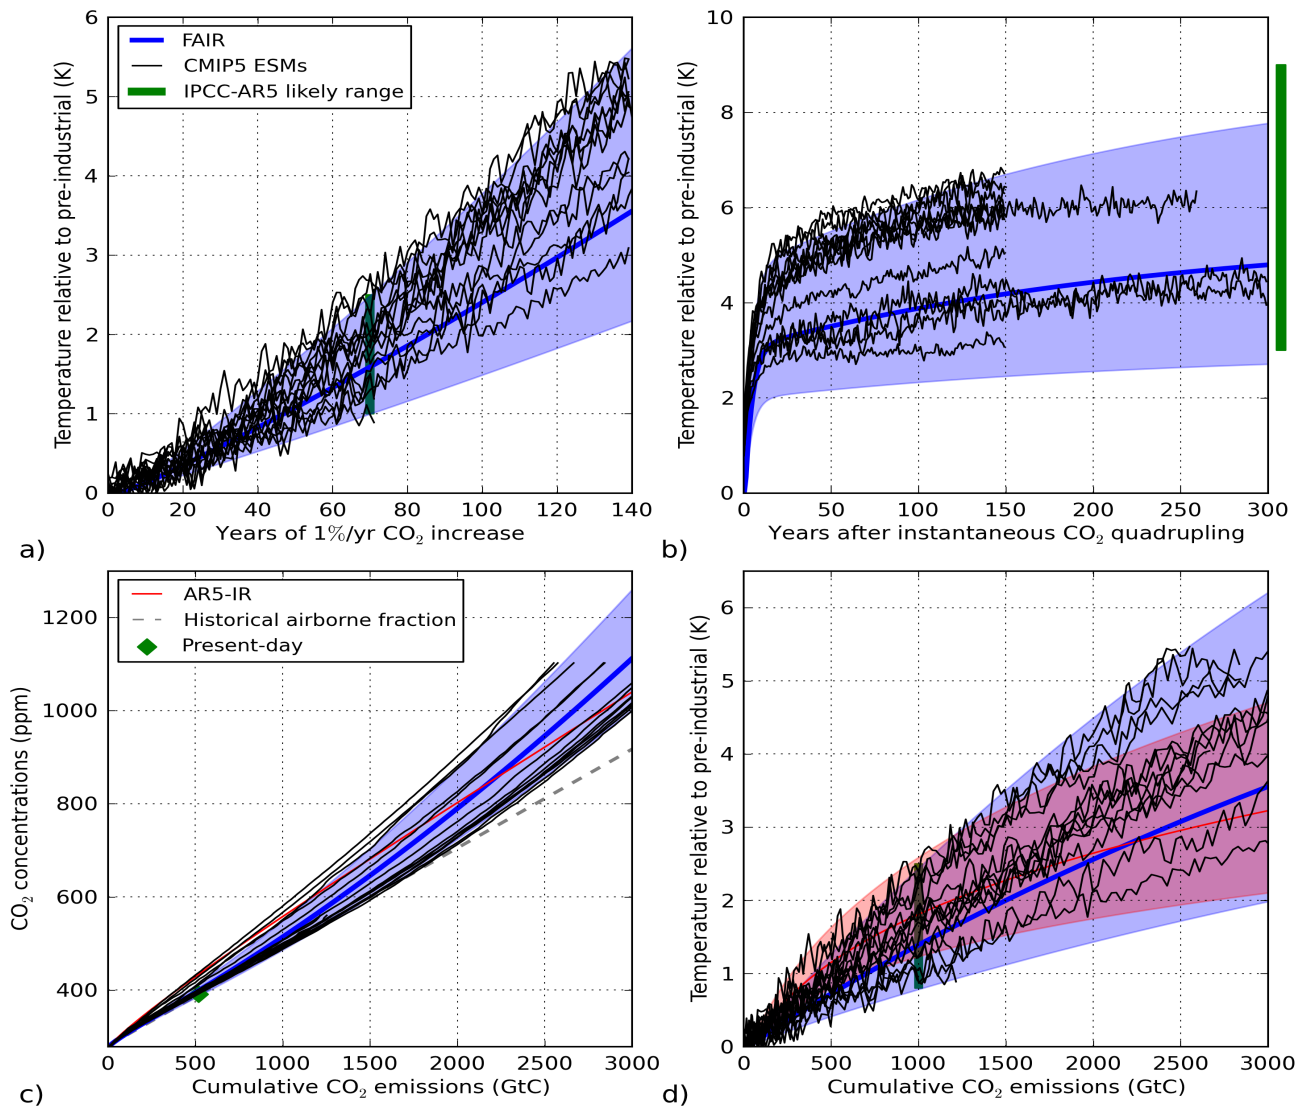


Figure 7. Climate response uncertainties in the FAIR (blue), AR5-IR (red) and CMIP5 (black) models. Panel a) shows the temperature responses to a 1%/yr concentration increase scenario. The green bar indicates the IPCC-AR5 TCR likely range. The blue shading in panels a) and b) shows the response of FAIR under IPCC-AR5 upper and lower likely TCR and ECS ranges. Panel b) shows the responses to an instantaneous quadrupling of atmospheric CO₂ which is held fixed subsequently. The 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. The plumes in panels c) and d) show the response for the IPCC-AR5 likely TCR and ECS ranges, with an additional 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 of the climate system (green diamond). Panel d) shows warming as a function of cumulative emissions in the 1%/yr scenario. The green bar shows the IPCC-AR5 likely 0.8-2.5K/TtC assessed range for TCRE.

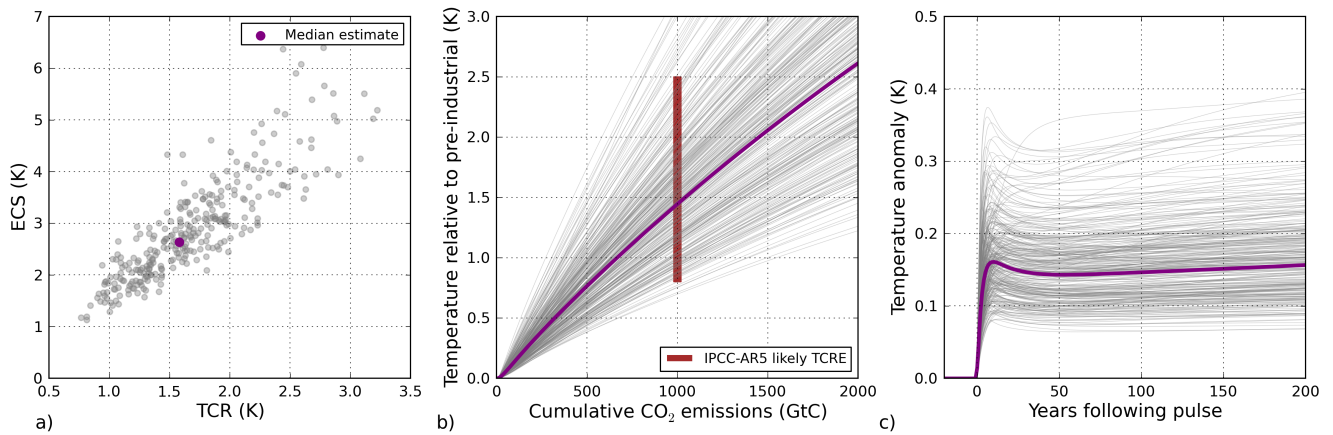


Figure 8. Probabilistic sampling in the FAIR model. Grey lines show 300 random draws from the input parameter distributions, as described in section 3.4. Panel (a) shows the joint distribution of TCR and ECS. Panel (b), warming as a function of cumulative emissions in the $1\%yr^{-1}$ concentration increase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