



## Decomposition of the Zero Emissions Commitment into thermal warming pipeline and carbon buffering components

Benjamin M. Sanderson<sup>1</sup>, Marit Sandstad<sup>1</sup>, Alejandro Romero-Prieto<sup>2,3</sup>, Stuart Jenkins<sup>4</sup>, Charles Koven<sup>5</sup>, Glen P. Peters<sup>1</sup>, Roland Séférian<sup>6</sup>, Andrew H. MacDougall<sup>7</sup>, Ashwin Seshadri<sup>8</sup>, Victor Brovkin<sup>9</sup>, John Dunne<sup>10</sup>, Abigail L. S. Swann<sup>11</sup>, Torben Koenigk<sup>12</sup>, Rosie A. Fisher<sup>1</sup>, David Hohn<sup>13</sup>, Tatiana Ilyina<sup>14</sup>, Chris D. Jones<sup>3,16</sup>, Hongmei Li<sup>9,15</sup>, Peter Lawrence<sup>17</sup>, Spencer Liddicoat<sup>3</sup>, Nadine Mengis<sup>13</sup>, Anastasia Romanou<sup>18,19</sup>, Lori T. Sentman<sup>10</sup>, Chris Smith<sup>20</sup>, Norman J. Steinert<sup>1</sup>, Jerry Tjiputra<sup>21</sup>, and Tilo Ziehn<sup>22</sup>

<sup>1</sup>CICERO International Center for Climate Research, Oslo, Norway

<sup>2</sup>Priestley Centre for Climate Futures, University of Leeds, Leeds, United Kingdom

<sup>3</sup>Met Office Hadley Centre, Exeter, UK

<sup>4</sup>University of Oxford, Oxford, UK

<sup>5</sup>Lawrence Berkeley National Laboratory, Berkeley, CA, USA

<sup>6</sup>Centre National de Recherches Météorologiques (CNRM), Toulouse, France

<sup>7</sup>St. Francis Xavier University, Antigonish, NS, Canada

<sup>8</sup>Indian Institute of Science, Bangalore, India

<sup>9</sup>Max Planck Institute for Meteorology, Hamburg, Germany

<sup>10</sup>NOAA/OAR/Geophysical Fluid Dynamics Laboratory, Princeton, NJ, USA.

<sup>11</sup>Atmospheric and Climate Science, University of Washington, Seattle, WA, USA

<sup>12</sup>Swedish Meteorological and Hydrological Institute (SMHI), Norrköping, Sweden

<sup>13</sup>GEOMAR, Helmholtz Centre for Ocean Research, Kiel, Germany

<sup>14</sup>Max Planck Institute for Marine Microbiology, Bremen, Germany

<sup>15</sup>Helmholtz-Zentrum Hereon, Geesthacht, Germany

<sup>16</sup>School of Geographical Sciences, University of Bristol, Bristol, UK

<sup>17</sup>NSF National Center for Atmospheric Research (NCAR), Boulder, CO, USA

<sup>18</sup>NASA Goddard Institute for Space Studies, New York, NY, USA

<sup>19</sup>Applied Physics and Applied Mathematics, Columbia University, New York, USA

<sup>20</sup>International Institute for Applied Systems Analysis, Laxenburg, Austria

<sup>21</sup>NORCE Norwegian Research Centre, Bjerknes Centre for Climate Research, Bergen, Norway

<sup>22</sup>CSIRO Environment, Aspendale, Australia

**Correspondence:** Benjamin M. Sanderson (benjamin.sanderson@cicero.oslo.no)

**Abstract.** The Zero Emissions Commitment (ZEC), the residual warming after anthropogenic CO<sub>2</sub> emissions cease, remains poorly constrained and directly affects remaining carbon budget calculations. We fit a simple coupled carbon–climate model to ten Earth System Models (ESMs) in the flat10 Model Intercomparison Project. Near-zero multi-model mean ZEC can be decomposed into unrealised ocean warming contributing roughly +0.2 K at 50 years post-cessation, which is almost exactly offset by ~−0.2 K of cooling from atmospheric CO<sub>2</sub> drawdown into land and ocean carbon sinks. Yet the thermal term carries more than twice the inter-model spread of the carbon term; as such, in the context of current ESMs, reducing ZEC uncertainty depends more strongly on equilibrium climate sensitivity and the slow ocean heat uptake timescale than on carbon-cycle parameters. Individual models split into a carbon-dominated majority (eight of ten, negative ZEC) and a thermal-dominated



minority (positive ZEC), distinguished primarily by their realised warming fraction and equilibrium climate sensitivity. The simple model yields a compact analytical formula for ZEC in terms of component climate and carbon-cycle parameters, reproducing ESM-simulated values to within  $\sim 0.03$  K (less than 5% of the  $\sim 0.6$  K inter-model spread). The decomposition also reveals that the carbon cycle buffers ZEC against uncertainty in committed warming to an existing energetic imbalance: recent downward revisions of the realised warming fraction (the ratio of transient to equilibrium warming) imply substantially more committed warming at fixed atmospheric composition, yet this additional heat is strongly attenuated before it appears in ZEC. A reduction of 10 per cent in the realised warming fraction would increase the radiatively committed warming by  $\sim 0.5$  K at cessation, but this shifts the 50-year ZEC by only  $\sim 0.05$  K. This allows potential for process-level constraints on climate and carbon-cycle components on ZEC to narrow carbon budgets. Out-of-sample testing on the flat10-cdr experiment shows that  $\text{CO}_2$  predictions diverge from ESMs within  $\sim 50$  years of the emission trajectory departing from the training rate, so the framework does not extend to full climate reversibility under sustained negative emissions. Structural limitations, including under-sampled land biosphere diversity across ESMs, may further emerge at higher warming levels or on multi-centennial timescales.

## 1 Introduction

One of the most robust results across Earth System Models is that global mean surface temperature is approximately proportional to cumulative  $\text{CO}_2$  emissions (Matthews et al., 2009; Allen et al., 2009; Eby et al., 2009; Zickfeld et al., 2009). This near-linear relationship, quantified by the Transient Climate Response to Cumulative Emissions (TCRE), underpins remaining carbon budget calculations. However, when anthropogenic  $\text{CO}_2$  emissions reach net zero (Allen et al., 2022), some models exhibit a residual temperature change, the Zero Emissions Commitment (ZEC), which characterises the imperfect compensation of two opposite-sign processes: the continued delivery of unrealised warming from the deep ocean, and the multi-decadal to millennial drawdown of atmospheric  $\text{CO}_2$  by land and ocean sinks (MacDougall et al., 2020; Palazzo Corner et al., 2023). The sign and magnitude of ZEC in principle affect remaining carbon budget calculations (though current assessments assume  $\text{ZEC} \approx 0$ ; Rogelj et al. 2019): a positive ZEC reduces the available budget, while a negative ZEC would allow positive residual emissions to remain compatible with stabilising temperatures at a given level.

The realised warming fraction and equilibrium climate sensitivity together determine the remaining warming for a given atmospheric composition. Recent work has focused on this residual warming under the heading of “warming in the pipeline.” (Hansen et al., 2023). Although this phrase has been used both for committed warming at fixed atmospheric composition and for the net post-cessation temperature change including carbon-cycle feedbacks (as in the title of MacDougall et al., 2020), throughout this paper we use the term “thermal pipeline” exclusively in the former, thermal sense, corresponding to the warming that would result from holding atmospheric carbon dioxide concentrations fixed after a period of emissions.

Hansen et al. (2023) estimate that Earth’s current energy imbalance implies several tenths of a degree of additional warming at fixed composition at present day concentrations, using a two-box model with an ocean heat uptake efficacy parameter to capture the pattern effect on radiative feedbacks. Armour et al. (2024) show that observed sea-surface temperature patterns



differ from free-running models and may have suppressed the effective climate sensitivity (EffCS) inferred from recent decades. Forcing models that have a high climate sensitivity with observed SSTs suggests that recent temperature evolution provides little evidence on equilibrium response; the implication is that the thermal pipeline of unrealised warming, set by the gap  
45 between transient and equilibrium warming ( $1 - r$  in our notation), could be substantially larger than energy-budget estimates of ECS from recent decades would suggest, because the pattern effect biases those estimates low. Myhre et al. (2025) reinforce this by demonstrating that the observed strengthening of Earth's energy imbalance over 2001–2023 is best reproduced by models with climate sensitivity above 2.5 K, providing a lower bound on the thermal pipeline. These studies use different tools, from simple energy balance models (Hansen et al., 2023) to pattern-effect analyses of CMIP ensembles (Armour et al.,  
50 2024) and observational energy-budget constraints (Myhre et al., 2025), but the phenomenon they identify is the same: state-dependent radiative feedbacks driven by evolving sea-surface temperature patterns cause the effective climate sensitivity in recent decades to fall well below the equilibrium value. The two-box energy balance model with a deep-ocean heat-uptake efficacy parameter (Held et al., 2010; Geoffroy et al., 2013; Winton et al., 2010) captures this behaviour in a simple framework, separating the fast surface adjustment from slow deep-ocean equilibration; it has been shown to reproduce the transient response  
55 of comprehensive ESMs across a range of forcing scenarios (Geoffroy et al., 2013). These studies focus on the thermal response in isolation. To determine the full ZEC, we also need to account for the post-cessation drawdown of CO<sub>2</sub> by land and ocean sinks, the process that offsets the thermal pipeline in the coupled carbon–climate system.

The Zero Emissions Commitment Model Intercomparison Project (ZECMIP; MacDougall et al. 2020) used an experimental design branching from the 1pctCO<sub>2</sub> experiment when diagnosed cumulative CO<sub>2</sub> emissions reached 1000 PgC, corresponding  
60 to model-specific concentrations in the range 500–600 ppm. They found ZEC<sub>50</sub> (the temperature change 50 years after cessation of emissions) ranges from approximately  $-0.3$  to  $+0.3$  K across participating ESMs, with a multi-model mean close to zero. This near-zero mean conceals large inter-model disagreement, and the physical drivers of this spread have not been fully identified, though MacDougall et al. (2020) presented a partial decomposition into ocean heat uptake and CO<sub>2</sub> drawdown terms. Koven et al. (2023) showed that much of the ZEC manifests before net zero is reached in gradual emission reduction  
65 scenarios, arguing that ZEC is better understood as the deviation from TCRE proportionality rather than as post-cessation warming alone. That is, the imbalance between unrealised thermal warming and carbon drawdown that determines ZEC builds up during the transition to net zero, not only after emissions cease.

The flat10 Model Intercomparison Project (flat10MIP; Sanderson et al. 2025a), motivated by the need for emissions-driven diagnostics in CMIP7 (Sanderson et al., 2024; Dunne et al., 2025), provides a clean experimental framework for assessing  
70 ZEC: constant emissions of 10 PgC yr<sup>-1</sup> for 100 years followed by abrupt cessation, avoiding the model-specific compatible emissions inherent in the 1% yr<sup>-1</sup> CO<sub>2</sub> concentration-driven protocol.

Several studies have attempted to bridge the thermal-pipeline and ZEC literatures analytically. Jenkins et al. (2022) characterises the post-cessation temperature evolution using the rate of adjustment to zero emissions (RAZE), combining linearised thermal and carbon response rates into a single net rate (spanning  $-0.24$  to  $+0.17\%$  yr<sup>-1</sup> in their constrained FaIR ensemble)  
75 that determines the residual CO<sub>2</sub> emissions compatible with halting warming. Jenkins et al. (2022) demonstrate deterministic FaIR emulations of five ZECMIP ESMs and produce an observationally-constrained 10 000-member ensemble spanning the



aggregate RAZE distribution. Although their framework can disaggregate the thermal and CO<sub>2</sub> drawdown contributions to the RAZE, they report a single scalar rate and do not decompose the inter-model spread into its component drivers, nor do they produce per-model probabilistic calibrations. The RAZE derivation also collapses the deep-ocean heat-uptake efficacy into the aggregate thermal adjustment rate, preventing its influence on ZEC from being isolated, and linearises the carbon cycle without an explicit carbon–climate feedback.

Williams et al. (2025) introduced a normalised decomposition of ZEC into thermal, radiative, and carbon contributions applied to nine ZECMIP ESMs and a large ensemble of the Warming Acidification and Sea-level Projector (WASP) simple climate model (Goodwin, 2016). They find that the thermal component explains 58% of inter-model variance and the carbon component 40%, with the radiative contribution negligible. But this decomposition is diagnostic: it does not yield closed-form analytical expressions linking ZEC to standard climate metrics such as ECS or the realised warming fraction, and it does not fit individual ESMs probabilistically.

Brovkin et al. (2025) take a more direct approach to the same flat10MIP data, building on earlier work by MacDougall and Friedlingstein (2015) who showed that the near-proportionality between warming and cumulative emissions arises from a cancellation between declining radiative efficiency of CO<sub>2</sub> and increasing airborne fraction, and identified the conditions under which TCRE breaks down. Brovkin et al. note that most ESMs exhibit a quasi-linear relationship between cumulative land and ocean carbon uptakes ( $C_l \approx k C_o$ ) and a quasi-linear relationship between cumulative ocean carbon and heat uptakes ( $C_o \approx \eta H$ ) during the emissions phase (a linearisation that may in part be facilitated by the constant emission rate of the flat10 protocol). Together, these two empirical linearisations reduce the coupled temperature–CO<sub>2</sub> system to a single ordinary differential equation for the atmospheric carbon amount, which admits an analytical solution in the linear-forcing limit and a numerical solution under the more realistic logarithmic forcing. An insight from their numerical solution (consistent with MacDougall and Friedlingstein (2015) and the original demonstration by Matthews et al. 2009 that TCRE arises from offsetting effects) is that the approximate constancy of TCRE arises from a compensation between the weakening radiative forcing per unit CO<sub>2</sub> and the concomitant rise in the airborne fraction, a cancellation that is invisible in the analytical (linear-forcing) case. However, these frameworks considered the thermal/carbon partitioning of the transient response, without assessing whether the relationships between these terms would be sustained after the cessation of emissions.

The present study aims to reconcile some of these lines of inquiry. We use a coupled carbon–climate model with 11 physically identifiable parameters that embeds the same two-box energy balance framework used in the thermal pipeline literature, retaining an explicit deep-ocean heat-uptake efficacy (Held et al., 2010; Geoffroy et al., 2013; Winton et al., 2010), while simultaneously coupling it to an impulse-response carbon cycle with a carbon–climate feedback. This structure lets us connect the thermal pipeline (governed by the gap between the realised warming fraction and equilibrium, modulated by efficacy) to the atmospheric carbon drawdown amount that offsets it, within a single analytical expression for ZEC. Our framework resolves the thermal and carbon sub-components of the RAZE while retaining efficacy and the carbon–climate feedback as independent parameters, and calibrates probabilistically against each of ten flat10MIP ESMs with out-of-sample validation on the flat10-zec and flat10-cdr experiments. Relative to Williams et al. (2025), it provides closed-form analytical solutions amenable to direct sensitivity analysis, model-specific history-matched parameter ensembles that propagate parametric uncertainty, and



out-of-sample testing that demonstrates predictive skill beyond the training data. By decomposing ZEC into thermal and carbon components, we show that the near-zero total reflects a partial compensation between a thermal warming of +0.216 K and a carbon cooling of  $-0.231$  K across ten ESMs, confirming the earlier finding of MacDougall et al. (2020) and extending it to emissions-driven flat10MIP experiments. This compensation arises because the rate of deep-ocean heat delivery and the rate of CO<sub>2</sub> drawdown cooling happen to be comparable, a physical coincidence identified in the review of Palazzo Corner et al. (2023) as a priority for further investigation, and that underpins the approximate constancy of the TCRE near net zero.

The paper proceeds as follows. Section 2 describes the 12-parameter coupled model. Section 3 presents the fitting procedure, including the generation of history-matched parameter ensembles. Section 4 derives the analytical decomposition of ZEC into thermal and carbon components, culminating in a compact Level 2 formula. Section 5 presents results: model fits, ZEC distributions, parameter drivers, analytical validation, and out-of-sample tests. Section 6 discusses implications and remaining uncertainties, and Section 7 concludes.

## 2 Model description

We use a 12-parameter simple climate model that couples a carbon cycle to a climate response (for a review of such models, see Romero-Prieto et al., 2026). This forward model is integrated numerically for all parameter fitting and ensemble generation; the analytical approximations developed in section 4 are used only for physical interpretation. The carbon cycle determines atmospheric CO<sub>2</sub> for a given emission pathway; the climate model converts the resulting radiative forcing into a temperature response. The two components interact through a carbon–climate feedback in which warming reduces carbon uptake efficiency.

Radiative forcing follows the commonly used logarithmic approximation (Myhre et al., 1998):

$$F = a \ln\left(\frac{C_0 + \Delta\text{CO}_2}{C_0}\right), \quad a = 5.35 \text{ W m}^{-2}, \quad C_0 = 284.3 \text{ ppm}, \quad (1)$$

giving  $F_{2\times} = a \ln 2 \approx 3.71 \text{ W m}^{-2}$  for CO<sub>2</sub> doubling.

### 2.1 Three-pool impulse-response carbon cycle

Atmospheric CO<sub>2</sub> (in ppm) responds to emissions  $E(t)$  (PgCyr<sup>-1</sup>) via convolution with an impulse-response function (IRF; Maier-Reimer and Hasselmann, 1987; Joos et al., 1996, 2013):

$$\Delta\text{CO}_2(t) = \frac{1}{\kappa} \int_0^t E(t') \text{IRF}(t-t', t') dt', \quad (2)$$

where  $\kappa = 2.124 \text{ PgC ppm}^{-1}$  converts mass to mixing ratio. The IRF comprises a permanent airborne fraction  $f_\infty$  and three exponential uptake pools:

$$\text{IRF}(s, t') = f_\infty^*(t') + \sum_{i=1}^3 f_i^*(t') e^{-s/\tau_i}, \quad (3)$$



**Table 1.** Model parameters and their physical interpretation

Symbol	Component	Meaning	Typical range
<i>Carbon cycle (7 parameters)</i>			
$f_\infty$	Permanent	Long-term airborne fraction	0.15–0.30
$f_1, \tau_1$	Very fast	Surface ocean equilibration	$f$ : 0.2–0.5, $\tau$ : 3–15 yr
$f_2, \tau_2$	Fast	Land + mixed-layer ocean	$f$ : 0–0.4, $\tau$ : 15–60 yr
$\tau_3$	Slow	Deep ocean uptake ( $f_3 = 1 - f_\infty - f_1 - f_2$ )	100–400 yr
$\zeta$	Feedback	Fractional uptake reduction per K	0.01–0.06 K <sup>-1</sup>
<i>Climate (5 parameters)</i>			
$C_f$	Fast box	Atmosphere + upper-ocean heat capacity	5–15 W yr m <sup>-2</sup> K <sup>-1</sup>
$C_s$	Slow box	Deep-ocean heat capacity	50–300 W yr m <sup>-2</sup> K <sup>-1</sup>
$\lambda_f$	Feedback	Climate feedback parameter	0.5–3 W m <sup>-2</sup> K <sup>-1</sup>
$\gamma$	Exchange	Heat transfer between boxes	0.1–1.5 W m <sup>-2</sup> K <sup>-1</sup>
$\varepsilon$	Efficacy	Deep-ocean heat uptake efficacy	1.0–1.8

where  $s = t - t'$  is the age of the pulse. The pool fractions  $\{f_\infty, f_1, f_2, f_3\}$  satisfy mass conservation  $f_\infty + \sum_i f_i = 1$ , and the timescales  $\tau_1 < \tau_2 < \tau_3$  represent progressively slower uptake reservoirs that can be broadly associated with mixed-layer ocean and rapid biosphere exchange, intermediate biosphere and thermocline uptake, and deep-ocean ventilation, respectively (table 1), though the mapping between IRF pools and specific Earth-system reservoirs is approximate (Joos et al., 2013).

The dependence of the pool fractions on the emission time  $t'$  introduces a carbon–climate feedback: warming reduces the efficiency of carbon sinks. Following Smith et al. (2018), we parameterise this as a linear scaling of uptake with temperature:

$$f_i^*(t') = f_i(1 - \zeta T(t')), \quad f_\infty^*(t') = 1 - \sum_i f_i^*(t'), \quad (4)$$

where  $\zeta$  (K<sup>-1</sup>) is the carbon–climate feedback parameter. Setting  $\zeta = 0$  recovers the standard time-invariant IRF.

## 2.2 Two-box energy balance climate model

The temperature response to radiative forcing  $F(t)$  is governed by a two-box energy balance model (Held et al., 2010; Geoffroy et al., 2013; Allen et al., 2022). In its most general form, the model includes an *efficacy* parameter  $\varepsilon$  on the ocean heat uptake term (Winton et al., 2010; Geoffroy et al., 2013; Seshadri, 2017):

$$C_f \frac{dT_f}{dt} = F - \lambda_f T_f - \varepsilon \gamma (T_f - T_s), \quad (5)$$

$$C_s \frac{dT_s}{dt} = \gamma (T_f - T_s), \quad (6)$$

where  $T_f$  and  $T_s$  are the temperature anomalies of the fast (atmosphere + upper ocean) and slow (deep ocean) boxes. The observable global mean surface temperature is  $T = T_f$ , the upper-box temperature. Note that these boxes are effective constructs:



155 the heat capacities  $C_f$  and  $C_s$  are fitted parameters that implicitly absorb structural features such as mixed-layer depth and its variation with climate state. Because the model is linear, these effective parameters could in principle shift at a different base-state forcing level; in practice, however, the close fit to both temperature and CO<sub>2</sub> trajectories across the full flat10 ramp of increasing emissions (fig. 1) suggests the linear approximation is adequate for the forcing range considered here.

The equilibrium response follows directly from eqs. (5) and (6): setting  $dT_f/dt = dT_s/dt = 0$  forces  $T_f = T_s$ , so the heat-  
160 exchange terms in both equations vanish and the surface energy balance reduces to  $\lambda_f T = F$ . The equilibrium climate sensitivity is therefore  $ECS = F_{2\times}/\lambda_f$ , independent of  $\varepsilon$ ,  $\gamma$ , and  $C_s$ . This result will be useful below when relating the transient and equilibrium responses.

The efficacy parameter captures the fact that the spatial pattern of warming associated with deep-ocean heat uptake differs from the equilibrium warming pattern, and these different patterns excite different radiative feedbacks (Winton et al., 2010; Armour et al., 2013). When  $\varepsilon > 1$ , deep-ocean heat uptake is associated with warming patterns (e.g. Southern Ocean, eastern tropical Pacific) that produce stronger negative radiative feedbacks than the global-mean feedback  $\lambda_f$ , meaning that the transient climate response is suppressed relative to what the equilibrium sensitivity would imply. Setting  $\varepsilon = 1$  recovers the standard two-box model. We note that the substitution  $\gamma' = \varepsilon\gamma$ ,  $C'_s = \varepsilon C_s$  could further reduce eqs. (5) and (6) to the standard  
165 four-parameter two-box form, so when only temperature is constrained,  $\varepsilon$  does not add independent degrees of freedom to the fit. This degeneracy is broken when top-of-atmosphere radiative flux  $N = F - \lambda_f T_f - (\varepsilon - 1)\gamma(T_f - T_s)$  is also constrained, since  $N$  depends on  $\varepsilon$  independently of  $\gamma$ . We retain  $\varepsilon$  as an explicit parameter because it isolates the pattern effect as a distinct physical mechanism with its own observational literature (Winton et al., 2010; Armour et al., 2013), and because it appears transparently in the analytical ZEC expression derived in section 4, making the influence of the pattern effect on the thermal pipeline directly readable.

175 The eigenstructure of this system is most easily found by writing eqs. (5) and (6) in matrix form. Defining the state vector  $\mathbf{T} = (T_f, T_s)^\top$  and collecting the linear terms on the right-hand side, the homogeneous system is

$$\frac{d\mathbf{T}}{dt} = \mathbf{A}\mathbf{T} + \mathbf{b}F, \quad \mathbf{A} = \begin{pmatrix} -(\lambda_f + \varepsilon\gamma)/C_f & \varepsilon\gamma/C_f \\ \gamma/C_s & -\gamma/C_s \end{pmatrix}, \quad (7)$$

where  $\mathbf{b} = (1/C_f, 0)^\top$ . In the system matrix  $\mathbf{A}$ , the fast-box diagonal element combines radiative loss to space ( $\lambda_f/C_f$ ) with heat transfer to the deep ocean ( $\varepsilon\gamma/C_f$ ), while the slow-box diagonal reflects only the inter-box exchange ( $\gamma/C_s$ ); the  
180 off-diagonal elements return heat from one box to the other.

Although the two boxes are coupled (heat flows from the upper ocean into the deep ocean and vice versa), the system can be rewritten as the sum of two independent, non-interacting modes by finding the eigenvalues and eigenvectors of  $\mathbf{A}$  (the eigendecomposition approach is laid out in a step-by-step derivation in the Supplementary Material, Section S4). The practical payoff of this decomposition is substantial: instead of tracking two coupled temperatures ( $T_f$  and  $T_s$ ) that evolve together in a  
185 complicated way, we obtain two decoupled exponential decays, each with its own characteristic timescale and amplitude. The observable surface temperature  $T = T_f$  can then be written as a sum of these two independent eigenmode contributions. This is the mathematical step that ultimately lets us write ZEC as a transparent sum of thermal and carbon terms (section 4).



The eigenvalues of  $\mathbf{A}$  define two natural response timescales of the climate system: a fast mode (years to decades), representing the rapid adjustment of the atmosphere and upper ocean to a forcing change, and a slow mode (centuries), representing the gradual equilibration of the deep ocean. Because  $\mathbf{A}$  is a  $2 \times 2$  matrix, the eigenvalues follow from the trace and determinant:

$$\mu_{1,2} = \frac{1}{2} \left[ \text{tr}(\mathbf{A}) \pm \sqrt{\text{tr}(\mathbf{A})^2 - 4 \det(\mathbf{A})} \right], \quad (8)$$

with  $\text{tr}(\mathbf{A}) = -(\lambda_f + \varepsilon\gamma)/C_f - \gamma/C_s$  and  $\det(\mathbf{A}) = \lambda_f \gamma / (C_f C_s)$ . The discriminant can be rewritten as  $[(\lambda_f + \varepsilon\gamma)/C_f - \gamma/C_s]^2 + 4\varepsilon\gamma^2 / (C_f C_s)$ , which is always positive, so both eigenvalues are real. Since  $\text{tr}(\mathbf{A}) < 0$  and  $\det(\mathbf{A}) > 0$ , both eigenvalues are also negative, indicating that the system is stable: any perturbation decays back toward equilibrium. Although the exact eigenvalues follow from eq. (8), their physical content is more transparent after exploiting  $C_s \gg C_f$  in a Taylor expansion (cf. Seshadri, 2017); the corresponding timescales  $\tau_j^{\text{clim}} = -1/\mu_j$  are approximately

$$\tau_1^{\text{clim}} \approx \frac{C_f}{\lambda_f + \varepsilon\gamma}, \quad \tau_2^{\text{clim}} \approx \frac{C_s}{\gamma} \frac{\lambda_f + \varepsilon\gamma}{\lambda_f}. \quad (9)$$

We label the eigenmodes with numeric subscripts ( $j = 1$  for the fast mode,  $j = 2$  for the slow mode) to distinguish them from the physical box quantities, which retain the subscripts  $f$  (surface) and  $s$  (deep). The eigenmode timescales are close to, but not identical to, the box timescales  $C_f/(\lambda_f + \varepsilon\gamma)$  and  $C_s/\gamma$ , with the correspondence becoming exact in the limit  $C_s/C_f \rightarrow \infty$  (the eigenvectors and their relation to the box temperatures are derived explicitly in the Supplementary Material, Section S4). Physically, the fast timescale  $\tau_1^{\text{clim}}$  is set by how quickly the upper ocean adjusts to a change in forcing (typically 3–8 years), while the slow timescale  $\tau_2^{\text{clim}}$  reflects how long the deep ocean takes to reach thermal equilibrium (typically 200–600 years in two-box fits to centennial-scale CMIP experiments; Geoffroy et al., 2013). Coupled modelling studies that integrate to full equilibrium find deep-ocean adjustment timescales closer to millennial (Stouffer and Manabe, 2003; Li et al., 2013; Jeevanjee et al., 2025), so the two-box slow mode is best understood as an effective timescale for the centennial response rather than a true equilibration time. It is this slow mode that sustains the “warming in the pipeline” after emissions cease.

To see concretely how the decomposition works, consider the temperature response to a sudden (step) increase in radiative forcing. Rather than solving the full coupled system, we can write the response as a sum of two exponentials: a fast component that rises quickly and accounts for most of the warming within a few years, and a slow component that continues to grow over centuries as the deep ocean gradually warms. The weight of each component (how much warming it ultimately contributes) is given by the eigenmode fractions  $q_1$  and  $q_2$ , with  $q_1 + q_2 = 1$ . In typical climate models, the fast mode delivers roughly 60–80% of the equilibrium warming ( $q_1 \approx 0.6$ – $0.8$ ), while the remaining 20–40% resides in the slow mode and emerges only over centuries. It is this slow, unrealised fraction that constitutes the thermal “pipeline.”

### 2.2.1 Exponential-decay structure and the link to the carbon cycle

Because both eigenvalues are real and negative, the general solution of the two-box system is a sum of decaying exponentials: each eigenmode relaxes independently toward equilibrium with its own characteristic timescale  $\tau_j^{\text{clim}}$ . This is the same mathematical form as the carbon-cycle IRF (eq. (3)), where each uptake pool draws down atmospheric  $\text{CO}_2$  as an independent



220 exponential with timescale  $\tau_i$ . The parallel is not a coincidence: both models describe the relaxation of a linear system back toward equilibrium, decomposed into independent modes.

This shared exponential structure is what enables the analytical ZEC framework developed in section 4. After emissions cease, the thermal pipeline delivers unrealised warming as a sum of eigenmode decays, while the carbon sinks simultaneously remove CO<sub>2</sub> as a sum of uptake-pool decays. Because both responses are sums of exponentials, their competition can be written  
225 as a single closed-form expression in which each term pairs one thermal mode with one carbon pool.

An important constraint on the eigenvalue spectrum comes from the determinant of  $\mathbf{A}$ :  $\mu_1 \mu_2 = \det(\mathbf{A}) = \lambda_f \gamma / (C_f C_s)$ . Because the determinant depends only on the baseline climate parameters  $\lambda_f$ ,  $\gamma$ ,  $C_f$ , and  $C_s$  — and not on the efficacy  $\varepsilon$  — the product of the eigenvalues is independent of  $\varepsilon$ . Since the timescales are the negative reciprocals of the eigenvalues,  $\tau_j^{\text{clim}} = -1/\mu_j$ , their product follows from the determinant as

$$230 \quad \tau_1^{\text{clim}} \tau_2^{\text{clim}} = \frac{1}{\mu_1 \mu_2} = \frac{C_f C_s}{\lambda_f \gamma}. \quad (10)$$

This invariance means that increasing  $\varepsilon$  speeds up one mode while slowing the other: it widens the separation between fast and slow timescales but cannot change their product. The constraint is useful because it limits how efficacy can reshape the thermal response: a higher  $\varepsilon$  shifts weight toward the slow mode (by reducing  $q_1$ ) and lengthens its timescale, but the product rule ensures the fast mode simultaneously speeds up. In the ZEC framework, this redistribution determines the balance between  
235 pipeline warming (dominated by the slow mode) and the timescale of thermal decay relative to carbon uptake.

### 2.2.2 Quantifying thermal 'warming in the pipeline'

As noted above, the equilibrium sensitivity  $\text{ECS} = F_{2\times} / \lambda_f$  is independent of  $\varepsilon$ . To quantify the thermal pipeline (the unrealised warming fraction that drives ZEC), we need to know how much equilibrium warming each eigenmode misses during a transient forcing ramp. The step-response solution provides this building block: any forcing history (including the emissions ramp-up)  
240 can be constructed from superpositions of steps, and the step-response makes the eigenmode weights and timescales directly visible. The step-response temperature is

$$T_{\text{step}}(t) = \frac{F}{\lambda_{\text{eff}}} \left[ 1 - \sum_j q_j e^{-t/\tau_j^{\text{clim}}} \right], \quad (11)$$

where  $\lambda_{\text{eff}} = \lambda_f$  for the observable  $T = T_f$ . The Transient Climate Response (TCR) metric is derived using the standard 1 % yr<sup>-1</sup> CO<sub>2</sub> experiment, where there is a linearly ramping forcing over 70 years. This can be calculated by integrating  
245 equation 11 (cf. Millar et al., 2017):

$$\text{TCR} = \text{ECS} \left[ 1 - \sum_j q_j \frac{\tau_j^{\text{clim}}}{70} \left( 1 - e^{-70/\tau_j^{\text{clim}}} \right) \right], \quad (12)$$

The realised warming fraction  $r = \text{TCR}/\text{ECS}$  is the bracketed term in eq. (12). Each eigenmode contributes a correction proportional to  $(\tau_j^{\text{clim}}/70)(1 - e^{-70/\tau_j^{\text{clim}}})$ , which grows with timescale: modes much faster than 70 yr contribute negligibly, while modes slower than 70 yr contribute a correction close to  $q_j$ . Increasing  $\varepsilon$  lengthens  $\tau_2^{\text{clim}}$  (eq. (9)), boosting the slow



250 mode's correction and reducing  $r$ . The fast mode shortens simultaneously (the product  $\tau_1^{\text{clim}}\tau_2^{\text{clim}}$  is fixed; eq. (10)), but since  $\tau_1^{\text{clim}} \ll 70$  yr its contribution is already near zero and barely changes. The net effect is that higher efficacy enlarges the “warming in the pipeline” ( $1 - r$ ), which is the thermal driver of ZEC (a detailed analysis is provided in Supplementary Material, Section S4).

### 3 Fitting procedure

255 We fit the 11 model parameters (see Table 1) to ESM output from the flat10MIP experimental protocol (Sanderson et al., 2025a), in which CO<sub>2</sub> emissions are held constant at 10 PgC yr<sup>-1</sup> for 100 years, then set to zero. These experiments are part of the CMIP7 fast track (Dunne et al., 2025). Two experiments are used jointly:

1. **flat10**: continued emissions for 100–200 years (constrains the carbon cycle during buildup);
2. **flat10-zec**: emissions cease at year 100 (constrains the post-cessation carbon drawdown and thermal relaxation).

260 The fitting proceeds in two sequential stages to reduce the dimensionality of each optimisation problem.

#### 3.1 Stage 1: Carbon cycle parameters

Seven carbon cycle parameters  $\{f_\infty, f_1, \tau_1, f_2, \tau_2, \tau_3, \zeta\}$  are initially fitted by minimising the mean root mean square error (RMSE) of  $\Delta\text{CO}_2$  across both experiments:

$$\mathcal{L}_{\text{CO}_2} = (\text{RMSE}_{\text{flat10}} + \text{RMSE}_{\text{flat10-zec}}). \quad (13)$$

265 Equal weighting between the two experiments is somewhat arbitrary but reflects their comparable durations (100 years each) and treats both the buildup and drawdown phases as equally constraining. Fitting to both experiments simultaneously is important: the continued-emission trajectory constrains the uptake timescales and the airborne fraction during buildup, while the zero-emission trajectory constrains the post-cessation drawdown behaviour and the permanent fraction. The original ESM temperature output (smoothed with an 11-year running mean to remove interannual variability) is supplied to the carbon–climate  
270 feedback term (eq. (4)).

Optimisation uses L-BFGS-B (Byrd et al., 1995) with physically motivated bounds (table 1) and 10 initial-guess restarts to provide robustness against finding local minima.

#### 3.2 Stage 2: Climate parameters

With the carbon cycle parameters fixed, the fitted CO<sub>2</sub> trajectory from Stage 1 is converted to a forcing timeseries via eq. (1).

275 Four climate parameters  $\{C_f, C_s, \lambda_f, \gamma\}$  are then fitted by minimising the mean temperature RMSE across both experiments:

$$\mathcal{L}_T = \frac{1}{2}(\text{RMSE}_{\text{flat10}}^T + \text{RMSE}_{\text{flat10-zec}}^T), \quad (14)$$

where the target is the 11-year running-mean ESM temperature anomaly. For optimisation, L-BFGS-B with 5 restarts is used.



### 3.3 Stage 3: Ensemble generation

Deterministic fits provide a single parameter vector per ESM but do not quantify parametric uncertainty. To propagate un-  
280 certainty, we generate history-matched parameter ensembles using an adaptive Latin Hypercube Sampling approach with per-  
channel filtering.

For each ESM, the procedure operates in waves of 5 000 samples:

1. **Wave 1** draws samples from a broad Latin Hypercube Sampling across the full parameter bounds (table 1).
2. **Subsequent waves** use a hybrid strategy: 60% of samples are generated by local perturbation around known good  
285 parameter vectors (the deterministic fit plus any already-retained members), while the remaining 40% come from a  
focused Latin Hypercube Sampling with bounds tightened to the 2nd–98th percentile range of the top-scoring 10% of  
valid samples from previous waves. The perturbation scale decreases across waves (from 15% to 2% of the parameter  
range), shifting from broad exploration to fine-tuning.

Each candidate parameter vector is evaluated by running the forward model on both the flat10 and flat10-zec experiments and  
290 computing per-channel RMSEs for CO<sub>2</sub> and temperature separately. A sample is retained only if *both* its CO<sub>2</sub> RMSE and  
its temperature RMSE fall below 3× the corresponding deterministic-fit RMSE (with minimum floors of 1 ppm and 0.05 K  
to avoid degenerate thresholds where the deterministic fit is highly constrained). Sampling continues until 10 members are  
retained for each model.

Each accepted ensemble member defines a complete 12-parameter model from which derived quantities (ZEC, ECS, TCR  
295 etc.) can be computed (ECS,TCR are defined in Section 2.2.2, ZEC is defined in Section 4). Pooling ensembles across all  
ten ESMs yields a multi-model distribution that captures both parametric uncertainty (within a single ESM) and structural  
uncertainty (across ESMs).

## 4 Analytical framework

As such, ZEC is defined as the temperature anomaly relative to the temperature at cessation:

$$300 \quad ZEC(\Delta t) = T(t_0 + \Delta t) - T(t_0), \quad (15)$$

where  $t_0$  is the cessation time and  $\Delta t$  is the time elapsed since. Note that ZEC is implicitly conditional on the climate state at  
cessation (i.e. the cumulative emissions and warming level); the sensitivity of ZEC to these conditioning variables is explored  
in section 5.7. We decompose ZEC into a positive thermal component (unrealised warming delivery) and a negative carbon  
component (CO<sub>2</sub> drawdown cooling).

305 The derivation proceeds in three steps, each grounded in the flat10 experiment: (i) constant emissions at rate  $E_0 = 10 \text{ PgC yr}^{-1}$   
for  $t_0 = 100$  years build up atmospheric CO<sub>2</sub> and a forcing profile of increasing emissions; (ii) the two-box climate system  
responds to this forcing, leaving a thermal disequilibrium at cessation; (iii) after emissions are set to zero, the carbon sinks



continue drawing down CO<sub>2</sub>, reducing the forcing and partially offsetting the continued ocean warming. The analytical solution below tracks how  $E_0$  and  $t_0$  propagate through both the carbon cycle and the climate eigenstructure to determine the sign and magnitude of ZEC.

#### 4.1 Thermal and carbon decomposition

The matrix form of the equations (eqs. (7) and (8)) provides the foundation for an analytical decomposition of ZEC. While in Jenkins et al. (2022) these expressions were Taylor-expanded to determine aggregate responses over multi-decadal timescales, here we attempt to separate the post-cessation temperature evolution into two additive general formulas keeping the underlying physical parameters explicit: one that depends only on the climate eigenstructure, and one that couples the climate eigenmodes to the carbon cycle.

##### 4.1.1 Thermal component (unrealised warming)

During the emission phase, constant emissions  $E_0$  over  $t_0$  years build up atmospheric CO<sub>2</sub> and radiative forcing gradually, reaching  $F_0$  at cessation. Because each eigenmode responds on its own timescale, the faster mode (timescale  $\tau_1^{\text{clim}}$ , carrying fraction  $q_1$  of the equilibrium response) nearly equilibrates to the forcing while the slower mode (timescale  $\tau_2^{\text{clim}}$ , fraction  $q_2 = 1 - q_1$ ) accumulates a large deficit, the “warming in the pipeline.” Formally, the deficit of mode  $j$  at cessation is proportional to  $q_j (\tau_j^{\text{clim}}/t_0)(1 - e^{-t_0/\tau_j^{\text{clim}}})$  (Supplementary Material, eq. S3). To illustrate with typical values from the ESM ensemble: in the flat10 experiment with  $t_0 = 100$  yr, a fast mode of  $\tau_1^{\text{clim}} \sim 4$  yr has  $t_0/\tau_1^{\text{clim}} \approx 25$ , so the exponential vanishes and the deficit is only  $\sim 4\%$  of the fast mode’s equilibrium share, so this mode has kept pace with the ramp. By contrast, a slow mode of  $\tau_2^{\text{clim}} \sim 250$  yr has  $t_0/\tau_2^{\text{clim}} \approx 0.4$ , giving a deficit of  $\sim 80\%$  of its equilibrium share: most of the slow mode’s warming remains unrealised. This behaviour was described by Allen et al. (2022) and shown in their figure 1.

After emissions cease, if atmospheric CO<sub>2</sub> were held fixed, each mode would relax toward its equilibrium share at rate  $1/\tau_j^{\text{clim}}$ . The resulting warming delivery is the thermal component of ZEC (Supplementary Material, eq. S4):

$$\text{ZEC}_{\text{thermal}}(\Delta t) = T_{\text{eq}} \sum_j q_j \frac{\tau_j^{\text{clim}}}{t_0} (1 - e^{-t_0/\tau_j^{\text{clim}}}) (1 - e^{-\Delta t/\tau_j^{\text{clim}}}), \quad (16)$$

where  $T_{\text{eq}} = F_0/\lambda_{\text{eff}}$  is the equilibrium temperature change for the cessation forcing. This term is always positive, and is dominated by the slow eigenmode because  $\tau_2^{\text{clim}} \gg \tau_1^{\text{clim}}$ , which implies the slow mode retains most of the unrealised warming.

##### 4.1.2 Carbon component (CO<sub>2</sub> drawdown cooling)

Jenkins et al. (2022) take a similar impulse-response approach but immediately linearise over multi-decadal intervals; here we retain the full eigenmode–pool structure before simplifying. In reality, atmospheric CO<sub>2</sub> is not fixed after cessation: it declines as the ocean and land sinks continue to absorb carbon. The amount of CO<sub>2</sub> available for drawdown is set by the emission history: constant emissions  $E_0$  over  $t_0$  years load each uptake pool  $i$  (with fraction  $f_i$  and timescale  $\tau_i$ ; eq. (3)) with excess airborne carbon proportional to  $f_i \tau_i (1 - e^{-t_0/\tau_i})$ . After cessation, each pool relaxes exponentially on its own



timescale, drawing down CO<sub>2</sub> and reducing the radiative forcing. Linearising the logarithmic forcing around the cessation concentration  $C(t_0)$  (Supplementary Material, eq. S7), the forcing perturbation from pool  $i$  is  $\delta F_i \propto -E_0 f_i \tau_i (1 - e^{-t_0/\tau_i})(1 - e^{-\Delta t/\tau_i})/[C_0 + C(t_0)]$ . The emission rate  $E_0$  thus enters the carbon component directly: higher emissions produce more excess airborne CO<sub>2</sub> and hence a stronger drawdown forcing, but they also raise  $C(t_0)$ , which dilutes the forcing sensitivity per ppm (because of the logarithmic relationship). This partial compensation is why the carbon cooling scales sub-linearly with cumulative emissions.

The climate-carbon feedback  $\zeta$  (eq. (4)) further modulates the drawdown: a factor  $(1 - \zeta \bar{T})$ , where  $\bar{T} \approx T_0/2$  is the mean emission-phase temperature change, represents the reduced effective uptake fractions by carbon sinks in a warmer world. The carbon component is always negative (cooling) and depends on the removable airborne fraction  $1 - f_\infty$  (where  $f_\infty$  is the permanent airborne fraction from eq. (3)) as well as the drawdown rate and this feedback-mediated sink efficiency.

The climate system's response to this gradually building forcing change is again governed by the two eigenmodes: the fast mode responds rapidly (on year-to-decade timescales), while the slow mode accumulates a partial response over centuries. The interplay between carbon timescale  $\tau_i$  and climate timescale  $\tau_j^{\text{clim}}$  is captured by a *mixed response function*  $\Phi_{ij}(\Delta t)$ , which gives the fraction of the equilibrium temperature response of climate mode  $j$  that has been realised at time  $\Delta t$  when driven by a forcing that itself builds exponentially on timescale  $\tau_i$ . Physically,  $\Phi_{ij}$  arises from convolving the climate mode's impulse response (a decaying exponential on timescale  $\tau_j^{\text{clim}}$ ) with the gradually saturating forcing from carbon pool  $i$  (Supplementary Material, eq. S5):

$$\Phi_{ij}(\Delta t) = 1 - \frac{\tau_j^{\text{clim}}}{\tau_j^{\text{clim}} - \tau_i} e^{-\Delta t/\tau_j^{\text{clim}}} + \frac{\tau_i}{\tau_j^{\text{clim}} - \tau_i} e^{-\Delta t/\tau_i}. \quad (17)$$

The function interpolates between two familiar limits depending on which timescale is slower. When the carbon pool draws down CO<sub>2</sub> quickly relative to the climate mode ( $\tau_i \ll \tau_j^{\text{clim}}$ ), the forcing looks like a step change and  $\Phi_{ij}$  reduces to the climate mode's own step-response,  $1 - e^{-\Delta t/\tau_j^{\text{clim}}}$ . When the climate mode is fast relative to the drawdown ( $\tau_i \gg \tau_j^{\text{clim}}$ ), the temperature tracks the forcing quasi-statically and the response timescale is set by the carbon pool, giving  $1 - e^{-\Delta t/\tau_i}$ . At long times both exponentials decay and  $\Phi_{ij} \rightarrow 1$  (full equilibration).

In practice, the pairing of the slowest carbon pool with the slowest climate mode ( $\Phi_{3,2}$  in our three-pool, two-mode framework) tends to dominate ZEC<sub>50</sub> because both timescales are long relative to the 50-year horizon: the fast climate mode has already largely equilibrated, while the slow mode continues to deliver pipeline warming as the slow carbon pool gradually draws down CO<sub>2</sub>. At asymptotically long times ( $\Delta t \gg \tau_2^{\text{clim}}$ ), all modes fully equilibrate and ZEC<sub>∞</sub> is determined solely by the permanent airborne fraction  $f_\infty$  (i.e., how much CO<sub>2</sub> remains after all removable pools have drawn down). However, the 100-year duration of the flat10-zec experiment is insufficient to reliably distinguish slow carbon uptake from truly permanent retention, making  $f_\infty$  poorly constrained: the fitted ensemble values span a wide range (Supplementary Material, section S7 and table S1), reflecting the fact that the slow carbon pool ( $\tau_3 \sim 200\text{--}400$  years) has not yet equilibrated by year 100. The implied ZEC<sub>∞</sub> is consequently dominated by slow-process carbon uptake and is itself highly uncertain. The extent to which ZEC<sub>50</sub> approaches this asymptotic limit thus varies considerably across ESMs depending on their thermal and carbon timescales, but across most models, significant adjustment occurs within the first half-century.



## 4.2 Levels of approximation

These approximate formulae are not used for fitting or for computing ZEC values reported elsewhere in the paper; they serve to identify which physical parameters control ZEC and why. The full analytical ZEC expression has many terms (two climate eigenmodes  $\times$  three carbon pools), so we present it at successive levels of simplification. *Level 1* retains all 12 parameters and all mode–pool pairings, giving maximum accuracy; *Level 2* collapses the sums into aggregate climate metrics ( $r$ ,  $\tau_2^{\text{clim}}$ ,  $\tau_c$ ,  $q_{\text{eff}}$ ), yielding a compact two-term formula suitable for physical interpretation. Level 2 is designed for 50–100 yr post-cessation (needs  $dt \ll \tau_2^{\text{clim}}$ , where  $\tau_2^{\text{clim}}$  is 200–500 yr, see Supplemental section S2). Level 1 is valid at all times.

### 4.2.1 The Level 1 formula

Summing the thermal and carbon components over both eigenmodes ( $j \in \{f, s\}$ , with timescales  $\tau_1^{\text{clim}}$ ,  $\tau_2^{\text{clim}}$  and response fractions  $q_1$ ,  $q_2$ ; eq. (9)) and all three carbon pools ( $i = 1, 2, 3$ , with fractions  $f_i$  and timescales  $\tau_i$ ; eq. (3)) gives

$$\text{ZEC}(\Delta t) = \underbrace{T_{\text{eq}} \sum_j q_j \frac{\tau_j^{\text{clim}}}{t_0} (1 - e^{-t_0/\tau_j^{\text{clim}}}) (1 - e^{-\Delta t/\tau_j^{\text{clim}}})}_{\text{thermal (unrealised warming, } > 0)} + \underbrace{\frac{1}{\lambda_f} \sum_i \delta F_{i,\infty} \sum_j q_j \Phi_{ij}(\Delta t)}_{\text{carbon (CO}_2 \text{ drawdown cooling, } < 0)} \quad (18)$$

where  $T_{\text{eq}} = F_0/\lambda_f$  is the equilibrium temperature for the cessation forcing,  $\delta F_{i,\infty}$  is the equilibrium forcing perturbation from pool  $i$  (Supplementary Material, eq. S7), and  $\Phi_{ij}$  is the mixed response function defined above. This expression uses all 12 model parameters (table 1) and reproduces numerical  $\text{ZEC}_{50}$  to within 0.018–0.064 K across ten ESMs, corresponding to  $\sim 8$ –30% of the individual thermal or carbon component magnitudes ( $\sim 0.2$  K each). The discrepancy includes a small systematic positive bias, the origin of which is discussed in section 5.2 (Supplementary Material, Section S1 gives the full derivation). The sign of ZEC is set by the competition between these two terms: ZEC is positive when eigenmode relaxation delivers heat faster than the carbon sinks can cool, and negative when the strength of the carbon sinks dominates.

### 4.3 A simplified expression for ZEC (Level 2)

By replacing the eigenmode sums with aggregate climate metrics: the realised warming fraction  $r = \text{TCR}/\text{ECS}$  (which absorbs  $q_1$ ,  $q_2$ ,  $\tau_1^{\text{clim}}$ , and  $\lambda_f$  into a single ratio) and an effective carbon timescale  $\tau_c$  that condenses the three pool timescales into one, the Level 1 expression simplifies (Supplementary Material, Section S2). The thermal term further linearises the slow-mode exponential for  $\Delta t \ll \tau_2^{\text{clim}}$ , the same multi-decadal expansion that underlies the RAZE derivation of Jenkins et al. (2022) (see section 4.4 for the explicit connection). However, we propose here that the carbon term should not be linearised:

$$\text{ZEC}(\Delta t) \approx T_0 \underbrace{\frac{(1/r - 1) \Delta t}{\tau_2^{\text{clim}}}}_{\text{thermal warming}} + \underbrace{\frac{\delta F_{\text{tot},\infty}}{\lambda_{\text{eff}}} q_{\text{eff}} (1 - e^{-\Delta t/\tau_c})}_{\text{carbon drawdown cooling}} \quad (19)$$

This expression is valid for  $\Delta t \approx 50$ –200 years, where the effective single-timescale approximation holds and the permanent fraction  $f_\infty$  has not yet fully dominated the carbon uptake term. The thermal term is linearised in  $\Delta t/\tau_2^{\text{clim}}$  (accurate because



400  $\tau_2^{\text{clim}} \sim 200\text{--}500 \text{ yr} \gg \Delta t$ ), while the carbon term retains the full exponential because  $\tau_c \sim 30\text{--}80 \text{ yr}$  is comparable to the evaluation horizon  $\Delta t$ . The parameters are:

- $T_0$  is the warming at cessation;
- $r = \text{TCR}/\text{ECS}$  is the realised warming fraction (the ratio of the transient response to the equilibrium sensitivity). Although TCR and ECS are conventionally defined in concentration-driven experiments, their ratio  $r$  is a property of the energy balance eigenstructure and is independent of the forcing protocol;
- 405 –  $\tau_2^{\text{clim}}$  is the slow climate eigenmode timescale;
- $\delta F_{\text{tot},\infty}/\lambda_{\text{eff}}$  is the equilibrium cooling that would result from complete drawdown of all non-permanent airborne  $\text{CO}_2$  (Supplementary Material, eq. S7; always negative);
- $q_{\text{eff}}$  is the effective climate response fraction (defined below);
- $\tau_c$  is the effective carbon uptake timescale (an emission-weighted mean of the pool timescales, Supplementary Material, eq. S6).
- 410

The effective climate response fraction accounts for the slow climate mode’s partial equilibration to the gradually building drawdown forcing. Because  $\text{CO}_2$  is drawn down gradually rather than instantaneously, the slow climate mode has time to partially track the evolving forcing, amplifying the temperature response beyond what the fast mode alone would produce:

$$q_{\text{eff}}(\Delta t) = q_1 + q_2 \frac{\Phi(\Delta t, \tau_c, \tau_2^{\text{clim}})}{1 - e^{-\Delta t/\tau_c}}, \quad (20)$$

415 where  $\Phi$  is the mixed response function (Supplementary Material, eq. S5). Without this correction ( $q_{\text{eff}} \rightarrow q_1$ ), the carbon cooling is systematically underestimated by  $\sim 0.030 \text{ K}$ . The correction captures the fact that the carbon drawdown forcing builds gradually on timescale  $\tau_c$ , giving the slow climate mode time to partially equilibrate. In practice,  $q_{\text{eff}}$  interpolates between  $q_1$ , where only the fast mode responds to the drawdown, and  $q_1 + q_2 = 1$  corresponding to the limit where both modes fully equilibrate. The slower the carbon drawdown relative to the deep-ocean timescale, the closer  $q_{\text{eff}}$  approaches unity and  
420 the larger the carbon cooling term becomes. This correction is small but systematic.

#### 4.4 Physical interpretation

Equation (19) expresses ZEC as the sum of two competing terms:

1. **Thermal warming:**  $T_0(1/r - 1) \cdot \Delta t/\tau_2^{\text{clim}}$ . The unrealised warming “pipeline”  $T_0(1/r - 1) = T_{\text{eq}} - T_0$ , delivered at rate  $\Delta t/\tau_2^{\text{clim}}$ . Models with low  $r$  (large pipeline) or short  $\tau_2^{\text{clim}}$  (fast delivery) have larger thermal ZEC.
- 425 2. **Carbon drawdown cooling:**  $(\delta F_{\text{tot},\infty}/\lambda_{\text{eff}}) q_{\text{eff}}(1 - e^{-\Delta t/\tau_c})$ . The equilibrium drawdown cooling, scaled by the effective climate response fraction and the fraction of drawdown realised by time  $\Delta t$ . A larger removable  $\text{CO}_2$  fraction ( $1 - f_{\infty}^{\text{eff}}$ , i.e. the share of emitted  $\text{CO}_2$  that sinks can eventually absorb), shorter  $\tau_c$ , and larger  $q_{\text{eff}}$  all strengthen the cooling.



Both terms depend on the warming at cessation,  $T_0$ : the thermal term explicitly through the pipeline  $T_{\text{eq}} - T_0$ , and the carbon  
430 term implicitly through the accumulated  $\text{CO}_2$  that determines  $\delta F_{\text{tot},\infty}$ . Because these dependencies differ (the thermal pipeline  
scales with  $T_{\text{eq}}$  while the drawdown fraction saturates), ZEC is not simply proportional to  $T_0$ , and the thermal–carbon balance  
shifts with emission rate and duration (fig. 9).

ZEC<sub>50</sub>  $\approx 0$  when these two rates balance, the condition under which TCRE remains approximately constant through net zero.  
The near-compensation observed across ESMs is unlikely to be mere coincidence. It may reflect a partial constraint from  
435 the physical coupling of ocean heat and carbon uptake processes (the ‘heat-carbon nexus’; S  ferian et al. 2024), which share  
common timescales set by ocean ventilation and mixing. However, our simple model structure treats these as independent  
elements, and several factors complicate this interpretation: land carbon uptake (which contributes substantially to the carbon  
term, particularly on multi-decade intervals) is not coupled to ocean heat uptake via the nexus; the post-cessation drawdown  
regime differs from the accumulation phase where the nexus is most clearly expressed; and the weak correlation between  
440 thermal and carbon components across our ensemble (fig. 5) suggests that other factors dominate the inter-model spread.

The carbon–climate feedback ( $\zeta$ ) pushes ZEC to become positive in two ways: it increases  $\text{CO}_2$  at cessation (raising  $F_0$  and  
hence the thermal term) and weakens post-cessation drawdown (the factor  $1 - \zeta\bar{T}$ ). Because  $\zeta$  enters the Level 2 carbon term  
only through the  $(1 - \zeta\bar{T})$  factor in  $\delta F_{\text{tot},\infty}$ , the feedback correction scales as  $\zeta T_0^2$ ; in the flat10 regime ( $T_0 \sim 1$  K) it is modest  
( $\sim 0.01$ – $0.02$  K), but would grow roughly tenfold for  $T_0 \sim 3$ – $4$  K. The role of ocean heat uptake efficacy ( $\epsilon$ ), which introduces a  
445 structural asymmetry between the buildup and relaxation phases, is analysed in Supplementary Material Section S3. Extension  
of the Level 2 formula to longer post-cessation timescales is discussed in Supplementary Material Section S4.

#### 4.4.1 Connection to RAZE

Jenkins et al. (2022) define the Remaining Adjustment to Zero Emissions (RAZE) as the fractional rate at which global tem-  
perature evolves after emissions cease: positive RAZE implies continued warming (and hence residual emission removals are  
450 needed to halt warming), while negative RAZE implies cooling (temperature overshoots before declining). In their FaIR ensem-  
ble, the RAZE spans  $-0.24$  to  $+0.17\% \text{ yr}^{-1}$ , determining the residual  $\text{CO}_2$  emission rate compatible with stable temperature.  
Their derivation further splits the RAZE parameter into two additive terms ( $\text{RAZE} = \rho_F - \rho_E$ ), with the former component  
defining the rate of adjustment of the thermal cycle to historical forcing after net zero ( $\rho_F$ ), and the latter defining the rate of  
adjustment of the carbon cycle to historical emissions after net zero ( $\rho_E$ ).

455 Equation (19) makes the physical content of these two RAZE components explicit. Identifying the thermal adjustment rate as  
 $\rho \approx (1/r - 1)/\tau_2^{\text{clim}}$ , the rate at which unrealised thermal pipeline warming is delivered (as shown by Allen et al. (2022)); this  
corresponds to the limiting case of  $\rho_F$  under the assumption that negligible slow-mode thermal re-equilibration has occurred  
over multi-decadal timescales. The effective carbon uptake rate  $1/\tau_c^* \approx -(\delta F_{\text{tot},\infty}/\lambda_{\text{eff}}) q_{\text{eff}}/(T_0 \tau_c)$  as the fractional rate  
at which  $\text{CO}_2$  drawdown cooling accrues. Our Level 2 expression thus decomposes the RAZE into its constituent physical  
460 components and highlights their parameter dependencies. Evaluating these rates across our pooled probabilistic ensemble  
gives a thermal adjustment rate  $\rho = +0.21 \pm 0.22\% \text{ yr}^{-1}$  and an effective carbon uptake rate  $1/\tau_c^* = +0.13 \pm 0.13\% \text{ yr}^{-1}$   
(mean  $\pm$  one standard deviation), yielding a RAZE of  $+0.08 \pm 0.27\% \text{ yr}^{-1}$  (5th–95th percentile:  $-0.33$  to  $+0.51\% \text{ yr}^{-1}$ ). This



range encloses the Jenkins et al. interval ( $-0.24$  to  $+0.17\% \text{ yr}^{-1}$ ) but is substantially wider, reflecting both the propagation of full parameter uncertainty per ESM and the structural diversity of ten independent models, and an ensemble which is not  
465 constrained against historical observations (as the FaIR ensemble’s underlying parameter distributions were in Jenkins et al. (2022)).

## 5 Results

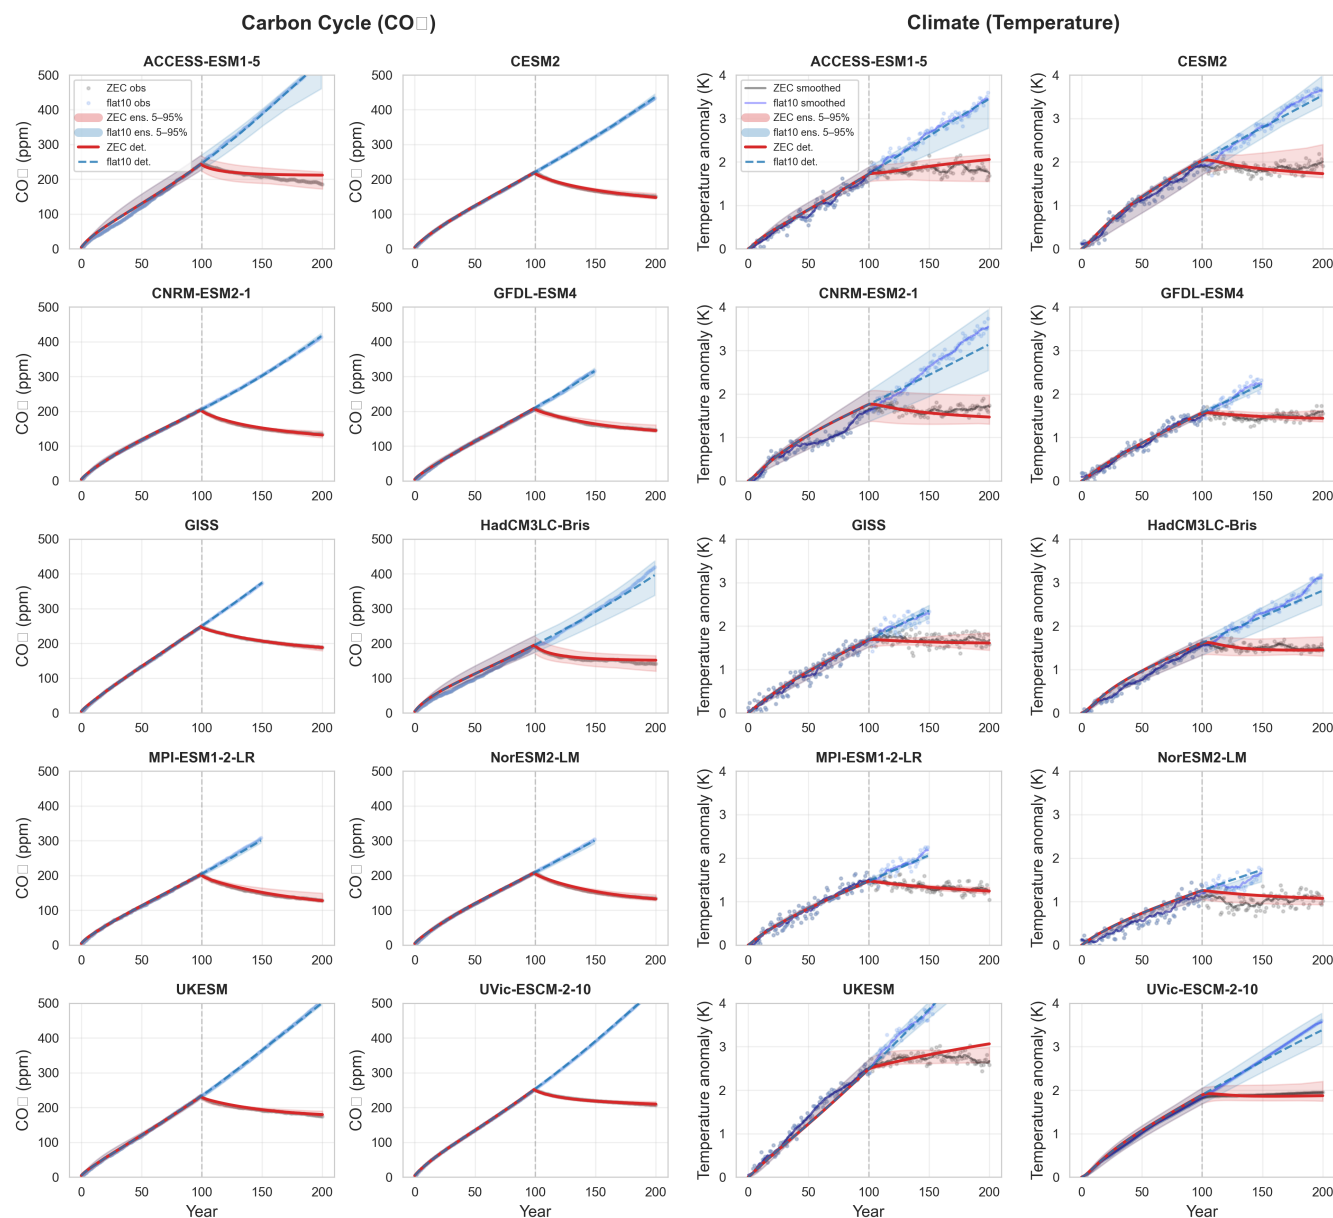
### 5.1 Model fits and parameter uncertainty

We fit the 12-parameter model to output from ten ESMs participating in the flat10MIP experiment (Sanderson et al., 2025a):  
470 ACCESS-ESM1-5 (Ziehn et al., 2020), CESM2 (Danabasoglu et al., 2020), CNRM-ESM2-1<sup>1</sup>, GFDL-ESM4 (Dunne et al., 2020), GISS (Kelley et al., 2020), HadCM3LC-Bris (Valdes et al., 2017), MPI-ESM1-2-LR (Mauritsen et al., 2019), NorESM2-LM (Seland et al., 2020), UKESM (Sellar et al., 2019), and UVic-ESCM-2.10 (Mengis et al., 2020). Figure 1 shows the deterministic fits for all ten models. CO<sub>2</sub> trajectories are reproduced with RMSE  $< 3$  ppm for most models, and temperature trajectories reproduced with RMSE  $< 0.050$  K after 11-year smoothing. The three-pool carbon cycle captures both the buildup  
475 and drawdown phases, while the two-box climate model reproduces the transient warming and post-cessation relaxation.

Although the model has 12 fitted parameters, the training data comprise  $\sim 200$  annual time points each for CO<sub>2</sub> and temperature across two experiments ( $\sim 800$  independent data points per ESM), giving a data-to-parameter ratio well above 60:1. Moreover, the parameters enter through a coupled ODE system whose structure severely constrains the fitted functions’ degrees of freedom: the carbon cycle must simultaneously reproduce the buildup trajectory, the drawdown trajectory, and the transition  
480 between them, while the climate parameters must be consistent with both transient and post-cessation temperature evolution. The physically motivated prior bounds (table 1) further restrict the parameter space. Taken together, the model is not in the high-parameter, low-data regime that would invite overfitting.

Figure 2 shows the posterior parameter distributions from the history-matched ensembles for each of the ten ESMs individually, plus the pooled multi-model distribution. The pooled ensemble combines all ensemble members across models, capturing  
485 both parametric uncertainty (within-model spread) and structural uncertainty (across-model differences). Across the pooled ensemble ( $n = 100$ ; 10 members per model), fitted ECS has a median of 3.8 K (5th–95th percentile: 2.1–12.4 K), demonstrating that the upper bound for equilibrium warming response to a forcing is only weakly constrained by the experiments. The slow climate eigenmode timescale  $\tau_2^{\text{clim}}$  (derived from the fitted  $C_f$ ,  $C_s$ ,  $\lambda_f$ , and  $\gamma$ ) has a pooled median of  $\sim 490$  years (5th–95th: 60–2230 years), reflecting the difficulty of constraining deep-ocean adjustment from the 200-year flat10 experi-  
490 ment. The carbon–climate feedback  $\zeta$  is relatively tightly constrained with a median of  $0.17 \text{ K}^{-1}$  (5th–95th:  $0.11$ – $0.20 \text{ K}^{-1}$ ). The effective permanent airborne fraction  $f_\infty$  has a median of 0.11 (5th–95th: 0.00–0.23), reflecting substantial uncertainty in whether a genuinely permanent fraction of emitted CO<sub>2</sub> persists on the timescales resolvable by the experiment.

<sup>1</sup>The flat10 simulations labelled CNRM-ESM2-1 were performed with an updated model version, CNRM-ESM2-2 (S  ferian et al., 2019; Bossert et al., 2025); we retain the original label for consistency with the flat10MIP data archive.



**Figure 1.** Deterministic model fits to ten ESMs from the flat10MIP ensemble (Sanderson et al., 2025a). Left: atmospheric CO<sub>2</sub> anomaly (ppm). Right: global mean surface temperature anomaly (K). Grey dots: raw ESM simulation output. Solid lines (temperature panels): 11-year running-mean smoothed ESM output. Dashed lines: fitted least-squares model (red: flat10-ZEC phase; blue: flat10 phase). Vertical grey line: cessation at year 100. Coloured regions: 5th–95th percentiles of fitted ensemble range.



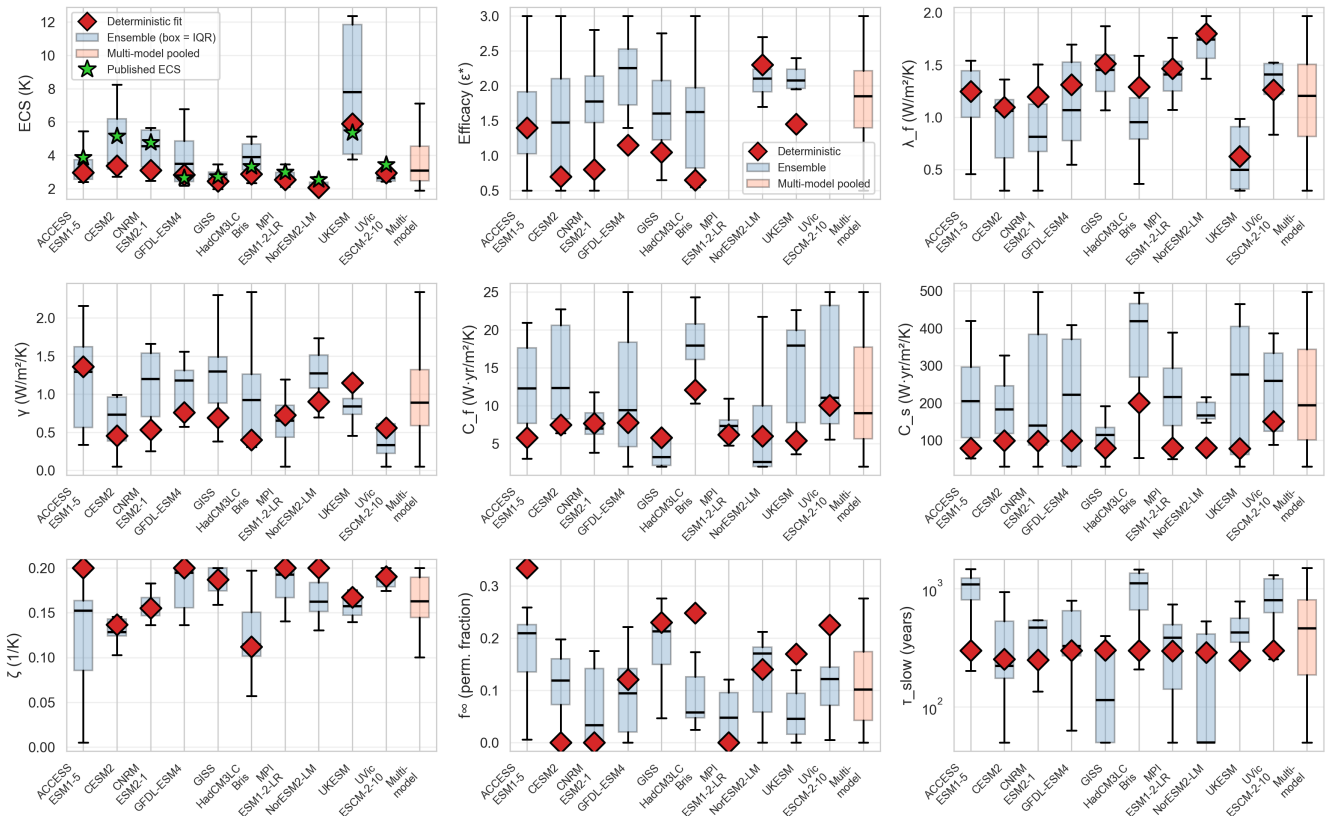
Two features of fig. 2 merit discussion. First, the deterministic fit (red diamond) does not always lie near the centre of the ensemble distribution and sometimes falls outside the 5th–95th percentile range (e.g.,  $C_s$  for three models,  $f_\infty$  for three models). This is expected: the deterministic fit minimises the *joint* CO<sub>2</sub>-plus-temperature cost and finds a single best-fitting parameter vector, whereas the ensemble retains all samples whose per-channel RMSEs fall below  $3\times$  the deterministic value. With 12 free parameters and strong trade-offs between them (e.g., a larger  $C_s$  can partly compensate for a smaller  $\lambda_f$ , or a larger  $f_\infty$  can mimic a slower  $\tau_3$ ), many different parameter combinations produce acceptable fits. The ensemble explores this full region of near-equivalent solutions, so for any single parameter the spread can be asymmetric and shifted relative to the optimum, so the single best fit and the typical ensemble member need not agree.

Second, the slow carbon timescale  $\tau_3$  exhibits strikingly consistent deterministic values across models (standard deviation  $\sim 16$  years; all ten fits cluster near  $\sim 300$  years) yet wide and model-dependent ensemble spreads (interquartile ranges of 130–855 years). The explanation is a practical identifiability problem: the 200-year experiment constrains timescales shorter than  $\sim 200$  years well, but for  $\tau_3 \gtrsim 200$  years the CO<sub>2</sub> trajectory over the observable window becomes nearly insensitive to the exact value: a timescale of 300 and 1000 years produce similar drawdown curves over this horizon. The deterministic optimiser consistently finds the sharp cost minimum near  $\sim 300$  years, but the RMSE-threshold ensemble reveals that the cost surface is essentially flat beyond this point, so the ensemble fills much of the prior range [50, 1500] years. The degree of ensemble spread varies across models depending on how well other parameters (particularly  $f_\infty$  and  $f_2$ ) compensate: when  $f_\infty$  is tightly constrained (e.g., near zero for CESM2),  $\tau_3$  becomes better identified because it must account for more of the residual drawdown. More generally, the carbon cycle pool fractions ( $f_1$ ,  $f_2$ ,  $f_\infty$ ) and their associated timescales are jointly underdetermined: trading uptake between pools while co-adjusting timescales leaves the aggregate CO<sub>2</sub> trajectory nearly unchanged, producing wide marginal distributions even when the predicted observables are tightly constrained.

## 5.2 Analytical validation

Before examining probabilistic ensembles, we validate the analytical expressions from section 4 against the numerical forward model using the deterministic fits (fig. 3).


The Level 1 expression (Supplementary Material, eq. S11) reproduces numerical ZEC<sub>50</sub> to within  $+0.018$  to  $+0.064$  K across all ten models (MAE  $\sim 0.029$  K), amounting to  $\sim 8$ – $30\%$  of either the thermal or carbon component individually ( $\sim 0.2$  K each); percentage errors relative to the near-zero total ZEC are uninformative given the compensation. In context,  $0.029$  K is less than  $5\%$  of the  $\sim 0.6$  K inter-model spread, and expressed as a RAZE-equivalent rate ( $\sim 0.03\% \text{ yr}^{-1}$  at  $T_0 \approx 2$  K) it is an order of magnitude smaller than the inter-model range reported by Jenkins et al. (2022). Omitting the  $q_{\text{eff}}$  correction (eq. (20)) and using  $q_1$  alone in the carbon term roughly doubles the error (MAE  $\sim 0.060$  K), but restoring  $q_{\text{eff}}$  brings the Level 2 MAE back to  $\sim 0.026$  K, comparable to Level 1. The systematic positive bias of  $\sim 0.025$  K is shared across all three approximation levels and all models. It arises because the analytical derivation assumes forcing ramps linearly during the emission phase, whereas the actual forcing is concave (logarithmic in CO<sub>2</sub>, which itself rises sub-linearly as sinks strengthen); this overestimates the forcing at cessation and hence the thermal deficit. A secondary contribution comes from the linearisation of the slow-mode exponential (a small effect,  $\sim 0.003$  K for typical  $\tau_2^{\text{clim}} \approx 300$  yr), which assumes negligible slow-mode re-



**Figure 2.** Ensemble parameter distributions for ten ESMs plus pooled multi-model distribution. Box plots show the interquartile range (IQR) for each model’s history-matched ensemble (blue) and the pooled multi-model ensemble (coral, rightmost column); whiskers extend to the 5th and 95th percentiles. Plots show results for, top row: ECS,  $\epsilon$ ,  $\lambda_f$ , middle row:  $\gamma$ ,  $C_f$ ,  $C_s$ , bottom row:  $\zeta$ ,  $f_\infty$ ,  $\tau_3$ . Red diamonds mark deterministic fits; green stars (ECS panel) show published ECS values. The multi-model distribution captures both parametric uncertainty within individual ESMs and structural uncertainty across models, yielding substantially wider ranges than single-model ensembles. Key features: pooled ECS has median 3.8 K (5th–95th percentile: 2.1–12.4 K; encompasses published values); efficacy  $\epsilon$  has pooled median 2.1 (5th–95th: 0.5–3.0; available for 8 models, as MPI-ESM1-2-LR and UVic-ESCM-2-10 lack the required top-of-atmosphere radiation data); the wide  $\epsilon$  ranges reflect the  $\gamma$ – $\epsilon$  degeneracy inherent in diagnosing efficacy from centennial experiments, but this has limited impact on ZEC itself (section 6.3); carbon–climate feedback  $\zeta$  median  $0.17 \text{ K}^{-1}$  (5th–95th:  $0.11$ – $0.20 \text{ K}^{-1}$ ); permanent airborne fraction  $f_\infty$  median 0.11 (5th–95th:  $0.00$ – $0.23$ ). Diamonds show best deterministic fits per models, blue bar show ensemble fit IQR, coral bar shows the pooled multi-model IQR, in the top left plot star symbols show published ECS values.



530 equilibration during the emission phase. In principle the bias could be reduced by using the actual concave forcing trajectory in the analytical integration, but doing so sacrifices the closed-form simplicity that makes the Level 2 expression (eq. (19)) physically transparent. Since the offset is small relative to the component magnitudes, systematic in sign, and well understood in origin, we retain the linear-ramp approximation and note the resulting positive offset when interpreting results.



**Figure 3.** Analytical versus numerical  $ZEC_{50}$  for ten ESMs at three levels of approximation. Level 1 (full parametric):  $MAE \approx 0.029$  K. Level 2 with  $q_1$  only:  $MAE \approx 0.060$  K. Level 2 with  $q_{eff}$ :  $MAE \approx 0.026$  K. Dashed line: 1:1. The  $q_{eff}$  correction eliminates the systematic bias of the fast-mode-only approximation.

### 5.3 ZEC decomposition and uncertainty

535 Figure 4 presents the pooled multi-model distributions of total, thermal, and carbon ZEC at 50 and 100 years post-cessation. In every model the thermal component is positive and the carbon component negative, so some degree of compensation is guaranteed by their opposite signs. The thermal component grows monotonically as the deep ocean delivers pipeline warming; the carbon component (in absolute value) also grows as  $CO_2$  drawdown accumulates. The degree of compensation varies substantially across models, however. For the deterministic fits, the residual total ranges from  $-0.19$  K (CNRM-ESM2-1, carbon dominated) to  $+0.33$  K (UKESM, thermal dominated), with eight of ten models showing negative (carbon-dominated)  $ZEC_{50}$ . The majority of models are carbon-dominated (negative ZEC: CESM2, CNRM-ESM2-1, GFDL-ESM4, GISS, HadCM3LC-Bris, MPI-ESM1-2-LR, NorESM2-LM, UVic-ESCM-2.10), with only two models showing thermal-dominated (positive) ZEC (ACCESS-ESM1-5, UKESM). Notably, models with coupled carbon–nitrogen cycles (ACCESS-ESM1-5, CESM2, NorESM2-LM, UKESM) appear in both groups, suggesting that nitrogen limitation does not systematically determine the sign of ZEC, contrary to the hypothesis of MacDougall et al. (2020), once the thermal and carbon contributions are separated. Consistent with the Level 2 formula (eq. (19)), the thermal-dominated group has a lower realised warming fraction (mean  $r = 0.57$ , i.e. more unrealised pipeline warming) and higher ECS (mean 4.4 K), while the carbon-dominated group has higher  $r$  (mean 0.74) and lower ECS (mean 2.8 K). The near-zero  $ZEC_{50}$  is a property of the *multi-model mean*, not of every individual model: the

540

545



pooled thermal component averages  $+0.216$  K (std  $0.137$  K) and the carbon component averages  $-0.231$  K (std  $0.062$  K), yielding a multi-model mean of only  $-0.014$  K. Because the two components are nearly uncorrelated across the ensemble ( $\rho \approx 0$ ; fig. 5), their variances *add* rather than cancel, and the total standard deviation ( $0.145$  K; 5th–95th percentile:  $-0.206$  to  $+0.246$  K) exceeds that of either component individually. This partial compensation of multi-model means, without reduction  
550 in spread, reflects the opposite signs of the two terms rather than any systematic co-variation that constrains their sum. Whether the compensation of means reflects a physical coupling between ocean heat and carbon uptake is discussed in section 4.4. The carbon component exhibits roughly half the variance of the thermal component, indicating that inter-model spread in ZEC is dominated by disagreement over the thermal term.

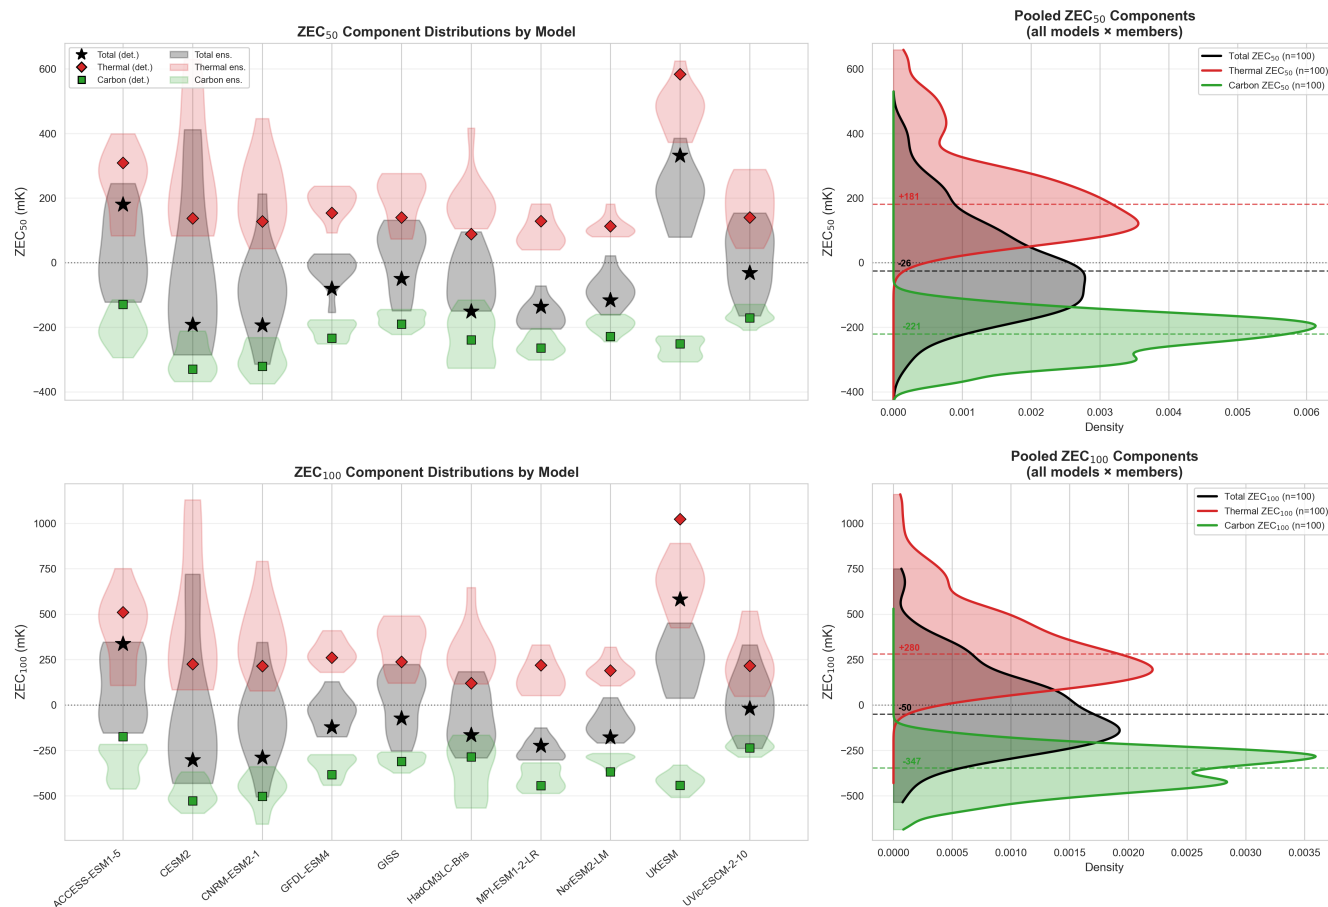
At  $ZEC_{100}$ , the total distribution widens (mean  $-0.022$  K, std  $0.221$  K) as the thermal component grows to  $+0.338$  K while  
555 the carbon component reaches  $-0.360$  K. The compensation persists but the wider spread reflects the greater sensitivity of the thermal term to the slow climate timescale at longer horizons. Individual models diverge further: the total  $ZEC_{100}$  now spans  $-0.30$  K (CESM2) to  $+0.58$  K (UKESM).

Figure 5 shows the joint distribution of the thermal and carbon components across the full multi-model probabilistic ensemble. At both  $ZEC_{50}$  and  $ZEC_{100}$ , the thermal component is universally positive (unrealised warming delivered by the ocean  
560 thermal pipeline) and the carbon component universally negative ( $CO_2$  drawdown into land and ocean sinks), confirming the partial compensation visible in fig. 4. However, in contrast to the strong linear relationship between cumulative carbon and heat uptakes documented for the warming phase by Brovkin et al. (2025), the relationship between the thermal and carbon components of ZEC is remarkably weak: the Pearson correlation across ensemble members is close to zero at both horizons. This decoupling arises because, after emission cessation, the thermal pipeline and the carbon drawdown are governed by largely  
565 independent timescale structures: the climate eigenmode timescales  $\tau_1^{\text{clim}}$  and  $\tau_2^{\text{clim}}$  for the thermal component, and the carbon pool decay timescales for the carbon component whose parameter uncertainties are only weakly correlated across the fitted ensembles. The joint  $CO_2$ -plus-temperature fitting can in principle imprint such correlations on the parameters, but the model structure does not enforce them (the thermal and carbon timescales are independent aside from  $\zeta$ ), so weak fitted correlations do not rule out stronger physical coupling in the underlying ESMs.

570 Despite this weak overall correlation, the deterministic fits reveal distinct model groupings in thermal–carbon space. The eight carbon-dominated ESMs span a range of negative  $ZEC_{50}$  values, from near zero (UVic-ESCM-2.10:  $-0.03$  K; GISS:  $-0.05$  K) to clearly negative (CNRM-ESM2-1:  $-0.19$  K; CESM2:  $-0.19$  K). ACCESS-ESM1-5 and UKESM form a separate thermally dominated cluster well above the compensation line at both horizons, consistent with their high ECS and low realised warming fraction that sustains a large thermal pipeline. These groupings are broadly consistent with the Level 2 formula  
575 (eq. (19)): models with lower  $r$  and higher ECS develop a larger thermal surplus, while those with higher  $r$  and lower ECS allow carbon drawdown to dominate.

#### 5.4 Parameter drivers of ZEC

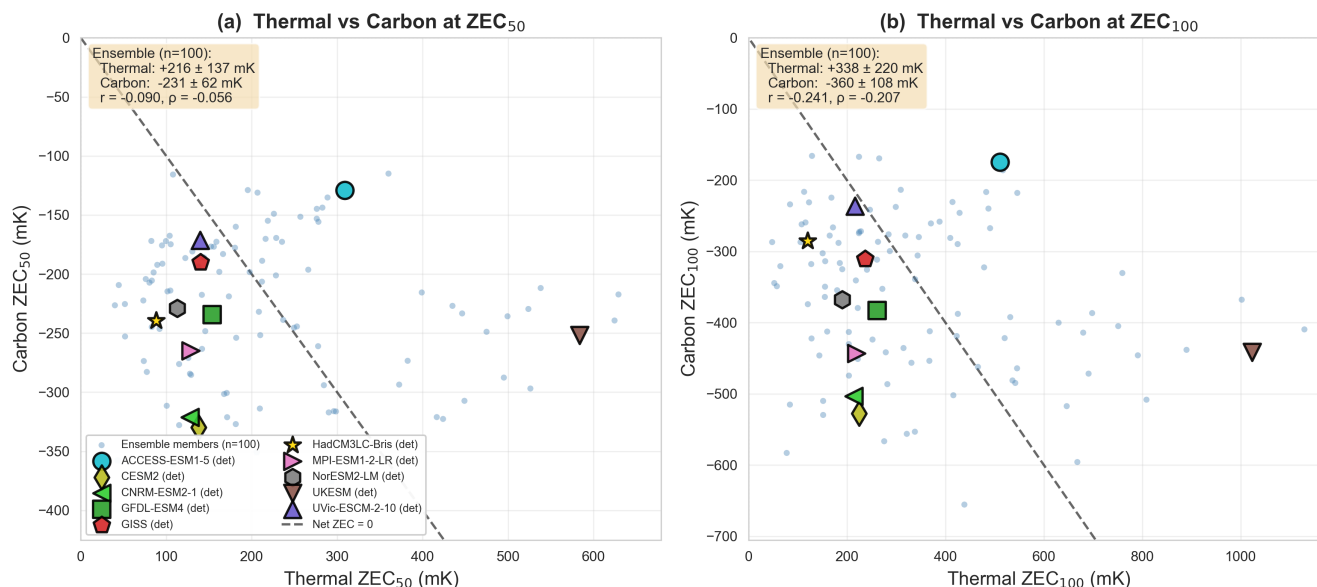
Figure 6 shows the correlation between ensemble parameters and ZEC at three horizons (25, 50, and 100 years), computed from the pooled multi-model ensemble. The dominant positive correlations are with ECS (higher sensitivity  $\rightarrow$  more thermal



**Figure 4.** Ensemble distributions of ZEC at 50 years (top row) and 100 years (bottom row) post-cessation, decomposed into total, thermal, and carbon components. Violins show individual ESM ensemble distributions; the black curve shows the pooled kernel density estimate. Horizontal dashed lines mark the pooled mean. The opposite-sign thermal and carbon components produce a total ZEC distribution whose mean lies between them, close to zero.




### Thermal vs Carbon Components of ZEC (Multi-model probabilistic ensemble + deterministic fits)



**Figure 5.** Joint distribution of thermal and carbon ZEC components at (a) 50 and (b) 100 years post-cessation. Blue dots: individual ensemble members pooled across all ESMs. Coloured markers: deterministic best-fit values for each ESM. The dashed line marks net ZEC = 0 (perfect thermal–carbon compensation). Despite the strong linear relationship between cumulative ocean carbon and heat uptakes during the warming phase (Brovkin et al., 2025), the thermal and carbon components of ZEC show only a very weak correlation after emission cessation.

580 pipeline → larger ZEC), the ocean heat exchange coupling  $\gamma$  (stronger coupling → faster heat delivery), and the carbon–climate feedback  $\zeta$  (weaker sinks → less carbon cooling). The mixed-layer heat capacity  $C_f$  is negatively correlated with ZEC: a larger  $C_f$  lengthens  $\tau_1^{\text{clim}}$  and slows the climate response.

The correlation structure shifts with horizon. ECS is the strongest positive correlate at every horizon ( $r \approx 0.55$ ), but the composition of secondary drivers changes: at  $ZEC_{25}$ ,  $C_f$  is the second-strongest positive driver ( $r = 0.46$ ), reflecting the  
585 dominance of the fast climate mode on short timescales. By  $ZEC_{100}$ ,  $\gamma$  and  $\varepsilon$  overtake  $C_f$  as the slow mode delivers an increasing share of the thermal pipeline warming. The heat-uptake efficacy  $\varepsilon$  enters the eigenstructure only through the product  $\varepsilon\gamma$ , so it is partially degenerate with  $\gamma$  in the flat10 calibration. Its normalized sensitivity to ZEC is nonetheless nontrivial, comparable to that of  $\gamma$  itself, suggesting that independent constraints on  $\varepsilon$  (for example through analysis of observed global radiative imbalance, ocean heat content, and surface temperature patterns) could materially reduce ZEC uncertainty. The full  
590 scatter plots underlying this heatmap are shown in the Supplementary Material (Fig. S1).



**Figure 6.** Pearson correlation coefficients between ensemble parameters and ZEC at three post-cessation horizons (25, 50, 100 years), pooled across nine ESMs. Warm colours: positive correlation (parameter increases ZEC); cool colours: negative. Stars denote statistical significance (\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ ).

## 5.5 Implications for the RAZE

The RAZE decomposition in section 4.4 reveals why inter-model spread in ZEC is dominated by the thermal term. The thermal delivery rate  $\rho = +0.20 \pm 0.16\% \text{ yr}^{-1}$  has a larger absolute standard deviation than the effective carbon uptake rate  $1/\tau_c^* = +0.13 \pm 0.13\% \text{ yr}^{-1}$  (both expressed as mean  $\pm$  one standard deviation across the pooled ensemble). Since ZEC is proportional to the difference  $\rho - 1/\tau_c^*$ , the larger absolute spread of the thermal rate propagates more variance into ZEC than the carbon rate does, consistent with the factor-of-two asymmetry in component standard deviations noted above. Although the thermal term dominates the *spread*, the carbon term dominates the *mean*: eight of the ten ESMs produce negative ZEC (carbon-dominated), so the multi-model mean is near zero or weakly negative, consistent with both ZECMIP (MacDougall et al., 2020) and the RAZE analysis of Jenkins et al. (2022). The thermal spread is governed by uncertainty in ECS,  $\tau_2^{\text{clim}}$ , and the ocean heat exchange coupling  $\gamma$  (quantities that the 200-year flat10 experiment constrains only weakly), whereas the carbon uptake rate is anchored by the directly observable  $\text{CO}_2$  drawdown trajectory. The observational implications of this asymmetry are discussed in section 6.4.

## 5.6 Out-of-sample validation: flat10-cdr

The flat10-cdr experiment provides an out-of-sample test for climate reversibility: in flat10-cdr, after 100 years of constant emissions,  $\text{CO}_2$  emissions become negative at the same rate ( $-10 \text{ PgC yr}^{-1}$ ). The model, fitted only to the flat10 and flat10-zec experiments, is integrated under this novel scenario.

Figure 7 compares predicted and observed  $\text{CO}_2$  and temperature trajectories for all ten ESMs. For most models, the deterministic prediction (red dashed) tracks the ESM (black) well during the constant-emission phase and reproduces the initial response to negative emissions, while the 5th–95th percentile ensemble band (blue shading) generally encloses the ESM tra-

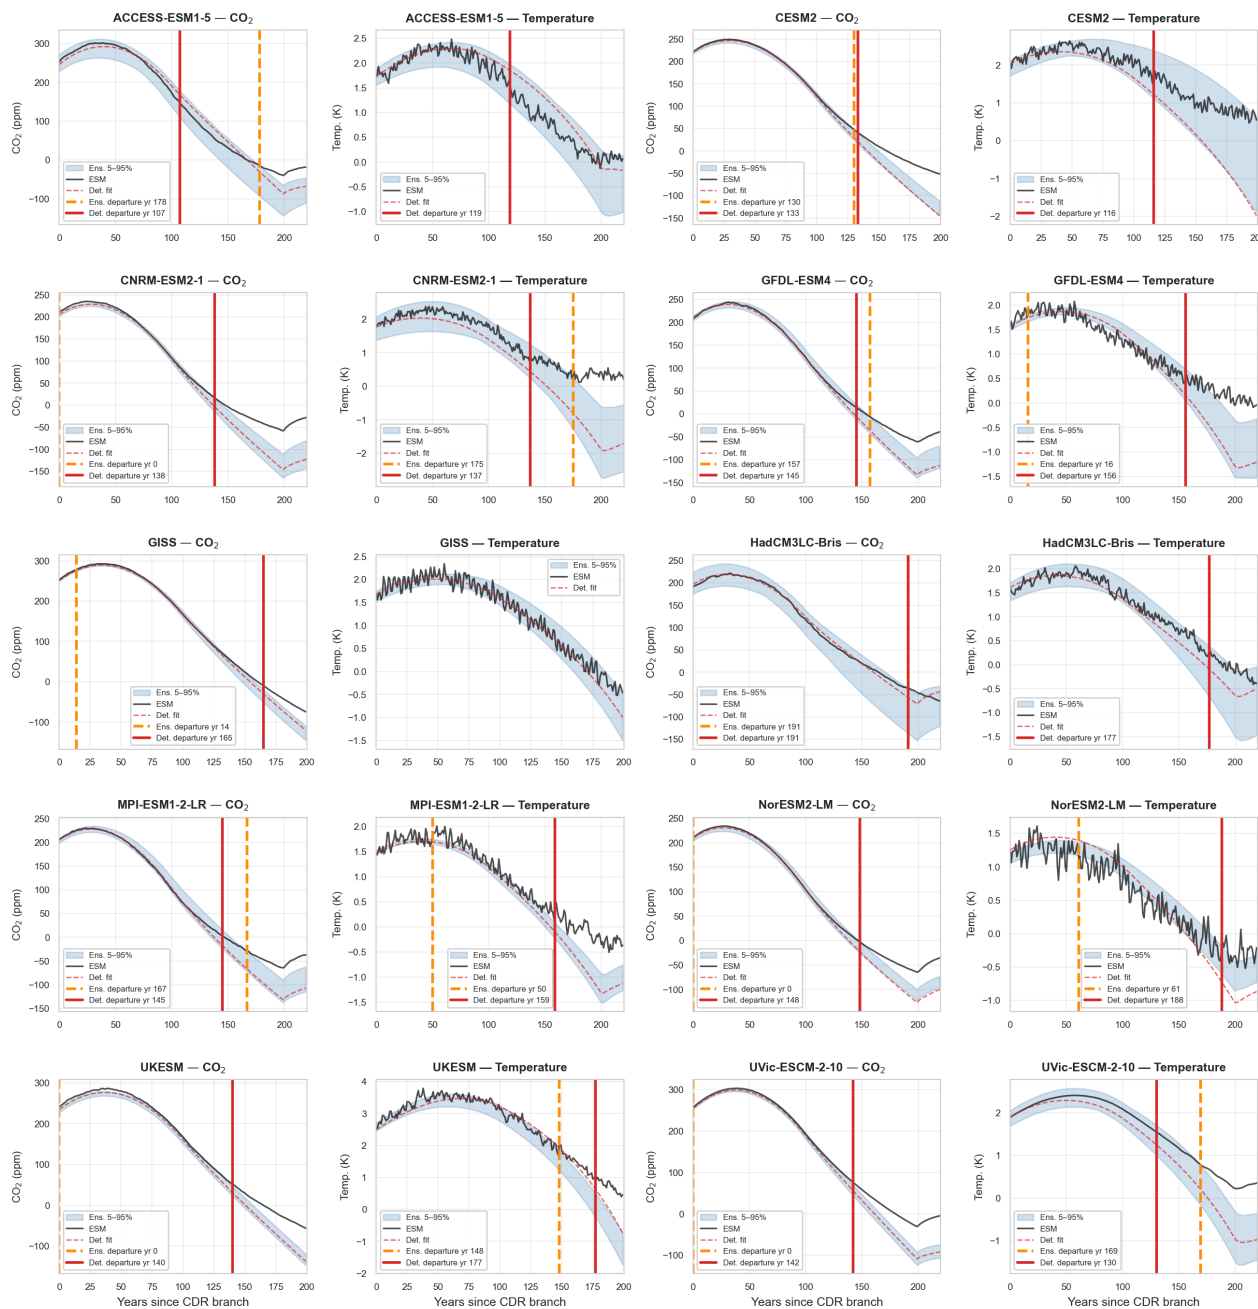


610 jectory. Performance degrades at longer times as cumulative negative emissions grow large, with CO<sub>2</sub> predictions diverging before temperature in most models, visually apparent as earlier departure of the black curve from the blue band in the left (CO<sub>2</sub>) panels than the right (temperature) panels.

To quantify this, we compute the first year at which each ESM trajectory persistently (for  $\geq 10$  consecutive years) leaves the 5th–95th percentile ensemble envelope, separately for CO<sub>2</sub> and temperature (orange dashed lines in fig. 7). If the carbon cycle parameterisation is the primary source of out-of-sample degradation, CO<sub>2</sub> should depart the envelope before temperature. This is indeed the case: in 6 of 10 models the CO<sub>2</sub> trajectory leaves the envelope first, with lead times of 32–189 years over the temperature departure. In a another model (ACCESS-ESM1-5), CO<sub>2</sub> departs but temperature never persistently leaves the envelope. Of the three remaining models where temperature departs first, GISS (year 4), NorESM2-LM (year 21), and HadCM3LC-Bris (year 88) all show temperature departing the ensemble envelope before CO<sub>2</sub>, with narrowly calibrated ensemble spreads contributing to the early departure in GISS and NorESM2-LM.

We complement the envelope departure metric with a second measure of *structural divergence*: the first year at which the absolute error of the deterministic prediction exceeds a fixed threshold (20 ppm for CO<sub>2</sub>, 0.3 K for temperature) for  $\geq 10$  consecutive years (red solid lines in fig. 7). The gap between the two vertical lines reveals the failure mode for each model. Where the ensemble departure (orange) occurs much earlier than the deterministic departure (red), as in GISS, CESM2, CNRM-ESM2-1, GFDL-ESM4, MPI-ESM1-2-LR, and UKESM, the ensemble spread is narrower than the true uncertainty, but the deterministic prediction remains qualitatively accurate until well past year 130. In contrast, where both lines appear at similar times (as in ACCESS-ESM1-5), the deterministic prediction itself has structurally diverged from the ESM. Even with this relatively strict temperature threshold, for 2 of 10 models (GISS and HadCM3LC-Bris) the temperature error never exceeds 0.3 K, confirming that the thermal component of the model remains robust even when the carbon cycle prediction has degraded.

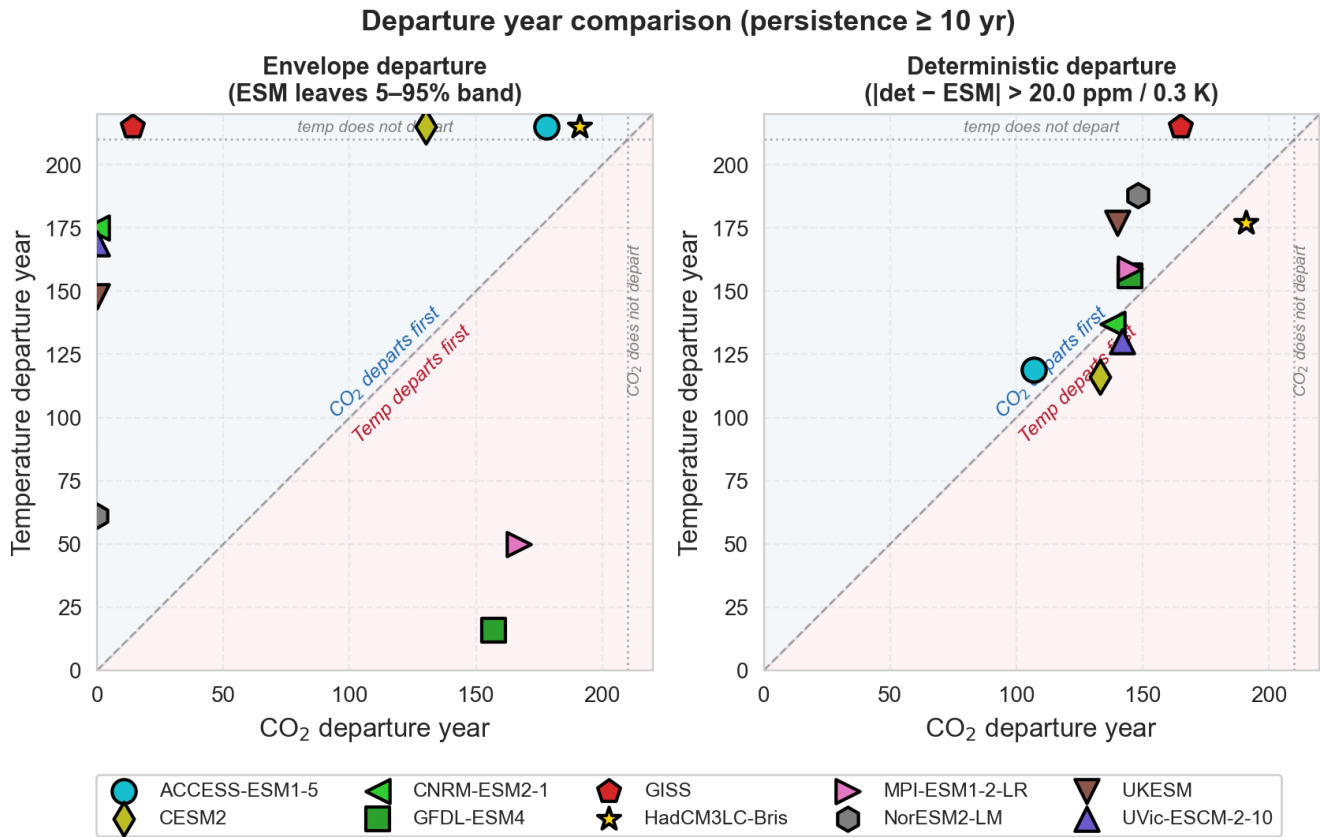
The consistent priority of CO<sub>2</sub> departure supports the interpretation that the carbon cycle feedback parametrisation, not the thermal model, is the primary source of out-of-sample error under large negative emissions. Notably, if the model is trained including the flat10-cdr experiment data, performance for flat10 and flat10-zec degrades, suggesting that the model structure cannot simultaneously fit all three experiments. This suggests the linear feedback parametrisation is insufficient to capture the dynamical response to large negative emissions (Schwinger and Tjiputra, 2018), a finding with broader implications. The deterministic CO<sub>2</sub> prediction diverges from the ESM within 33–89 years of the emission trajectory departing from the constant training rate (median 46 years), with most models departing while emissions are still positive during the ramp-down phase. This connects to the finding of Koven et al. (2023) that ZEC emerges during the emission reduction phase: while our decomposition successfully captures that early emergence within the flat10 framework, the out-of-sample test reveals that the linear carbon–climate feedback cannot generalise to emission pathways that differ substantially from those used for training. The thermal–carbon decomposition presented here thus provides useful physical insight into the drivers of ZEC, but is not a substitute for ESM simulations when evaluating the climate response under novel emission pathways.



**Figure 7.** Out-of-sample validation against the flat10-cdr experiment for ten ESMs. Each row shows one model, with CO<sub>2</sub> anomaly (ppm, left) and temperature anomaly (K, right). Black lines: ESM output; red dashed lines: deterministic prediction (trained on flat10 and flat10-zec only); blue shading: 5th–95th percentile ensemble band. Orange dashed vertical lines mark the first year of persistent ( $\geq 10$  yr) departure from the ensemble envelope; red solid vertical lines mark the first year when the deterministic prediction error exceeds 20 ppm (CO<sub>2</sub>) or 0.3 K (temperature). The gap between the two lines indicates whether out-of-sample failure reflects an underestimated ensemble spread (large gap) or genuine structural divergence of the deterministic fit (small gap or overlapping lines).



645 Figure 8 summarises the departure asymmetry: for both the envelope and deterministic metrics, CO<sub>2</sub> departure precedes temperature departure in the majority of models, indicating that the carbon cycle parameterisation is the leading source of out-of-sample error rather than the thermal model.



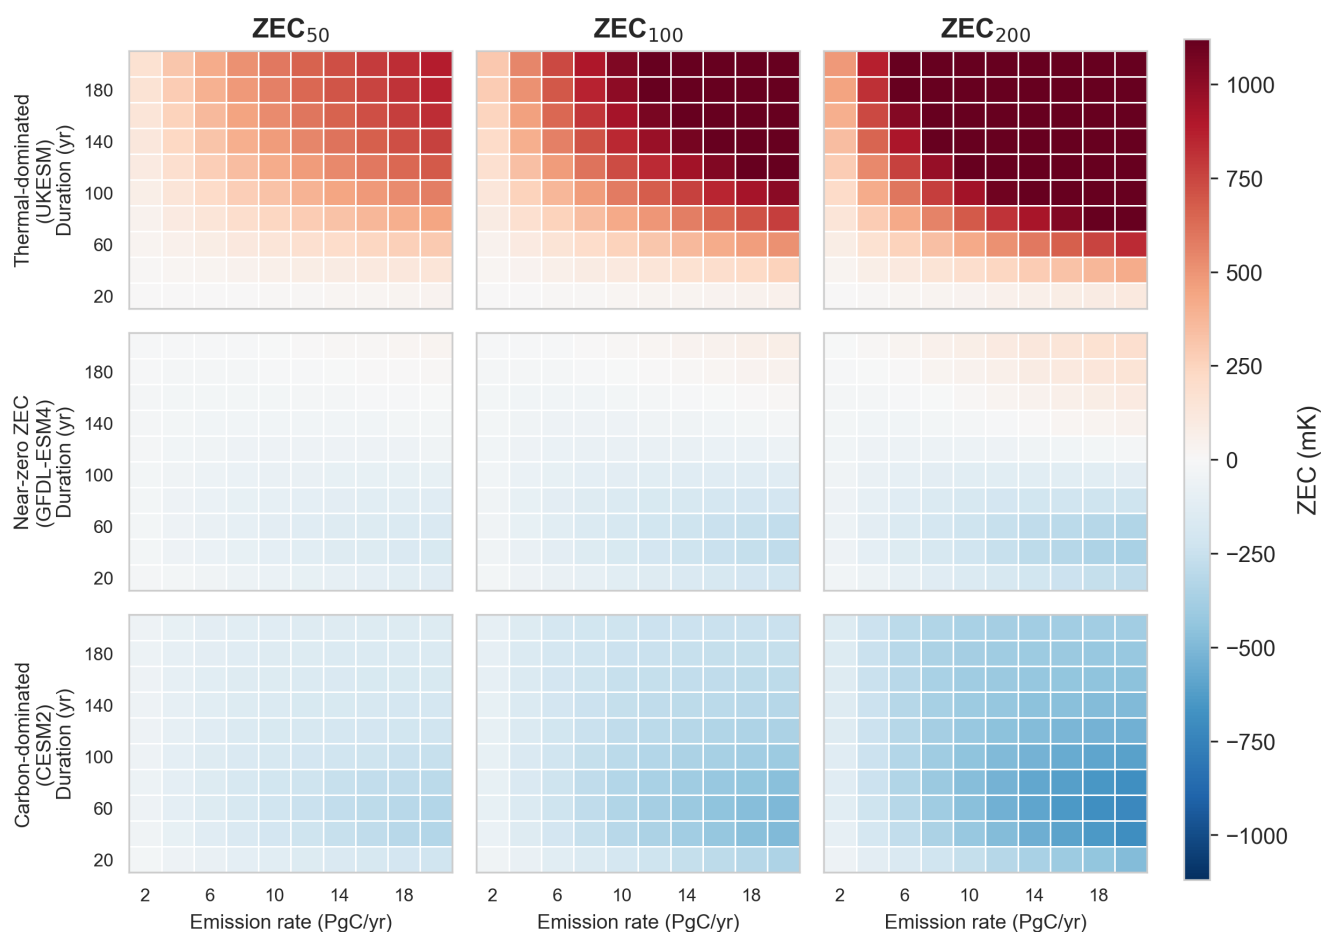
**Figure 8.** Departure year comparison for CO<sub>2</sub> versus temperature in the flat10-cdr out-of-sample test. Left: year at which the ESM first persistently ( $\geq 10$  yr) leaves the 5th–95th percentile ensemble envelope. Right: year at which the deterministic prediction error first persistently exceeds 20 ppm (CO<sub>2</sub>) or 0.3 K (temperature). Points below the 1:1 line indicate that CO<sub>2</sub> departs before temperature; points in the grey strips along the edges indicate that one variable never departs. Marker colours and shapes follow fig. 5.

### 5.7 ZEC sensitivity to emission scenarios

650 The Level 2 formula (eq. (19)) predicts that ZEC over multi-decadal timescales following net zero depends on the warming at cessation  $T_0$  (which is scenario-dependent) and on intrinsic system properties ( $r$ ,  $\tau_2^{\text{clim}}$ ,  $\tau_c$ , etc.). Figure 9 shows ZEC<sub>50</sub> as a function of cumulative emissions and emission rate for the pooled ensemble, revealing the sensitivity of the thermal–carbon compensation to the scenario.



At low cumulative emissions, both thermal and carbon terms are small and the overall effect is small. At higher cumulative emissions, the thermal term grows faster than the carbon term (because the thermal warming pipeline scales with  $T_{eq}$  while the drawdown fraction saturates), producing a net positive ZEC. The emission rate matters because faster emissions shift the relative timing of thermal equilibration and carbon drawdown.



**Figure 9.** Sensitivity of ZEC to emission duration and rate for three representative models: UKESM (thermal-dominated, positive  $ZEC_{50}$ ), GFDL-ESM4 (near-zero  $ZEC_{50}$ ), and CESM2 (carbon-dominated, negative  $ZEC_{50}$ ). Columns show ZEC at 50, 100, and 200 years after emission cessation, in millikelvin. Longer emission durations and faster rates shift ZEC positive because the pipeline warming term (which scales with  $T_{eq}$ ) grows faster than the carbon drawdown term. The thermal-dominated group also includes ACCESS-ESM1-5; the carbon-dominated group also includes CNRM-ESM2-1, GISS, HadCM3LC-Bris, MPI-ESM1-2-LR, NorESM2-LM, and UVic-ESCM-2.10. The full ten-model version is shown in the Supplementary Material (Fig. S2).



## 655 6 Discussion

### 6.1 Near-compensation of thermal and carbon terms

The central finding is that  $ZEC_{50}$  is small not because both constituent processes are small, but because they are nearly equal in magnitude and evolve on comparable timescales. The thermal component (+0.216 K) and carbon component (−0.231 K) each exceed the total ZEC by more than an order of magnitude. This compensation is not structurally required by the model equations: it reflects a quantitative coincidence between the rate of deep-ocean heat delivery ( $\sim 50/\tau_2^{\text{clim}}$ ) and the rate of CO<sub>2</sub> drawdown cooling ( $\sim q_{\text{eff}}(1 - f_{\infty})/[(1 + C_0/C(t_0))\tau_c]$ , where  $\tau_c$  is an effective timescale combining land and ocean uptake) in the current climate system.

The asymmetry in uncertainty between the two components has direct implications. As such, since the carbon component shows less inter-model spread (std 0.062 K versus 0.137 K), and the most positive ZEC values are driven almost entirely by large unrealised warming, improving constraints on the thermal component (particularly ECS,  $\tau_2^{\text{clim}}$ , and the ocean heat exchange coupling) is the most effective target for reducing ZEC uncertainty. This parallels the finding of Jones and Friedlingstein (2020) that climate response uncertainty dominates over carbon cycle uncertainty in the TCRE budget. The carbon cycle parameters show less inter-model spread partly because atmospheric CO<sub>2</sub> directly records the net sink strength, and partly because the IRF pools aggregate land and ocean uptake into effective timescales that are constrained jointly by the atmospheric CO<sub>2</sub> trajectory.

The carbon–climate feedback ( $\zeta$ ) reinforces this asymmetry by pushing ZEC positive in two ways: (i) higher CO<sub>2</sub> at cessation increases  $F_0$  and hence the thermal component; (ii) the factor  $(1 - \zeta\bar{T})$  directly weakens post-cessation drawdown. For the flat10 experiment ( $T_0 \sim 1$  K), the correction is small ( $\mathcal{O}(0.001\text{--}0.003$  K)), but it scales as  $T_0^2$  and could reach  $\mathcal{O}(0.05$  K) for emission pathways that produce 2–3 K of warming before cessation.

### 675 6.2 Normalised ZEC frameworks

Williams et al. (2025) independently decompose ZEC into normalised thermal, radiative, and carbon contributions, finding that the thermal component drives 58% of inter-model variance and the carbon component 40%. Our results are quantitatively consistent: the thermal standard deviation (0.137 K) exceeds the carbon standard deviation (0.062 K) by a factor of 2.2, comparable to their variance ratio of 1.45. The qualitative agreement across protocols (Williams et al. use ZECMIP (1000 PgC from 1% yr<sup>−1</sup> CO<sub>2</sub>), while we use flat10MIP) supports the conclusion that partial compensation is a common feature of the current generation of ESMs, although its degree varies across models.

Both Williams et al.’s normalised framework and our Level 2 formula (eq. (19)) predict that ZEC should scale linearly with  $T_0$ , making a normalised (per-kelvin) decomposition natural. Our results partially support this: models with positive ZEC (where the thermal component dominates) cluster tightly along a line  $ZEC_{50} \propto T_0$  (fig. 5), consistent with the picture that unrealised warming is a fixed fraction of the warming at cessation. Yet models with negative ZEC (where C-cycle drawdown wins) deviate from strict proportionality. The reason is that the carbon cooling term in eq. (19) contains the factor  $(1 + C_0/C(t_0))^{-1}$ , where  $C(t_0)$  is the CO<sub>2</sub> anomaly at cessation. Because  $C(t_0)$  itself grows sub-linearly with cumulative emissions (due to the



logarithmic relationship between  $\text{CO}_2$  and forcing), the carbon term does not scale purely as  $T_0$ ; it retains a residual dependence on the emission pathway that breaks the normalisation. In practice, this means that a normalised decomposition works well when ZEC is positive, and thus when the thermal warming pipeline is the dominant signal, but becomes less accurate for negative-ZEC models where the carbon term introduces additional pathway dependence.

Brovkin et al. (2025) take a complementary approach to the same flat10MIP data, deriving a simplified ordinary differential equation for the airborne fraction by exploiting quasi-linear relationships between cumulative land and ocean carbon uptakes ( $C_l \approx k C_o$ ) and between ocean carbon and heat uptakes ( $C_o \approx \eta H$ ). Their reduction to a single ODE is analogous to our own use of impulse-response functions to decouple the thermal and carbon sub-systems, though the two approaches proceed in opposite directions: Brovkin et al. start from the global energy balance and derive the airborne fraction as a prognostic variable, while we start from calibrated impulse-response representations and derive ZEC as a sum of independently interpretable terms. Both approaches provide independent confirmation that the coupled climate–carbon dynamics in flat10MIP are well captured by low-dimensional analytical models.

The most direct point of contact concerns the TCRE constancy. Brovkin et al. show that under logarithmic forcing the airborne fraction increases after 30–40 years, compensating for the weakening radiative forcing per unit  $\text{CO}_2$  and thereby stabilising TCRE. In our framework, the same compensation appears as a near compensation between the thermal (+0.216 K) and carbon (−0.231 K) terms of ZEC: the residual warming after cessation is small precisely because the rate at which the ocean delivers stored heat to the surface is matched by the rate at which  $\text{CO}_2$  drawdown reduces radiative forcing. The two perspectives are thus independent views of the same underlying physical balance.

### 6.3 Pattern effect, efficacy, and the thermal warming pipeline

The efficacy parameter  $\varepsilon$  in our two-box model (eqs. (5) and (6)) plays the same role as the efficacy of ocean heat uptake in the thermal pipeline analyses of Hansen et al. (2023) and Armour et al. (2024): it captures the pattern-dependent relationship between ocean heat uptake and surface radiative feedbacks. Armour et al. show that spatially non-uniform warming patterns have depressed the effective climate sensitivity governing recent decades to about 2 K, well below equilibrium, implying a substantial reservoir of unrealised warming that will emerge as warming patterns evolve. Myhre et al. (2025) corroborate this from the energy-budget side, demonstrating that only models with climate sensitivity above 2.5 K reproduce the observed 2001–2023 trend in Earth’s energy imbalance.

#### 6.3.1 The carbon buffer

Our analytical framework makes explicit why the large thermal pipeline identified by these studies has a muted effect on ZEC. The committed warming at fixed atmospheric composition is  $T_0(1/r - 1)$ , where  $r$  is the realised warming fraction; for our pooled ensemble this quantity has a median of  $\sim 0.7$  K (fig. 10a, grey dashed curve). Two successive buffers reduce this thermal pipeline before it is expressed in ZEC:



720 1. **Slow delivery.** Only a fraction  $\Delta t/\tau_2^{\text{clim}}$  of the stored heat reaches the surface within  $\Delta t$  years. For the ensemble-median  $\tau_2^{\text{clim}} \approx 440$  yr and  $\Delta t = 50$  yr, approximately 11% of the thermal pipeline is delivered, yielding a thermal  $\text{ZEC}_{50} \approx +0.13$  K.

2. **Carbon drawdown.** Post-cessation  $\text{CO}_2$  uptake by the ocean and land sinks generates a roughly equal and opposite carbon  $\text{ZEC}_{50} \approx -0.06$  K (median), reducing the net  $\text{ZEC}_{50}$  to  $\sim 0.05$  K. This carbon offset depends on carbon-cycle parameters ( $\tau_c, f_\infty, \zeta$ ) and is essentially independent of  $r$ .

725 These two buffers are visible in fig. 10a: the grey dashed curve (committed warming) is steep in  $r$ , but the net ZEC (black solid curve) is nearly flat because the carbon offset (red) does not move with  $r$ . As the horizon lengthens to 100 years (dashed coloured curves), both thermal and carbon terms roughly double but remain closely matched, preserving the near-compensation. fig. 10b confirms this across the full ensemble: the running median of ZEC remains weakly positive ( $\lesssim 0.1$  K) even as thermal pipeline warming ranges from 0 to 3 K, far below the 1:1 line that would apply without carbon buffering.

### 730 6.3.2 Sensitivity to the realised warming fraction $r$

The marginal sensitivity of net ZEC to the realised warming fraction is


$$\frac{\partial \text{ZEC}_{\Delta t}}{\partial r} = -\frac{T_0 \Delta t}{r^2 \tau_2^{\text{clim}}}, \quad (21)$$

735 which for  $T_0 \approx 1.7$  K,  $r = 0.6$ , and  $\tau_2^{\text{clim}} = 440$  yr evaluates to approximately  $-0.54$  K per unit  $r$  at  $\Delta t = 50$  yr and  $-1.07$  K per unit  $r$  at  $\Delta t = 100$  yr. The corresponding thermal pipeline sensitivity is  $-T_0/r^2 \approx -4.7$  K per unit  $r$ , nearly an order of magnitude larger. In practical terms, if  $r$  were to decrease by 0.1 (for example, because ECS is revised upward), committed warming at fixed composition would increase by  $\sim 0.5$  K, but net ZEC would shift by only  $\sim +0.05$  K at 50 years and  $\sim +0.11$  K at 100 years (fig. 10c).

### 6.3.3 Implications

740 These results have two main consequences. First, the Earth energetic imbalance thermal warming pipeline debate primarily affects ZEC through the thermal channel, and even there the signal is attenuated by a factor  $\Delta t/\tau_2^{\text{clim}} \sim 0.1$  relative to committed warming at fixed composition. The carbon cycle acts as a buffer that strongly decouples ZEC uncertainty from uncertainty in  $r$  (and hence in ECS and the pattern effect). This buffering role is not unique to post-cessation dynamics: during ongoing emissions, the same sink processes keep the airborne fraction below unity and contribute to the near-constancy of TCRE. Second, the remaining uncertainty in ZEC is dominated by the thermal component, not because the carbon offset is uncertain, 745 but because its inter-model spread ( $\sigma = 0.06$  K for  $\text{ZEC}_{50}$ ) is smaller than that of the thermal term ( $\sigma = 0.13$  K). Tightening constraints on the slow ocean timescale  $\tau_2^{\text{clim}}$ , ECS, and ocean heat exchange coupling  $\gamma$  therefore offers the most leverage for narrowing ZEC bounds, consistent with the thrust of current observational work on the thermal pipeline.

The efficacy analysis (Supplementary Material, Section S4), performed for eight of the ten ESMs (MPI-ESM1-2-LR and UVic-ESCM-2-10 lack the required top-of-atmosphere radiation variables), reveals a structural asymmetry: during forcing



**Figure 10.** The carbon buffer between the thermal pipeline and ZEC. **(a)** Temperature components as a function of the realised warming fraction  $r$ . Grey dashed: committed warming at fixed composition,  $T_0(1/r - 1)$ . Coloured solid (dashed) curves: analytical thermal, carbon, and net ZEC evaluated at median parameters for  $\Delta t = 50$  yr ( $\Delta t = 100$  yr). Scatter points show individual ensemble members (148 across nine ESMs) for  $\Delta t = 50$  yr. Arrows highlight the two successive buffers (slow delivery and carbon drawdown) that reduce the thermal pipeline to the net ZEC. **(b)** Running median of ZEC<sub>50</sub> (solid) and ZEC<sub>100</sub> (dashed) as a function of thermal pipeline warming  $T_0(1/r - 1)$ , computed in a sliding  $\pm 0.5$  K window. The medians are weakly positive throughout, confirming that the thermal pipeline is almost entirely offset by carbon drawdown; the dotted 1:1 line shows the response that would obtain without any carbon buffering. **(c)** Probability distributions of the four quantities, illustrating how the broad thermal pipeline distribution (grey) is compressed to a narrow net ZEC (black). Solid (dashed) lines denote  $\Delta t = 50$  yr ( $\Delta t = 100$  yr).

750 buildup, the ocean heat uptake pattern effect suppresses transient warming ( $\varepsilon > 1$ ); after cessation, the warming pattern relaxes toward equilibrium and the effective feedback weakens. A model fitted with  $\varepsilon = 1$  absorbs this into a larger effective  $\lambda_f$ , then applies this (too-strong) feedback to the relaxation phase, potentially underestimating the thermal ZEC by  $\mathcal{O}(0.010\text{--}0.030)$  K.

This effect is small for the flat10 experiment but would grow for scenarios with larger cumulative emissions and longer post-cessation horizons. The sensitivity analysis over  $\varepsilon = 1.0\text{--}1.8$  (fig. 9) provides bounds on this structural uncertainty.

#### 755 6.4 Observational leverage on ZEC uncertainty

The dominance of the thermal component in ZEC uncertainty (section 5.5) points to specific observational targets. Of the three thermal parameters that most influence ZEC (ECS, the ocean heat exchange coupling  $\gamma$ , and the deep-ocean timescale  $\tau_2^{\text{clim}}$ ), ECS is increasingly well constrained by the combination of historical warming, radiative forcing, and paleoclimate evidence (Sherwood et al., 2020). The remaining uncertainty in ZEC therefore hinges largely on the ocean parameters  $\gamma$  and  $\tau_2^{\text{clim}}$ , which  
760 govern how quickly stored heat is delivered to the surface after emissions cease.

These parameters can be constrained by observations that resolve the vertical structure of ocean heat content. The Argo float network already provides global coverage of the upper 2000 m (Cheng et al., 2022; von Schuckmann et al., 2023), sufficient to estimate the partitioning of heat uptake between surface and intermediate layers and thereby inform the two-box model's



$C_f$  and  $\gamma$  (Geoffroy et al., 2013). Constraining  $\tau_2^{\text{clim}}$ , however, requires observations of deep-ocean heat content changes  
765 below 2000 m, a regime where data remain sparse. The extension of Argo to full-depth profiling (Deep Argo) and sustained  
monitoring of deep western boundary currents (e.g. through RAPID-MOCHA and OSNAP) would directly target this gap.  
Additionally, paleo-oceanographic constraints on ocean circulation timescales during past climate transitions could provide  
independent bounds on  $\tau_2^{\text{clim}}$ .

For the carbon component, atmospheric  $\text{CO}_2$  measurements already constrain the net drawdown rate well, but partition-  
770 ing between land and ocean sinks (which determines whether the post-cessation drawdown persists on decadal or centennial  
timescales) benefits from continued investment in ocean carbon observations (e.g. BGC-Argo, repeat hydrography) and terres-  
trial carbon monitoring.

## 6.5 Limitations

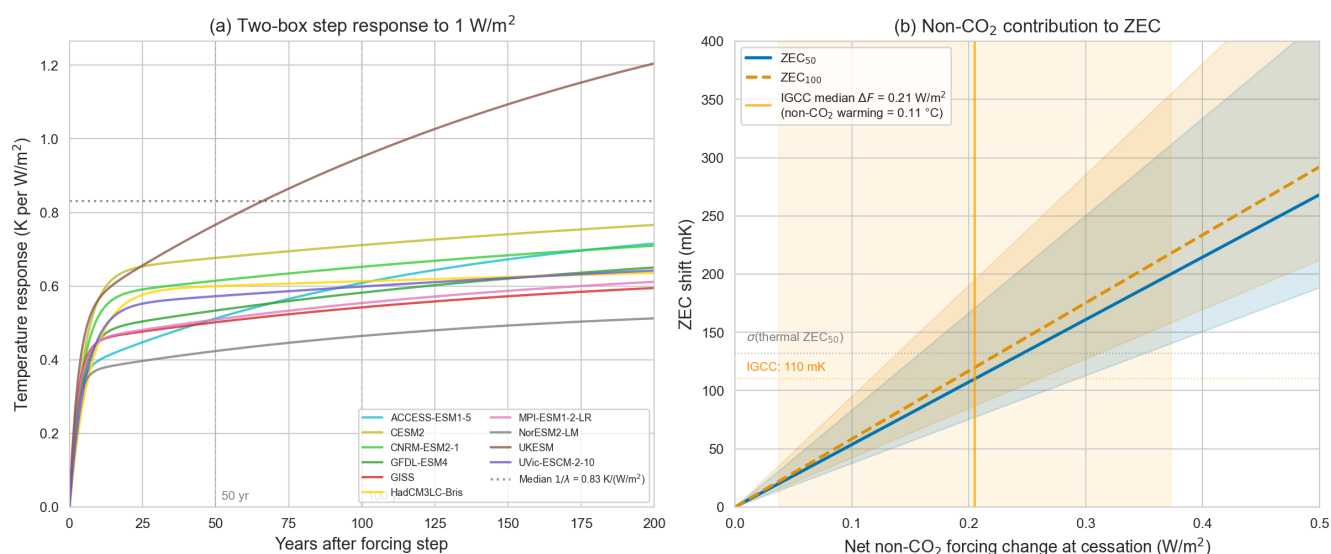
Several limitations qualify our results and merit reflection. The most consequential is the linear parameterisation of the carbon-  
775 climate feedback,  $f_i^*(t') = f_i(1 - \zeta T(t'))$ , which assumes a proportional relationship between warming and carbon uptake  
reduction. However, non-linear carbon-cycle feedbacks that could emerge at higher warming levels or on longer timescales  
(for example, permafrost carbon release, decreased ocean  $\text{CO}_2$  solubility, or unrepresented ecosystem responses in CMIP6-  
class models) could weaken the carbon buffer and increase ZEC. In addition, our own results suggest that strong negative  
emissions could break the underlying assumptions of the model structure, which could have implications for ZEC under large-  
780 scale carbon dioxide removal scenarios.

Our approximation performs well in the flat10 regime ( $\sim 1$  K of warming) but degrades for large temperature excursions or  
under sustained negative emissions, as the out-of-sample validation demonstrates (section 5.6). The flat10-cdr experiment is  
particularly instructive: although the model reproduces ZEC behaviour well in the first  $\sim 100$  years after emission cessation,  
it progressively diverges from ESM trajectories as cumulative negative emissions grow and the system moves into a regime  
785 of strong climate reversal. This demonstrates that a framework calibrated to explain ZEC (the response to zero emissions) is  
not sufficient to predict the response to large-scale carbon dioxide removal, where nonlinearities in ocean carbon chemistry,  
terrestrial carbon feedbacks, and the pattern effect under cooling may all behave asymmetrically relative to the warming  
trajectory. Training the model on the flat10-cdr data degrades performance for the flat10 and flat10-zec experiments, suggesting  
that the linear feedback parameterisation cannot simultaneously capture the dynamics of both warming and cooling trajectories  
790 (see also S  f  rian et al., 2024, who demonstrate multi-decadal lags between carbon cycle and climate responses to  $\text{CO}_2$  removal  
in CNRM-ESM2-1). Climate reversibility cannot therefore be assumed from ZEC constraints alone, and dedicated CDR-  
specific model evaluation is needed. Furthermore, Koven et al. (2022) showed that overshoot asymmetry in ESMs is closely  
linked to the ZEC at longer horizons ( $\text{ZEC}_{90}$ ), reinforcing the connection between post-cessation temperature evolution and  
the feasibility of temperature reversal through negative emissions. Existing simple climate model formulations used for large  
795 negative emission scenarios may need to be reassessed in light of these findings, as they often rely on similar linear feedback  
assumptions that may not hold under strong cooling conditions. Empirical evidence of this was presented in Sanderson et al.  
(2025a), where some SCM ensembles overestimated cooling in the flat10-cdr experiment, and this is further corroborated by the



800 out-of-sample validation results presented here. Experiments like flat10-cdr could be a valuable addition to reduced-complexity model intercomparisons such as RCMIP (Romero-Prieto et al., 2025) for systematically evaluating how SCMs perform under strong negative emissions.

805 The CO<sub>2</sub>-only scope of the model introduces a further limitation: concurrent changes in non-CO<sub>2</sub> forcing agents (aerosols, methane, nitrous oxide) at net-zero CO<sub>2</sub> would shift the baseline from which ZEC is measured (see also the non-CO<sub>2</sub> analysis in Jenkins et al., 2022), and permafrost carbon feedbacks, which Steinert and Sanderson (2025) estimate could increase ZEC<sub>100</sub> by ~0.08 K using FaIR with the flat10 protocol (consistent with an earlier estimate of ~0.06 K by MacDougall 2021), are not represented. To estimate the magnitude of the non-CO<sub>2</sub> forcing effect, we applied the two-box step response from the pooled ensemble to a sustained forcing perturbation. Remaining carbon budget frameworks assume a net non-CO<sub>2</sub> warming contribution of approximately 0.1°C at net-zero CO<sub>2</sub> (Forster et al., 2025), implying a forcing change of order 0.1–0.2 W m<sup>-2</sup>. A sustained perturbation of that magnitude would shift ZEC<sub>50</sub> by roughly 50–110 mK in our framework, comparable in size to the CO<sub>2</sub>-only ZEC values diagnosed here (fig. 11).



**Figure 11.** Estimated contribution of non-CO<sub>2</sub> forcing to ZEC. **(a)** Two-box temperature response to a sustained 1 W m<sup>-2</sup> step forcing, using deterministic parameters for each ESM. Horizontal dashed line marks the median equilibrium sensitivity 1/λ. Vertical grey lines indicate Δt = 50 and 100 yr horizons. **(b)** Implied ZEC shift (mK) as a function of the net non-CO<sub>2</sub> forcing change at emission cessation. Solid (dashed) lines show ZEC<sub>50</sub> (ZEC<sub>100</sub>); shading indicates the 5–95 % range across the pooled ensemble. The orange band marks the non-CO<sub>2</sub> forcing range implied by remaining-carbon-budget frameworks (Forster et al., 2025), with the vertical line at the median estimate.

810 The flat10 experimental design itself introduces scenario dependence. With 10 PgC yr<sup>-1</sup> for 100 years and roughly 1 K of warming, the protocol sits in a regime where carbon-cycle and climate-response nonlinearities remain modest. At higher cumulative emissions the thermal–carbon compensation documented here may evolve differently, and the linear impulse-response structure that underpins our analytical framework would require reassessment.



The two-box energy balance climate model, while a standard tool in reduced-complexity climate modelling (Held et al., 2010; Geoffroy et al., 2013), carries inherent structural limitations. The two-layer representation is known to be overly diffusive compared with the resolved ocean circulation in comprehensive ESMs (Raper et al., 2002; Jeevanjee et al., 2025), potentially smoothing the vertical structure of heat uptake and affecting the timescale separation between fast and slow response modes. It captures the dominant fast and slow response modes that emerge from the eigendecomposition of comprehensive ESMs, and with the efficacy parameter it reproduces the first-order effects of time-varying feedbacks documented in the pattern-effect literature (Armour et al., 2024; Myhre et al., 2025). However, it represents these effects through a single globally-averaged efficacy that is constant in time, whereas in reality the feedback pattern evolves continuously as the warming pattern shifts from transient (Southern Ocean-dominated) to equilibrium (Arctic-amplified) configurations (Andrews et al., 2015; Armour et al., 2013). The two-box structure also cannot represent intermediate-timescale processes such as Atlantic Meridional Overturning Circulation (AMOC) reorganisation, which may introduce non-exponential thermal relaxation behaviour after emissions cease, or state-dependent feedbacks (e.g. cloud and ice-albedo) that could alter the relationship between ocean heat release and surface warming at higher temperatures. Regional-scale responses that may be important for impacts-oriented carbon budget assessments are likewise beyond the scope of a globally-averaged framework. For the flat10 regime (1-2 K warming above preindustrial, 50–100 year horizons), these omissions are unlikely to dominate: the two-box model reproduces ESM temperature trajectories to within 0.05–0.15 K RMSE (section 3.2), and the analytical ZEC formulas closely match numerical integrations. At higher warming levels or on multi-centennial timescales, however, these structural simplifications would require careful reassessment. The current calibration targets atmospheric CO<sub>2</sub> and surface temperature; a natural extension would be to use the top-of-atmosphere energy imbalance  $N(t)$ , which the two-box model also predicts, as an additional validation or training target. Because  $N(t)$  directly constrains the ocean heat uptake coupling  $\gamma$  and the efficacy  $\varepsilon$ , incorporating it could tighten the  $\gamma$ - $\varepsilon$  degeneracy that currently limits the separate identification of these parameters. Additionally, the near-cancellation of thermal and carbon terms documented here is diagnosed from global-mean quantities under a single emission rate (10 PgC yr<sup>-1</sup>); whether it persists at very low emission rates, where the forced signal competes with internal climate variability, or at warming levels where ocean circulation reorganisation (e.g. AMOC weakening) breaks the linear timescale structure, remains an open question. The forthcoming Tipping Points Model Intercomparison Project (TIPMIP; Winkelmann et al., 2025), which explores ESM responses at higher warming levels, will provide a direct test of these predictions. Similarly, the global-mean cancellation may mask divergent regional responses, since models can produce similar global ZEC through very different spatial patterns of warming and carbon uptake.

The carbon component of ZEC in our ensemble may also under-sample the true range of possible responses, because CMIP6-era ESMs share similar representations of terrestrial carbon cycling. The multi-model spread in carbon-concentration and carbon-climate feedback parameters has not changed significantly between CMIP5 and CMIP6 (Arora et al., 2020), and land models continue to rely on similar soil decomposition structures (Varney et al., 2024; Gier et al., 2024), and the terrestrial sink (currently absorbing roughly 30% of anthropogenic emissions; Friedlingstein et al. 2022) remains subject to structural uncertainties that are shared across the model ensemble. Key missing or underrepresented processes include nutrient limitation feedbacks (where phosphorus cycling in particular remains rare across ESMs, and the omission of nutrient constraints has been



shown to bias remaining carbon budget estimates by 19–26%; De Sisto and MacDougall 2024; Davies-Barnard et al. 2020;  
850 Wieder et al. 2015), dynamic vegetation mortality, which requires demographic approaches that most current land surface  
schemes lack (Fisher et al., 2018; Fisher and Koven, 2022) and for which observational evidence indicates that climate sensi-  
tivity of tree mortality is a significant and poorly captured uncertainty in projecting the future forest carbon sink (Pugh et al.,  
2019; Needham et al., 2020), and peatland or wetland carbon dynamics, which remain largely absent from CMIP6 models  
despite peatlands storing roughly 25% of global soil carbon (Loisel et al., 2021; Müller and Joos, 2020). A wider spread in  
855 land carbon uptake would primarily affect the magnitude of the carbon ZEC term (which has a current inter-model standard  
deviation of only 0.062 K), but would be unlikely to change the qualitative finding that the thermal component dominates  
inter-model spread. This is because the post-cessation temperature response is governed primarily by differences in radiative  
response, planetary heat uptake and climate feedbacks (Williams et al., 2025; MacDougall et al., 2020; Palazzo Corner et al.,  
2023), while the carbon drawdown timescales are set largely by the rate of ventilated transfer from the surface ocean to the  
860 thermocline and deep ocean on centennial timescales (Katavouta et al., 2019).

## 7 Conclusions

We have shown that a near-zero multi-model mean estimate for Zero Emissions Commitment results from a near-perfect  
compensation between thermal warming (+0.216 K) and carbon cooling (−0.231 K) at 50 years post-cessation. As such, this  
is a compensation of means rather than a variance-reducing cancellation: the two components are nearly uncorrelated ( $\rho \approx 0$ ),  
865 so their uncertainties combine rather than offset, and the total spread (std 0.145 K) exceeds that of either component alone. This  
decomposition, derived from a 12-parameter coupled carbon–climate model fitted to ten flat10MIP ESMs (Sanderson et al.,  
2025a), explains both the near-zero multi-model mean (−0.014 K) and the large inter-model spread (std 0.145 K) in ZEC<sub>50</sub>.

The compact Level 2 formula (eq. (19)) expresses ZEC as the warming at cessation multiplied by the difference of two rates:  
the pace of deep-ocean heat delivery and the pace of CO<sub>2</sub> drawdown cooling. It reproduces numerical ZEC<sub>50</sub> with an MAE  
870 of ~0.026 K and identifies ECS, the ocean heat exchange coupling, the carbon–climate feedback, and the mixed-layer heat  
capacity as the dominant parameter drivers.

Three features of the decomposition have practical consequences:

1. **The carbon component has narrower inter-model spread than the thermal component** (std 0.062 versus 0.137 K),  
so reducing ZEC uncertainty requires improved constraints on ECS and the slow climate timescale rather than on carbon  
875 cycle parameters.
2. **The compensation persists at ZEC<sub>100</sub>** (thermal +0.338 K, carbon −0.360 K) but with wider total spread (std 0.221 K),  
indicating that the thermal–carbon balance evolves on centennial timescales.
3. **The  $q_{\text{eff}}$  correction** (a corrected fraction of the equilibrium warming associated with the fast climate mode, accounting  
for how much of the CO<sub>2</sub> drawdown cooling is realised at the surface given the slow mode’s finite response time)



880 eliminates  $\sim 0.030$  K of systematic bias in simplified ZEC estimates. A residual positive bias remains (Fig. 3), reflecting the fact that linearised analytical frameworks systematically overestimate ZEC relative to full numerical integration.

The analytical framework bridges ESM complexity and simple-model intuition. By resolving ZEC into physically interpretable components, it clarifies why TCRE is approximately constant near net zero, identifies the conditions under which this constancy breaks down, and offers a computationally inexpensive tool for propagating parametric and structural uncertainty  
885 into carbon budget assessments.

These results have direct implications for remaining carbon budget estimates. The dominance of the thermal component in driving inter-model spread suggests that carbon budget calculations can treat  $ZEC_{50}$  as approximately zero for the current generation of ESMs, consistent with the IPCC AR6 assessment (MacDougall et al., 2020). Still, the  $\pm 0.15$  K standard deviation implies that individual ESMs may exhibit ZEC values large enough to shift the remaining budget by  $\mathcal{O}(50)$  PgC in either  
890 direction. Moreover, because ZEC scales with the warming at cessation  $T_0$  (eq. (19)), higher cumulative emissions affect the remaining carbon budget through two channels: directly through the cumulative  $\text{CO}_2$  budget itself, and indirectly by increasing  $T_0$  and hence the magnitude of ZEC. For thermal-dominated models the indirect effect amplifies the post-cessation warming commitment, while for carbon-dominated models it weakens the carbon drawdown cooling (through the carbon–climate feedback term  $1 - \zeta\bar{T}$ ), in both cases shifting ZEC toward more positive values and reducing the remaining budget beyond  
895 what a fixed-ZEC assumption would imply. Delayed mitigation therefore compounds the challenge of meeting a temperature target: each additional year of high emissions simultaneously depletes the  $\text{CO}_2$  budget and increases the ZEC correction applied to it. Jones and Friedlingstein (2020) showed that TCRE uncertainty can be decomposed into contributions from three process-level quantities: the climate sensitivity ( $\alpha \equiv \Delta T / \Delta F$ , the warming per unit forcing), the carbon-cycle response to  $\text{CO}_2$  ( $\beta \equiv \Delta F / \Delta C_a$ , the forcing per unit change in atmospheric carbon, set by logarithmic  $\text{CO}_2$  forcing and the airborne frac-  
900 tion), and the climate–carbon feedback ( $\gamma$ , the modification of carbon uptake by warming), with  $\alpha$  and  $\beta$  jointly dominating. Our framework extends this logic beyond TCRE to ZEC itself: the analytical expressions identify which process-level parameters control ZEC uncertainty and, critically, reveal that the leverage points differ. Whereas TCRE depends on both climate and carbon-cycle parameters, ZEC uncertainty is governed primarily by ECS,  $\tau_2^{\text{clim}}$ , and the ocean heat exchange coupling, the same parameters that control the thermal pipeline. Combining constraints on these quantities with the component-level carbon-  
905 cycle constraints advocated by Jones and Friedlingstein (2020) would therefore tighten confidence bounds on both TCRE and ZEC simultaneously, offering a more complete pathway to narrowing remaining carbon budget estimates.

Finally, the decomposition reveals that the carbon cycle acts as a powerful buffer between the “warming in the pipeline” and ZEC. Recent work on the pattern effect (Hansen et al., 2023; Armour et al., 2024; Myhre et al., 2025) suggests that the realised warming fraction may be lower than previously thought, implying a large reservoir of committed warming at fixed  
910 atmospheric composition. But this thermal pipeline is strongly attenuated before it is expressed in ZEC: slow ocean delivery reduces it by a factor  $\Delta t / \tau_2^{\text{clim}} \sim 0.1$ , and the carbon drawdown offset, which depends on carbon-cycle parameters rather than on climate sensitivity, offsets most of what remains. A 0.1 decrease in the realised warming fraction increases committed warming by  $\sim 0.5$  K but net  $ZEC_{50}$  by only  $\sim 0.05$  K. Uncertainty in CMIP6-generation models considered here lies with the thermal component, and tightening constraints on ECS,  $\tau_2^{\text{clim}}$ , and the ocean heat exchange coupling is where the largest gains



915 in narrowing remaining carbon budget estimates lie. However, common structural assumptions across ESMs, particularly in the land carbon cycle, may lead to an underestimation of carbon-cycle uncertainty and thus of the carbon ZEC term and CMIP7 generation models with increased carbon cycle complexity may reveal a wider spread in carbon ZEC responses.

920 *Code and data availability.* The flat10MIP ESM and SCM data used in this study are archived at <https://doi.org/10.5281/zenodo.15267556> (Sanderson et al., 2025b), as described in Sanderson et al. (2025a). Code is made available to reviewers and will be publicly released upon publication.

*Author contributions.* BMS conceived the study, performed the analysis, and wrote the manuscript, building on earlier framework derived by SJ. MS, SJ, GP and ARP reviewed and assessed and helped revise the analytical solution. CK motivated carbon/thermal summary and out-of-sample failure analysis. All authors reviewed and edited the manuscript text. Nadir Jeevanjee is acknowledged for comments and suggestions on an internal review of the manuscript.

925 *Competing interests.* At least one of the (co-)authors is a member of the editorial board of Earth System Dynamics. The peer-review process was guided by an independent editor, and the authors also have no other competing interests to declare.

*Disclaimer.* Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims made in the text, published maps, institutional affiliations, or any other geographical representation in this paper. While Copernicus Publications makes every effort to include appropriate place names, the final responsibility lies with the authors.

930 *Acknowledgements.* Alejandro Romero Prieto was supported by the Leeds-York-Hull Natural Environment Research Council (NERC) Doctoral Training Partnership (DTP) Panorama under grant NE/S007458/1.



## 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, <https://doi.org/10.1038/nature08019>, 2009.
- 935 Allen, M. R., Friedlingstein, P., Girardin, C. A., Jenkins, S., Malhi, Y., Mitchell-Larson, E., Peters, G. P., and Rajamani, L.: Net zero: science, origins, and implications, *Annual Review of Environment and Resources*, 47, 849–887, <https://doi.org/10.1146/annurev-environ-112320-105050>, 2022.
- Andrews, T., Gregory, J. M., and Webb, M. J.: The dependence of radiative forcing and feedback on evolving patterns of surface temperature change in climate models, *Journal of Climate*, 28, 1630–1648, <https://doi.org/10.1175/JCLI-D-14-00545.1>, 2015.
- 940 Armour, K. C., Bitz, C. M., and Roe, G. H.: Time-varying climate sensitivity from regional feedbacks, *Journal of Climate*, 26, 4518–4534, <https://doi.org/10.1175/JCLI-D-12-00544.1>, 2013.
- Armour, K. C., Proistosescu, C., Dong, Y., Hahn, L. C., Blanchard-Wrigglesworth, E., Pauling, A. G., Wjffels, R. C. J., Andrews, T., Smith, M. F., Palter, S. P.-C., Middlemas, I., Forster, P. M., and Gregory, J. M.: Sea-surface temperature pattern effects have slowed global warming and biased warming-based constraints on climate sensitivity, *Proceedings of the National Academy of Sciences*, 121, e2312093 121, <https://doi.org/10.1073/pnas.2312093121>, 2024.
- 945 Arora, V. K., Katavouta, A., Williams, R. G., Jones, C. D., Brovkin, V., Friedlingstein, P., Schwinger, J., Bopp, L., Boucher, O., Cadule, P., Chamberlain, M. A., Christian, J. R., Delire, C., Fisher, R. A., Hajima, T., Ilyina, T., Joetzjer, E., Kawamiya, M., Koven, C. D., Krasting, J. P., Law, R. M., Lawrence, D. M., Lenton, A., Lindsay, K., Pongratz, J., Raddatz, T., Séférian, R., Tachiiri, K., Tjiputra, J. F., Wiltshire, A., Wu, T., and Ziehn, T.: Carbon–concentration and carbon–climate feedbacks in CMIP6 models and their comparison to CMIP5 models, *Biogeosciences*, 17, 4173–4222, <https://doi.org/10.5194/bg-17-4173-2020>, 2020.
- 950 Bossert, I., Baur, S., Cussac, M., et al.: CNRM-ESM2-2: development and evaluation of an updated version of the CNRM-CERFACS Earth System Model, *ESS Open Archive [preprint]*, <https://doi.org/10.22541/essoar.175977424.42948487/v1>, 2025.
- Brovkin, V., Sanderson, B. M., Brizuela, N. G., Hajima, T., Ilyina, T., Jones, C. D., Koven, C., Lawrence, D., Lawrence, P., Li, H., Liddicoat, S., Romanou, A., Séférian, R., Sentman, L. T., Swann, A. L. S., Tjiputra, J., Ziehn, T., and Winkler, A. J.: On a simplified solution of climate–carbon dynamics in idealized flat10MIP simulations, *Earth System Dynamics*, 16, 2021–2038, <https://doi.org/10.5194/esd-16-2021-2025>, 2025.
- 955 Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C.: A limited memory algorithm for bound constrained optimization, *SIAM Journal on Scientific Computing*, 16, 1190–1208, <https://doi.org/10.1137/0916069>, 1995.
- Cheng, L., Abraham, J., Trenberth, K. E., Fasullo, J., Boyer, T., Locarnini, R., Zhang, B., Yu, F., Wan, L., Chen, X., Song, X., Liu, Y., and Mann, M. E.: Another record: Ocean warming continues through 2021 despite La Niña conditions, *Advances in Atmospheric Sciences*, 39, 373–385, <https://doi.org/10.1007/s00376-022-1461-3>, 2022.
- 960 Danabasoglu, G., Lamarque, J.-F., Bacmeister, J., Bailey, D. A., DuVivier, A. K., Edwards, J., Emmons, L. K., Fasullo, J., Garcia, R., Gettelman, A., Hannay, C., Holland, M. M., Large, W. G., Lauritzen, P. H., Lawrence, D. M., Lenaerts, J. T. M., Lindsay, K., Lipscomb, W. H., Mills, M. J., Neale, R., Oleson, K. W., Otto-Bliesner, B., Phillips, A. S., Sacks, W., Tilmes, S., van Kampenhout, L., Vertenstein, M., Bertini, A., Dennis, J., Deser, C., Fischer, C., Fox-Kemper, B., Kay, J. E., Kinnison, D., Kushner, P. J., Larson, V. E., Long, M. C., Mickelson, S., Moore, J. K., Nienhouse, E., Polvani, L., Rasch, P. J., and Strand, W. G.: The Community Earth System Model version 2 (CESM2), *J. Adv. Model. Earth Syst.*, 12, e2019MS001 916, <https://doi.org/10.1029/2019MS001916>, 2020.



- Davies-Barnard, T., Meyerholt, J., Zaehle, S., Friedlingstein, P., Brovkin, V., Fan, Y., Fisher, R. A., Jones, C. D., Lee, H., Peano, D., Smith, B., Wiltshire, A., and Yue, X.: Nitrogen cycling in CMIP6 land surface models: progress and limitations, *Biogeosciences*, 17, 5129–5148, <https://doi.org/10.5194/bg-17-5129-2020>, 2020.
- De Sisto, M. L. and MacDougall, A. H.: Effect of terrestrial nutrient limitation on the estimation of the remaining carbon budget, *Biogeosciences*, 21, 4853–4873, <https://doi.org/10.5194/bg-21-4853-2024>, 2024.
- Dunne, J. P., Horowitz, L. W., Adcroft, A. J., Ginoux, P., Held, I. M., John, J. G., Krasting, J. P., Malyshev, S., Naik, V., Paulot, F., Shevliakova, E., Stock, C. A., Zadeh, N., Balaji, V., Blanton, C., Dunne, K. A., Dupuis, C., Durachta, J., Dussin, R., Gauthier, P. P. G., Griffies, S. M., Guo, H., Hallberg, R. W., Harrison, M., He, J., Hurlin, W., McHugh, C., Menzel, R., Milly, P. C. D., Nikonov, S., Paynter, D. J., Ploshay, J., Radhakrishnan, A., Rand, K., Reichl, B. G., Robinson, T., Schwarzkopf, D. M., Sentman, L. T., Underwood, S., Vahlenkamp, H., Winton, M., Wittenberg, A. T., Wyman, B., Zeng, Y., and Zhao, M.: The GFDL Earth System Model version 4.1 (GFDL-ESM4.1): overall coupled model description and simulation characteristics, *J. Adv. Model. Earth Syst.*, 12, e2019MS002015, <https://doi.org/10.1029/2019MS002015>, 2020.
- Dunne, J. P., Hewitt, H. T., Arblaster, J., Bonou, F., Boucher, O., Cavazos, T., Durack, P. J., Hassler, B., Juckes, M., Miyakawa, T., Mizielinski, M., Naik, V., Nicholls, Z., O'Rourke, E., Pincus, R., Sanderson, B. M., Simpson, I. R., and Taylor, K. E.: An evolving Coupled Model Intercomparison Project phase 7 (CMIP7) and Fast Track in support of future climate assessment, *Geoscientific Model Development*, 18, 6671–6696, <https://doi.org/10.5194/gmd-18-6671-2025>, 2025.
- Eby, M., Zickfeld, K., Montenegro, A., Archer, D., Meissner, K. J., and Weaver, A. J.: Lifetime of anthropogenic climate change: Millennial time scales of potential CO<sub>2</sub> and surface temperature perturbations, *Journal of Climate*, 22, 2501–2511, <https://doi.org/10.1175/2008JCLI2554.1>, 2009.
- Fisher, R. A. and Koven, C. D.: Dynamic Global Vegetation Models: searching for the balance between demographic process representation and computational tractability, *PLOS Climate*, 1, e0000068, <https://doi.org/10.1371/journal.pclm.0000068>, 2022.
- Fisher, R. A., Koven, C. D., Anderegg, W. R. L., Christoffersen, B. O., Dietze, M. C., Farrior, C. E., Holm, J. A., Hurtt, G. C., Knox, R. G., Lawrence, P. J., Lichstein, J. W., Longo, M., Matheny, A. M., Medvigy, D., Muber-Landau, H. C., Powell, T. L., Serbin, S. P., Sber, H., Weng, E. S., Wood, J. D., and Moorcroft, P. A.: Vegetation demographics in Earth System Models: a review of progress and priorities, *Global Change Biology*, 24, 35–54, <https://doi.org/10.1111/gcb.13910>, 2018.
- Forster, P. M., Smith, C. J., Walsh, T., Lamb, W. F., Lamboll, R., Hall, B., Hauser, M., Ribes, A., Rosen, D., Gillett, N. P., Palmer, M. D., Rogelj, J., von Schuckmann, K., Trewin, B., Allen, M., Andrew, R., Betts, R. A., Borber, A., Canadell, J., et al.: Indicators of Global Climate Change 2024: annual update of key indicators of the state of the climate system and human influence, *Earth Syst. Sci. Data*, 17, 2641–2724, <https://doi.org/10.5194/essd-17-2641-2025>, 2025.
- Friedlingstein, P., Jones, M. W., O'Sullivan, M., Andrew, R. M., Bakker, D. C. E., Hauck, J., Le Quéré, C., Peters, G. P., Peters, W., Pongratz, J., et al.: Global carbon budget 2021, *Earth System Science Data*, 14, 1917–2005, <https://doi.org/10.5194/essd-14-1917-2022>, 2022.
- Geoffroy, O., Saint-Martin, D., Olivié, D. J. L., 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, <https://doi.org/10.1175/JCLI-D-12-00195.1>, 2013.
- Gier, B. K., Schlund, M., Friedlingstein, P., Jones, C. D., Jones, C., Zaehle, S., and Eyring, V.: Representation of the terrestrial carbon cycle in CMIP6, *Biogeosciences*, 21, 5321–5360, <https://doi.org/10.5194/bg-21-5321-2024>, 2024.
- Goodwin, P.: How historic simulation–observation discrepancy affects future warming projections in a very large model ensemble, *Clim. Dyn.*, 47, 2219–2233, <https://doi.org/10.1007/s00382-015-2960-z>, 2016.



- Hansen, J. E., Sato, M., Simons, L., Nazarenko, L. S., Sangha, I., von Schuckmann, K., Loeb, N. G., Osman, M. B., Jin, Q., Tselioudis, G., Lacic, A., Ruedy, R., Russell, G., Cao, J., and Li, J.: Global warming in the pipeline, *Oxford Open Climate Change*, 3, kgad008, <https://doi.org/10.1093/oxfclm/kgad008>, 2023.
- 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, <https://doi.org/10.1175/2009JCLI3466.1>, 2010.
- Jeevanjee, N., Paynter, D. J., Dunne, J. P., Sentman, L. T., and Krasting, J. P.: A holistic view of climate sensitivity, *Annual Review of Earth and Planetary Sciences*, 53, <https://doi.org/10.1146/annurev-earth-040523-014302>, 2025.
- Jenkins, S., Sanderson, B. M., Peters, G. P., Aengenheyster, M., Smith, C. J., and Allen, M. R.: The multi-decadal response to net zero CO<sub>2</sub> emissions and implications for emissions policy, *Geophysical Research Letters*, 49, e2022GL101047, <https://doi.org/10.1029/2022GL101047>, 2022.
- Jones, C. D. and Friedlingstein, P.: Quantifying process-level uncertainty contributions to TCRE and carbon budgets for meeting Paris Agreement climate targets, *Environmental Research Letters*, 15, 074 019, <https://doi.org/10.1088/1748-9326/ab858a>, 2020.
- Joos, F., Bruno, M., Fink, R., Siegenthaler, U., Stocker, T. F., Le Quéré, C., and Sarmiento, J. L.: An efficient and accurate representation of complex oceanic and biospheric models of anthropogenic carbon uptake, *Tellus B*, 48, 397–417, <https://doi.org/10.3402/tellusb.v48i3.15921>, 1996.
- Joos, F., Roth, R., Fuglestedt, J. S., Peters, G. P., Enting, I. G., von Bloh, W., Brovkin, V., Burke, E. J., Eby, M., Edwards, N. R., Friedrich, T., Frölicher, T. L., Halloran, P. R., Holden, P. B., Jones, C., Kleinen, T., Mackenzie, F. T., Matsumoto, K., Meinshausen, M., Plattner, G.-K., Reisinger, A., Segschneider, J., Shaffer, G., Steinacher, M., Strassmann, K., Tanaka, K., Timmermann, A., and Weaver, A. J.: 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, <https://doi.org/10.5194/acp-13-2793-2013>, 2013.
- Katavouta, A., Williams, R. G., Goodwin, P., and Roussenov, V.: The effect of ocean ventilation on the transient climate response to emissions, *Journal of Climate*, 32, 5085–5105, <https://doi.org/10.1175/JCLI-D-18-0829.1>, 2019.
- Kelley, M., Schmidt, G. A., Nazarenko, L. S., Bauer, S. E., Ruedy, R., Russell, G. L., Ackerman, A. S., Aleinov, I., Bauer, M., Bleck, R., Canuto, V., Cesana, G., Cheng, Y., Clune, T. L., Cook, B. I., Cruz, C. A., Del Genio, A. D., Elsaesser, G. S., Faluvegi, G., Kiang, N. Y., Kim, D., Lacic, A. A., Leboissetier, A., LeGrande, A. N., Lo, K. K., Marshall, J., Matthews, E. E., McDermid, S., Mezuman, K., Miller, R. L., Murray, L. T., Oinas, V., Orbe, C., García-Pando, C. P., Perlwitz, J. P., Puma, M. J., Rind, D., Romanou, A., Shindell, D. T., Sun, S., Tausnev, N., Tsigaridis, K., Tselioudis, G., Weng, E., Wu, J., and Yao, M.-S.: GISS-E2.1: configurations and climatology, *J. Adv. Model. Earth Syst.*, 12, e2019MS002 025, <https://doi.org/10.1029/2019MS002025>, 2020.
- Koven, C. D., Arora, V. K., Cadule, P., Fisher, R. A., Jones, C. D., Lawrence, D. M., Lewis, J., Lindsay, K., Mathesius, S., Meinshausen, M., Mills, M., Nicholls, Z., Sanderson, B. M., Séférian, R., Swart, N. C., Wieder, W. R., and Ziehn, T.: Multi-century dynamics of the climate and carbon cycle under both high and net negative emissions scenarios, *Earth Syst. Dynam.*, 13, 885–909, <https://doi.org/10.5194/esd-13-885-2022>, 2022.
- Koven, C. D., Sanderson, B. M., and Swann, A. L. S.: Much of zero emissions commitment occurs before reaching net zero emissions, *Environmental Research Letters*, 18, 014 017, <https://doi.org/10.1088/1748-9326/acab1a>, 2023.
- Li, C., von Storch, J.-S., and Marotzke, J.: Deep-ocean heat uptake and equilibrium climate response, *Climate Dynamics*, 40, 1071–1086, <https://doi.org/10.1007/s00382-012-1350-z>, 2013.



- Loisel, J., Gallego-Sala, A. V., Amesbury, M. J., Magnan, G., Anshari, G., Beilman, D. W., Benavides, J. C., Blewett, J., Camill, P., Charman, D. J., et al.: Expert assessment of future vulnerability of the global peatland carbon sink, *Nature Climate Change*, 11, 70–77, <https://doi.org/10.1038/s41558-020-00944-0>, 2021.
- 1045 MacDougall, A. H.: Estimated effect of the permafrost carbon feedback on the zero emissions commitment to climate change, *Biogeosciences*, 18, 4937–4952, <https://doi.org/10.5194/bg-18-4937-2021>, 2021.
- MacDougall, A. H. and Friedlingstein, P.: The origin and limits of the near proportionality between climate warming and cumulative CO<sub>2</sub> emissions, *Journal of Climate*, 28, 4217–4230, <https://doi.org/10.1175/JCLI-D-14-00036.1>, 2015.
- MacDougall, A. H., Frölicher, T. L., Jones, C. D., Rogelj, J., Matthews, H. D., Zickfeld, K., Arora, V. K., Barrett, N. J., Brovkin, V., 1050 Burger, F. A., Eby, M., Eliseev, A. V., Hajima, T., Holden, P. B., Jeltsch-Thömmes, A., Koven, C., Mengis, N., Menviel, L., Michou, M., Mokhov, I. I., Oka, A., Schwinger, J., Séférian, R., Shaffer, G., Sokolov, A., Tachiiri, K., Tjiputra, J., Wiltshire, A., and Ziehn, T.: Is there warming in the pipeline? A multi-model analysis of the Zero Emissions Commitment from CO<sub>2</sub>, *Biogeosciences*, 17, 2987–3016, <https://doi.org/10.5194/bg-17-2987-2020>, 2020.
- Maier-Reimer, E. and Hasselmann, K.: Transport and storage of CO<sub>2</sub> in the ocean — an inorganic ocean-circulation carbon cycle model, 1055 *Climate Dynamics*, 2, 63–90, <https://doi.org/10.1007/BF01054491>, 1987.
- 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, <https://doi.org/10.1038/nature08047>, 2009.
- Mauritsen, T., Bader, J., Becker, T., Behrens, J., Bittner, M., Brokopf, R., Brovkin, V., Claussen, M., Crueger, T., Esch, M., Fast, I., Fiedler, S., Fläschner, D., Gayler, V., Giorgetta, M., Goll, D. S., Haak, H., Hagemann, S., Hedemann, C., Hohenegger, C., Ilyina, T., Jahns, 1060 T., Jimenez-de-la Cuesta, D., Jungclaus, J., Kleinen, T., Kloster, S., Kracher, D., Kinne, S., Kleberg, D., Lasslop, G., Kornbluh, L., Marotzke, J., Matei, D., Meraner, K., Mikolajewicz, U., Modali, K., Möbis, B., Müller, W. A., Nabel, J. E. M. S., Nam, C. C. W., Notz, D., Nyawira, S.-S., Paulsen, H., Peters, K., Pincus, R., Pohlmann, H., Pongratz, J., Popp, M., Raddatz, T. J., Rast, S., Redler, R., Reick, C. H., Rohrschneider, T., Schemann, V., Schmidt, H., Schnur, R., Schulzweida, U., Six, K. D., Stein, L., Stemmler, I., Stevens, B., von Storch, J.-S., Tian, F., Voigt, A., Vrese, P., Wieners, K.-H., Wilkenskjaeld, S., Winkler, A., and Roeckner, E.: Developments in the 1065 MPI-M Earth System Model version 1.2 (MPI-ESM1.2) and its response to increasing CO<sub>2</sub>, *J. Adv. Model. Earth Syst.*, 11, 998–1038, <https://doi.org/10.1029/2018MS001400>, 2019.
- Mengis, N., Keller, D. P., MacDougall, A. H., Eby, M., Wright, N., Meissner, K. J., Oeschli, A., Schmittner, A., MacIsaac, A. J., Matthews, H. D., and Zickfeld, K.: Evaluation of the University of Victoria Earth System Climate Model version 2.10 (UVic ESCM2.10), *Geosci. Model Dev.*, 13, 4183–4204, <https://doi.org/10.5194/gmd-13-4183-2020>, 2020.
- 1070 Millar, R. J., Nicholls, Z. R., Friedlingstein, P., and Allen, M. R.: A modified impulse-response representation of the global near-surface air temperature and atmospheric concentration response to carbon dioxide emissions, *Atmospheric Chemistry and Physics*, 17, 7213–7228, <https://doi.org/10.5194/acp-17-7213-2017>, 2017.
- Müller, J. and Joos, F.: Global peatland area and carbon dynamics from the Last Glacial Maximum to the present – a process-based model investigation, *Biogeosciences*, 17, 5285–5308, <https://doi.org/10.5194/bg-17-5285-2020>, 2020.
- 1075 Myhre, G., Highwood, E. J., Shine, K. P., and Stordal, F.: New estimates of radiative forcing due to well mixed greenhouse gases, *Geophysical Research Letters*, 25, 2715–2718, <https://doi.org/10.1029/98GL01908>, 1998.
- Myhre, G., Hodnebrog, Ø., Loeb, N., and Forster, P. M.: Observed trend in Earth energy imbalance may provide a constraint for low climate sensitivity models, *Science*, 388, 1210–1213, <https://doi.org/10.1126/science.adt0647>, 2025.



- Needham, J. F., Chambers, J., Fisher, R., Knox, R., and Koven, C. D.: Forest responses to simulated elevated CO<sub>2</sub> under alternate hypotheses of size-and age-dependent mortality, *Global Change Biology*, 26, 5734–5753, <https://doi.org/10.1111/gcb.15254>, 2020.
- 1080 Palazzo Corner, S., Siegert, M., Ceppi, P., Fox-Kemper, B., Frölicher, T. L., Gallego-Sala, A., Haigh, J., Hegerl, G. C., Jones, C. D., Knutti, R., Koven, C. D., MacDougall, A. H., Meinshausen, M., Nicholls, Z., Sallée, J.-B., Sanderson, B. M., Séférian, R., Turetsky, M., Williams, R. G., Zaehle, S., and Rogelj, J.: The Zero Emissions Commitment and climate stabilization, *Frontiers in Science*, 1, 1170 744, <https://doi.org/10.3389/fsci.2023.1170744>, 2023.
- 1085 Pugh, T. A. M., Arneth, A., Kautz, M., Poulter, B., and Smith, B.: Important role of forest disturbances in the global biomass turnover and carbon sinks, *Nature Geoscience*, 12, 730–735, <https://doi.org/10.1038/s41561-019-0427-2>, 2019.
- Raper, S. C. B., Gregory, J. M., and Stouffer, R. J.: The role of climate sensitivity and ocean heat uptake on AOGCM transient temperature response, *Journal of Climate*, 15, 124–130, [https://doi.org/10.1175/1520-0442\(2002\)015<0124:TROCSA>2.0.CO;2](https://doi.org/10.1175/1520-0442(2002)015<0124:TROCSA>2.0.CO;2), 2002.
- Rogelj, J., Forster, P. M., Kriegler, E., Smith, C. J., and Séférian, R.: Estimating and tracking the remaining carbon budget for stringent climate targets, *Nature*, 571, 335–342, <https://doi.org/10.1038/s41586-019-1368-z>, 2019.
- 1090 Romero-Prieto, A., Sandstad, M., Sanderson, B. M., Nicholls, Z. R. J., Steinert, N. J., Gasser, T., Mathison, C., Kikstra, J., Aubry, T. J., and Smith, C.: Reduced Complexity Model Intercomparison Project Phase 3: Experimental protocol for coordinated constraining and evaluation of Reduced-Complexity Models, *EGUsphere* [preprint], <https://doi.org/10.5194/egusphere-2025-5775>, 2025.
- Romero-Prieto, A., Mathison, C., and Smith, C.: Review of climate simulation by Simple Climate Models, *Geoscientific Model Development*, 19, 115–165, <https://doi.org/10.5194/gmd-19-115-2026>, 2026.
- 1095 Sanderson, B. M., Booth, B. B. B., Dunne, J., Eyring, V., Fisher, R. A., Friedlingstein, P., Gidden, M. J., Hajima, T., Jones, C. D., Jones, C. G., King, A., Koven, C. D., Lawrence, D. M., Lowe, J., Mengis, N., Peters, G. P., Rogelj, J., Smith, C., Snyder, A. C., Simpson, I. R., Swann, A. L. S., Tebaldi, C., Ilyina, T., Schleussner, C.-F., Séférian, R., Samset, B. H., van Vuuren, D., and Zaehle, S.: The need for carbon-emissions-driven climate projections in CMIP7, *Geoscientific Model Development*, 17, 8141–8172, <https://doi.org/10.5194/gmd-17-8141-2024>, 2024.
- 1100 Sanderson, B. M., Brovkin, V., Fisher, R. A., Hohn, D., Ilyina, T., Jones, C. D., Koenigk, T., Koven, C., Li, H., Lawrence, D. M., Lawrence, P., Liddicoat, S., MacDougall, A. H., Mengis, N., Nicholls, Z., O'Rourke, E., Romanou, A., Sandstad, M., Schwinger, J., Séférian, R., Sentman, L. T., Simpson, I. R., Smith, C., Steinert, N. J., Swann, A. L. S., Tjiputra, J., and Ziehn, T.: flat10MIP: an emissions-driven experiment to diagnose the climate response to positive, zero and negative CO<sub>2</sub> emissions, *Geoscientific Model Development*, 18, 5699–5724, <https://doi.org/10.5194/gmd-18-5699-2025>, 2025a.
- 1105 Sanderson, B. M., Koven, C., Sandstad, M., and Nicholls, Z.: benmsanderson/esm-deck: GMD revised version, <https://doi.org/10.5281/zenodo.15267556>, 2025b.
- Schwinger, J. and Tjiputra, J. F.: Ocean carbon cycle feedbacks under negative emissions, *Geophys. Res. Lett.*, 45, 5062–5070, <https://doi.org/10.1029/2018GL077790>, 2018.
- 1110 Séférian, R., Nabat, P., Michou, M., Saint-Martin, D., Voldoire, A., Colin, J., Decharme, B., Delire, C., Berthet, S., Chevallier, M., et al.: Evaluation of CNRM Earth System Model, CNRM-ESM2-1: Role of Earth System Processes in Present-Day and Future Climate, *J. Adv. Model. Earth Syst.*, 11, 4182–4227, <https://doi.org/10.1029/2019MS001791>, 2019.
- Séférian, R., Bossy, T., Gasser, T., Nichols, Z., Dorheim, K., Su, X., Tsutsui, J., and Santana-Falcón, Y.: Physical inconsistencies in the representation of the ocean heat-carbon nexus in simple climate models, *Communications Earth & Environment*, 5, 291, <https://doi.org/10.1038/s43247-024-01464-x>, 2024.
- 1115



- Seland, Ø., Bentsen, M., Olivie, D., Toniazzo, T., Gjermundsen, A., Graff, L. S., Debernard, J. B., Gupta, A. K., He, Y.-C., Kirkevåg, A., Schwinger, J., Tjiputra, J., Aas, K. S., Bethke, I., Fan, Y., Griesfeller, J., Grini, A., Guo, C., Ilıcak, M., Karset, I. H. H., Landgren, O., Liakka, J., Moseid, K. O., Nummelin, A., Spensberger, C., Tang, H., Zhang, Z., Heinze, C., Iversen, T., and Schulz, M.: Overview of the Norwegian Earth System Model (NorESM2) and key climate response of CMIP6 DECK, historical, and scenario simulations, *Geosci. Model Dev.*, 13, 6165–6200, <https://doi.org/10.5194/gmd-13-6165-2020>, 2020.
- 1120 Sellar, A. A., Jones, C. G., Mulcahy, J. P., Tang, Y., Yool, A., Wiltshire, A., O'Connor, F. M., Stringer, M., Hill, R., Palmieri, J., Woodward, S., Mora, L. d., Kuhlbrodt, T., Rumbold, S. T., Kelley, D. I., Ellis, R., Johnson, C. E., Walton, J., Abraham, N. L., Andrews, M. B., Andrews, T., Sherlock, V., Sherlock, C., Burke, E., Blockley, E., Carslaw, K., Dalvi, M., Edwards, J., Folberth, G. A., Gedney, N., Griffiths, P. T., Harper, A. B., Hendry, M. A., Hewitt, A. J., Johnson, B., Jones, A., Jones, C. D., Keeble, J., Liddicoat, S., Morgenstern, O., Parker, R. J.,
- 1125 Predoi, V., Robertson, E., Siahna, A., Smith, R. S., Swaminathan, R., Sherlock, M. T., Zeng, G., and Zerroukat, M.: UKESM1: description and evaluation of the UK Earth System Model, *J. Adv. Model. Earth Syst.*, 11, 4513–4558, <https://doi.org/10.1029/2019MS001739>, 2019.
- Seshadri, A. K.: Fast–slow climate dynamics and peak global warming, *Climate Dynamics*, 48, 2235–2253, <https://doi.org/10.1007/s00382-016-3202-8>, 2017.
- Sherwood, S. C., Webb, M. J., Annan, J. D., Armour, K. C., Forster, P. M., Hargreaves, J. C., Hegerl, G., Klein, S. A., Marvel, K. D., Rohling, E. J., et al.: An assessment of Earth's climate sensitivity using multiple lines of evidence, *Reviews of geophysics*, 58, e2019RG000678, <https://doi.org/10.1029/2019RG000678>, 2020.
- 1130 Smith, C. J., Forster, P. M., Allen, M., Leach, N., Millar, R. J., Passerello, G. A., and Regayre, L. A.: FAIR v1.3: a simple emissions-based impulse response and carbon cycle model, *Geoscientific Model Development*, 11, 2273–2297, <https://doi.org/10.5194/gmd-11-2273-2018>, 2018.
- 1135 Steinert, N. J. and Sanderson, B. M.: Normalizing the permafrost carbon feedback contribution to the Transient Climate Response to Cumulative Carbon Emissions and the Zero Emissions Commitment, *Earth System Dynamics*, 16, 1711–1721, <https://doi.org/10.5194/esd-16-1711-2025>, 2025.
- Stouffer, R. and Manabe, S.: Equilibrium response of thermohaline circulation to large changes in atmospheric CO<sub>2</sub> concentration, *Climate Dynamics*, 20, 759–773, <https://doi.org/10.1007/s00382-002-0302-4>, 2003.
- 1140 Valdes, P. J., Armstrong, E., Badger, M. P. S., Bradshaw, C. D., Bragg, F., Crucifix, M., Davies-Barnard, T., Day, J. J., Farnsworth, A., Gordon, C., Hopercoft, P. O., Kennedy, A. T., Lord, N. S., Lunt, D. J., Marzocchi, A., Parry, L. M., Pope, V., Roberts, W. H. G., Stone, E. J., Tourte, G. J. L., and Williams, J. H. T.: The BRIDGE HadCM3 family of climate models: HadCM3@Bristol v1.0, *Geosci. Model Dev.*, 10, 3715–3743, <https://doi.org/10.5194/gmd-10-3715-2017>, 2017.
- Varney, R. M., Chadburn, S. E., Burke, E. J., and Jones, C. D.: Soil carbon-concentration and carbon-climate feedbacks in CMIP6 Earth system models, *Biogeosciences*, 21, 2759–2776, <https://doi.org/10.5194/bg-21-2759-2024>, 2024.
- 1145 von Schuckmann, K., Minière, A., Gues, F., Cuesta-Valero, F. J., Kirchengast, G., Adusumilli, S., Straneo, F., Ablain, M., Allan, R. P., Barker, P. M., et al.: Heat stored in the Earth system 1960–2020: where does the energy go?, *Earth System Science Data*, 15, 1675–1709, <https://doi.org/10.5194/essd-15-1675-2023>, 2023.
- Wieder, W. R., Cleveland, C. C., Smith, W. K., and Todd-Brown, K.: Future productivity and carbon storage limited by terrestrial nutrient availability, *Nature Geoscience*, 8, 441–444, <https://doi.org/10.1038/ngeo2413>, 2015.
- 1150 Williams, R. G., Goodwin, P., Ceppi, P., Jones, C. D., and MacDougall, A. H.: A normalised framework for the Zero Emissions Commitment, *Biogeosciences*, 22, 7167–7189, <https://doi.org/10.5194/bg-22-7167-2025>, 2025.



- Winkelmann, R., Dennis, D. P., Donges, J. F., Loriani, S., Klose, A. K., Abrams, J. F., Alvarez-Solas, J., Albrecht, T., Armstrong McKay, D., Bathiany, S., Blasco Navarro, J., Brovkin, V., Burke, E., Danabasoglu, G., Donner, R. V., Drüke, M., Georgievski, G., Goelzer, H., Harper, A. B., Hegerl, G., Hirota, M., Hu, A., Jackson, L. C., Jones, C., Kim, H., Koenigk, T., Lawrence, P., Lenton, T. M., Liddy, H., Licón-Salaíz, J., Menthon, M. d., Montoya, M., Nitzbon, J., Nowicki, S., Otto-Bliesner, B., Pausata, F., Rahmstorf, S., Ramin, K., Robinson, A., Rockström, J., Romanou, A., Sakschewski, B., Schädel, C., Sherwood, S., Smith, R. S., Steinert, N. J., Swingedouw, D., Willeit, M., Weijer, W., Wood, R., Wyser, K., and Yang, S.: The Tipping Points Modelling Intercomparison Project (TIPMIP): assessing tipping point risks in the Earth system, EGU sphere [preprint], <https://doi.org/10.5194/egusphere-2025-1899>, 2025.
- 1155
- Winton, M., Takahashi, K., and Held, I. M.: Importance of ocean heat uptake efficacy to transient climate change, *Journal of Climate*, 23, 2333–2344, <https://doi.org/10.1175/2009JCLI3139.1>, 2010.
- 1160
- Zickfeld, K., Eby, M., Matthews, H. D., and Weaver, A. J.: Setting cumulative emissions targets to reduce the risk of dangerous climate change, *Proceedings of the National Academy of Sciences*, 106, 16 129–16 134, <https://doi.org/10.1073/pnas.0805800106>, 2009.
- Ziehn, T., Chamberlain, M. A., Law, R. M., Lenton, A., Bodman, R. W., Dix, M., Stevens, L., Wang, Y.-P., and Srbnovsky, J.: The Australian Earth System Model: ACCESS-ESM1.5, *J. South. Hemisph. Earth Syst. Sci.*, 70, 193–214, <https://doi.org/10.1071/ES19035>, 2020.
- 1165