

Sustained Neutralization of the Warming Response to Emissions through a Portfolio of GHG Mitigation Strategies

Eleanor Hughes^{1*}, Zeke Hausfather², Randall Spock¹, & Christopher Van Arsdale¹

¹Google LLC, Mountain View, CA, USA

²Stripe, LLC, South San Francisco, CA, USA

*Corresponding author

Abstract

Both emissions abatement and carbon dioxide removal will play a critical role in achieving a global net-zero scenario under which temperatures are stabilized. In the near term, organizations pursuing their own net-zero targets may use superpollutant emission abatement and both shorter- and longer-duration carbon dioxide removal to compensate for their remaining emissions. Because these mitigation strategies have climate impacts across distinct timescales, they can serve as powerful and complementary tools to address both the acute near-term and long-term equilibrium warming response to emissions. However, the distinct timescales of these impacts mean that, if accounting approaches are to ensure emissions remain compensated across all time horizons, they must explicitly evaluate the impact of each mitigation strategy as a function of time. We propose one such accounting approach, based on sustained warming neutralization. By translating the climate impacts of both emissions and mitigation into a single metric - the continuous global mean surface temperature response - the approach aims to ensure the neutralization of the warming response to an emission in both the near- and long-term. Ultimately, this warming neutralization approach structurally aligns emission compensation claims with global goals for temperature stabilization. We then introduce a specific method through which an organization can use the approach to design mitigation portfolios to compensate for its emissions. We invite future work to consider how the warming neutralization approach could be operationalized in a way that ensures impactful, accessible, and widely recognized emissions claims.

1. Introduction

Greenhouse gas (GHG) mitigation plays a critical role in achieving a net-zero scenario under which global temperatures are stabilized (Riahi et al. 2023). In the near term, ‘credits’ for GHG mitigation achieved outside of the scope of an organization’s operations provide important instruments for organizations to compensate for their remaining emissions and, in so doing, help to scale the long-term mitigation strategies that will be essential in global net-zero or overshoot scenarios. These credits include: carbon dioxide removal (CDR) with storage durability comparable to the atmospheric lifetime of CO₂; CDR with shorter or more uncertain storage durability; and superpollutant emission abatement. Superpollutants are a class of GHGs that possess significantly higher near-term warming potentials compared to CO₂. The subset of superpollutants with atmospheric lifetimes shorter than CO₂ – including methane and some F-gases – are commonly referred to as short-lived climate pollutants, or SLCPs.

Table 1. Atmospheric lifetimes and Global Warming Potentials (GWPs) of major greenhouse gases (Forster et al. 2023). GWP is a metric used to compare the integrated radiative forcing of a specific mass of a greenhouse gas to an equivalent mass of CO₂. GWP100 and GWP20 represent this integrated forcing relative to CO₂ over 100-year and 20-year time horizons, respectively. *No single atmospheric lifetime can be defined for CO₂ because it is removed via multiple processes operating on different timescales. Based on current ocean and land carbon sink strength, approximately 50 percent is removed within 30 years, an additional 30 percent within a few centuries, and the remaining 20 percent persists for thousands of years (Joos et al. 2013). Note that, at the scale of corporate emissions, the processes that remove atmospheric CO₂ apply approximately symmetrically to CDR (see Section A.1.4.3) (Zickfeld et al. 2021).

Greenhouse Gas (GHG)	Atmospheric Lifetime (Years)	GWP20	GWP100
Carbon dioxide (CO ₂)	Multiple*	1	1
Methane (CH ₄)	11.8	81.2	27.9
HFC-134a (CH ₂ FCHF ₃)	14	4140	1530
HFC-32 (CH ₂ F ₂)	5.4	2690	771
Nitrous oxide (N ₂ O)	109	273	273

One way to use GHG mitigation to compensate for hard-to-abate emissions is on a like-for-like mass basis (Allen et al. 2022). For example, long-lived fossil CO₂ emissions would be paired with an equivalent number of tonnes of CDR with durability comparable to the CO₂ emission’s lifetime (hereafter ‘long-lived’ CDR), and CH₄ emissions would be paired with an equivalent number of tonnes of CH₄ abatement. However, organizations commonly seek to compensate for hard-to-abate emissions – most often a ‘residual’ of fossil CO₂ emissions – through a portfolio of GHG mitigation strategies that are not strictly matched to the GHG composition of the emissions. A portfolio can incorporate superpollutant abatement and CDR with shorter or more uncertain durability relative to the CO₂ emission’s lifetime (hereafter ‘shorter-lived’ CDR) alongside long-lived CDR.

Within such a portfolio, long-lived CDR remains essential in addressing the millennial-duration warming driven by fossil CO₂ emissions. Due to the dynamics of ocean heat uptake, this warming persists significantly beyond the atmospheric residence of the CO₂ itself (MacDougall et al. 2020). However, the

projected supply of long-lived CDR is not expected to meet the quantity of hard-to-abate emissions in time to limit peak warming over the coming decades (S. M. Smith et al. 2024). Inclusion of SLCP mitigation and shorter-lived CDR can help to address this near-term capacity constraint, as both strategies can be delivered through comparably mature and scaled technologies relative to early-stage engineered CDR. Shorter-lived CDR, including many nature-based pathways, can fully neutralize the warming response to CO₂ emissions for the duration of their storage relative to a non-intervention counterfactual. Nature-based CDR may also have considerable ecological co-benefits (for example, through reforestation). Concurrently, because SLCPs have higher warming potentials and shorter atmospheric lifetimes than CO₂, the destruction of a given mass of SLCP emissions can have a greater impact on near-term warming than a CO₂-equivalent mass of CDR (Szopa et al. 2023) (Figure 1). This makes SLCP mitigation a highly effective strategy for neutralizing acute near-term warming, including the potential tipping points and rate-based system risks that near-term net-zero goals seek to avoid (Wang et al. 2023; Fischer et al. 2021). SLCP mitigation and shorter-lived CDR therefore form complementary strategies to long-lived CDR, neutralizing near-term warming while long-lived CDR achieves the scale ultimately required to neutralize long-term equilibrium warming (Matthews et al. 2022; Ocko et al. 2021; Streck et al. 2025).

Many organizations therefore opt for a mitigation portfolio that includes SLCP mitigation and shorter-lived CDR alongside long-lived CDR. However, using SLCP mitigation or shorter-lived CDR requires careful accounting, because their impacts are shorter lived than the warming response to a fossil CO₂ emission (Figure 1) (Allen et al. 2022; Pierrehumbert 2014; Brunner et al. 2024). To neutralize the long-term warming response, SLCP mitigation or shorter-lived CDR therefore need to be replaced – either by additional and repeated SLCP mitigation or shorter-lived CDR, or by a single instance of long-lived CDR – as their impacts dissipate. This requirement creates choices regarding the quantity of SLCP mitigation or shorter-lived CDR to purchase, and the schedule on which it is replaced with more SLCP mitigation, shorter-lived CDR, or long-lived CDR.

Here, we propose a warming neutralization approach to explicitly account for these different timescales of impact (Section 2). The approach aims to provide an accounting framework through which a portfolio of GHG mitigation strategies can achieve sustained neutralization of the warming response to emissions. We then introduce a specific method to demonstrate how organizations could design mitigation portfolios under the approach (Section 3). Finally, we briefly outline the financial structures that may be required alongside the approach (Section 4). We invite future work to consider how the approach could be operationalized in a way that ensures impactful, accessible, and widely recognized emission compensation claims.

2. The Warming Neutralization Approach

An accounting approach that accommodates a portfolio of mitigation strategies needs to have two core components: a specific measure of the climate impact of emissions and mitigation, and a defined time over which to evaluate that impact. Conventional frameworks have relied on radiative forcing (RF) – which represents the instantaneous energy imbalance in the Earth's atmosphere resulting from a change in GHG concentration – as the measure of climate impact. To evaluate impact over time, the RF values are then translated into a global warming potential (GWP; Table 1).

The use of GWP values originates in the Kyoto Protocol, where they were adopted as the standard metric for comparing GHGs. Each GWP represents the integrated RF of a given GHG relative to an equal mass of CO₂ up to a defined time horizon, most commonly 20 years (GWP20) or 100 (GWP100) years. The 20- and 100-year horizons have since emerged within carbon accounting, in effect establishing an implicit compromise between projects that target near-term versus long-term temperature stabilization (Arcusa and Hagood 2024). However, using a single time horizon for carbon accounting obscures the variation in the climate impact of emissions and each mitigation strategy through time (Figure 1). At short time horizons, the RF of SLCP mitigation is significantly greater than a GWP20- or GWP100-equivalent CO₂ emission, meaning that the emission is initially over-compensated. Conversely, as the RF of SLCP mitigation dissipates at longer time horizons, the RF of the GWP-equivalent CO₂ emission is no longer fully compensated (Hausfather 2025). Similarly, while shorter-lived CDR can provide full compensation for an equivalent mass of CO₂ emissions over the time period for which storage is durable, any return of CO₂ to the atmosphere beyond this would mean that the emission is no longer fully compensated (Cullenward 2023). Using a single time horizon for carbon accounting therefore means that, in reality, emissions may not be compensated across all time horizons. This creates a structural inconsistency between emission compensation claims and the Paris Agreement's goal of long-term stabilization of global surface temperatures.

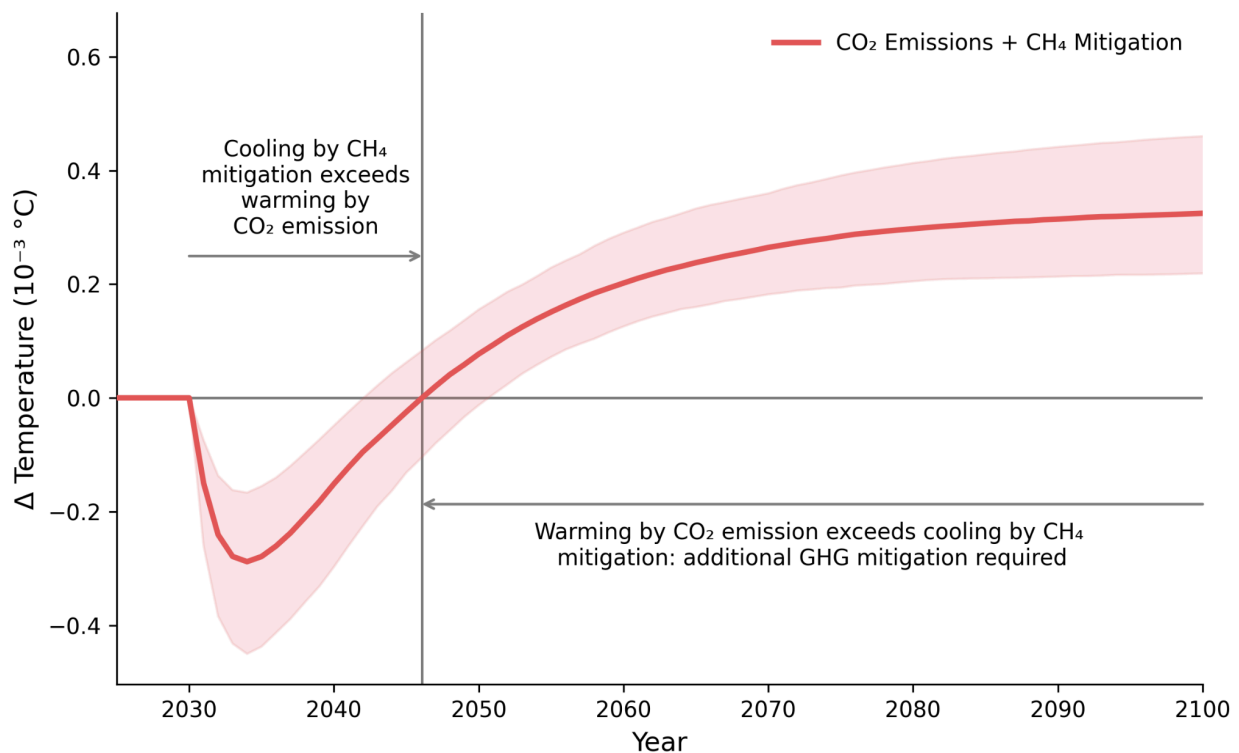


Figure 1. Modeled global mean surface temperature response to the combination of a 1 Gt CO₂ emissions pulse in 2030 and a GWP20-equivalent pulse of CH₄ mitigation in 2030 (1 Gt CO₂e, or 12.3 Mt CH₄). The thick line is the median (p50) estimate of GMST response, and shaded regions are the p05 to p95 range.

While conventional approaches based on GWP20 or GWP100 are simple and convenient, alignment between carbon accounting and global goals for long-term temperature stabilization requires an approach that explicitly represents the temperature impact of each mitigation strategy across time: SLCP mitigation to address near-term warming, shorter-lived CDR to address warming over decades to centuries, and long-lived CDR to address millennial-scale CO₂-induced warming.

To achieve this alignment, we propose a warming neutralization approach. The approach is built on two core components: a physically representative metric for climate impact, and the continuous evaluation of this impact across time. The selected metric is the global mean surface temperature (GMST) response. GMST is selected because, as a measure of the Earth system's integrated thermal response, it is coupled to climate outcomes, serves as a proxy for climate damage and reversibility, and provides a state variable that is directly aligned with global temperature goals (Lynch et al. 2020). By evaluating the combined GMST response across time, the approach allows the distinct climate outcomes of emissions and each mitigation strategy to be accounted for through a single metric. Ultimately, this enables emission compensation claims that represent the sustained neutralization of warming and are structurally consistent with global goals for long-term temperature stabilization. While GMST is the selected metric, other measures of impact including RF may also be accommodated within a sustained neutralization framework; this is discussed in Section A.2., and we invite future work to expand on this.

3. Designing GHG Mitigation Portfolios Under the Warming Neutralization Approach

Under the proposed warming neutralization approach, organizations could compensate for emissions through any combination of mitigation strategies that ensures the warming response to the emissions remains neutralized at all time horizons. Here we demonstrate how portfolios can be designed that achieve this outcome, using one possible method for implementing the approach. Multiple methods could ultimately be used, and we expand on alternatives in Section A.2. We welcome further exploration of the alternatives and their inherent trade-offs in computational efficiency, uncertainty, and operational simplicity.

3.1 A Method for Implementing the Warming Neutralization Approach

The method uses a reduced-complexity climate model (RCCM) to simulate the continuous GMST response to the combination of emissions and any existing mitigation. It thereby determines the quantity and timing of any additional mitigation required to ensure that the sustained GMST response to emissions and mitigation is $\leq 0^{\circ}\text{C}$, meaning that warming is neutralized across time. We use the Finite-amplitude Impulse Response (FaIR) RCCM for GMST simulations, but other models could be substituted without changing the underlying accounting framework. Section A.1 provides full descriptions of the method and model.

Portfolios designed using the method can combine any set of mitigation strategies, either concurrently or sequentially, that together achieve sustained warming neutralization. Because the method considers the response to emissions and mitigation across successive years, it ensures the warming response to the emissions remains neutralized without requiring year-by-year mass-based accounting. In practice, however, because a significant fraction of the warming response to a fossil CO₂ emission persists for > 1,000 years, long-term neutralization of these emissions will ultimately require an equivalent quantity of

long-lived CDR (Solomon et al. 2009; Joos et al. 2013; Archer et al. 2009). While repeat use of shorter-lived GHG mitigation could also be used for long-term neutralization, this is unlikely to be practicable on millennial timescales.

The method therefore enables two primary scenarios for portfolio design: ‘stacking’ and ‘bridging’. First, if the quantity and timing of shorter-lived GHG mitigation is known, the method can determine the schedule of replacement (or stacking) with long-lived CDR that ensures the long-term warming response to CO₂ emissions remains neutralized. Conversely, if the schedule of long-lived CDR is already defined but delayed relative to the time of an emission, the method can determine the quantity of shorter-lived mitigation required to neutralize the warming response to emissions in the interim (bridging). We outline these stacking and bridging scenarios in turn.

3.2 Stacking: Scheduling Long-Lived CDR to Replace Shorter-Lived GHG Mitigation

To demonstrate a stacking scenario, we use the simplest example: an organization uses SLCP mitigation to initially neutralize the warming response to a fossil CO₂ emission, and needs to determine the schedule for replacing this SLCP mitigation with long-lived CDR that ensures the warming remains neutralized in the long term. In this example, we arbitrarily set the emission at 1 Gt CO₂ in 2030 and the initial SLCP mitigation, also in 2030, as 1 Gt CO₂e of CH₄ mitigation (using GWP20; equates to 12.3 Mt CH₄).

To determine the schedule for replacement by long-lived CDR, the method first uses an RCCM to evaluate the continuous GMST response to the combination of the CH₄ mitigation and CO₂ emission (Figure 2). Initially, the CH₄ mitigation over-compensates for the warming response to the emission, resulting in net cooling (combined GMST response < 0°C). However, as the impact of the CH₄ mitigation dissipates, it no longer completely neutralizes the warming response (combined GMST response > 0°C). The first year that the combined GMST response crosses 0°C - using the 95th percentile estimate to account for uncertainty in the modeled GMST response (see Section A.1.4.1) - therefore defines the first year when long-lived CDR must be added to ensure the warming response to the emission remains fully neutralized.

Because the CH₄ mitigation continues to partially neutralize the warming response beyond this first year, the entire quantity of long-lived CDR does not need to be added immediately. Instead, for every year in which the combined GMST response is > 0°C, the minimum quantity of long-lived CDR required to neutralize the warming response in that year can be calculated. In this way, a schedule is built that defines the latest possible year at which a given quantity of long-lived CDR needs to be added. In the method, the schedule concludes in the year that the cumulative mass of long-lived CDR equals the mass of the CO₂ emission (see Section A.1.4.3). Any CDR procurement that meets, or exceeds, this schedule is then sufficient to achieve warming neutralization across all time horizons.

Beyond this simple example, portfolio design under a stacking scenario could include multiple types of SLCP mitigation, including CH₄ and HFC mitigation. For each SLCP there is a relationship between the magnitude of a CO₂ emission, the quantity and timing of SLCP mitigation, and the timescale over which it can fully neutralize the warming response to an emission.




Figure 2. Application of the warming neutralization approach to a stacking scenario. **Top.** Red: modeled global mean surface temperature response to the combination of a 1 Gt CO₂ emission and 1 Gt CO₂e (GWP20) of CH₄ mitigation, both in 2030. Blue: modeled global mean surface temperature response to the combination of a 1 Gt CO₂ emission in 2030, 1 Gt CO₂e (GWP20) of CH₄ mitigation in 2030, and the successive additions of long-lived CDR required to fully neutralize the warming response to the emission. The schedule of CDR additions was calculated using the 95th percentile of the model's ensemble members to account for uncertainty in the GMST response (see Section A.1.4.1). The thick lines are the median (p50) estimates of GMST response, and shaded regions are the p05 to p95 range. **Bottom.** The initial pulse of CH₄ mitigation and the cumulative schedule of CDR additions required to achieve warming neutralization using the p95 estimate. Note that the CDR schedule represents the latest time at which long-lived CDR must be added to sustain the neutralization of the warming response, and addition before this may be preferable (see Section 4).

Organizations may also use shorter-lived CDR to initially compensate for a CO₂ emission in a stacking scenario. This CDR can fully neutralize the warming response to an emission for as long as the storage, relative to a non-intervention counterfactual, remains durable (typically over decades to centuries). However, once that storage degrades, the CDR would need to be replaced with additional mitigation to sustain neutralization. In principle, the replacement time is set by the time at which CO₂ storage is no longer assured: either because a ‘reversal’ (partial or complete return of CO₂ to the atmosphere) is observed, or because active monitoring of the storage reservoir ceases and so the absence of a reversal can no longer be verified. When physical reversals do occur, CO₂ may return to the atmosphere either abruptly or progressively, with the potential for a fraction of the CO₂ to remain stored over centuries to millennia in certain reservoirs (Piñeiro-Juncal et al. 2025). In practice, the schedule for replacement of shorter-lived CDR is therefore determined by project-specific factors, including the time at which a reversal occurs or at which monitoring for reversals ceases, and the rate at which CO₂ is subsequently returned to the atmosphere.

Given these project-specific influences on CO₂ storage, credits originating from shorter-lived CDR projects are typically assigned a durability claim (also referred to as an estimated durability, a permanence period, or storage term) representing the time over which CO₂ storage is assured. For example, for nature-based CDR, durability claims often correspond to the duration of active monitoring; biogenic carbon storage reservoirs are vulnerable to reversal risks like land use change, and when monitoring ceases it is no longer possible to verify that a reversal has not occurred. Within the duration of the durability claim, projects may still be vulnerable to reversals, but in this case the risk can be managed by independent frameworks (e.g. buffer pools or monetary permanence reserves) that are beyond the scope of this paper.

Under a stacking scenario, durability claims therefore represent a convenient proxy for the time at which a given shorter-lived CDR credit needs to be replaced with additional mitigation to sustain the neutralization of warming. To implement this conservatively, the method requires the addition of an equivalent quantity of long-lived CDR in the year following the end of the durability claim. This ensures that, even if full ‘reversal’ of shorter-lived CDR occurred immediately, the warming response to the CO₂ emission would remain neutralized. Organizations may choose to partially or completely replace shorter-lived CDR on shorter timescales, and further work to standardize reservoir- or project-specific reversal rates could ultimately enable replacement schedules distributed across multiple years (Arcusa and Hagood 2024).

3.3 Bridging: Scheduling Shorter-Lived GHG Mitigation to Bridge to Delayed CDR

Previous work has suggested that the temporary impact of SLCP mitigation can be used to ‘bridge’ a delay in net CO₂ drawdown by an individual CDR project (Cabiyo et al. 2026). Within a project, delays can originate from effects including the physical processes responsible for CO₂ drawdown, the release of CO₂ in a project counterfactual, accelerated emissions resulting from project activities, or the embodiment of project emissions. These effects produce a time lag, or progressive ramp-up, in the net CO₂ drawdown relative to the initiation of the project. If used to compensate for an immediate emission, such a lag or ramp-up would leave near-term warming uncompensated and could even drive a temporary net temperature increase.

Beyond an individual CDR project, the expansion in the supply of CDR to a climate-relevant scale is currently projected to be outpaced by the rate of hard-to-abate CO₂ emissions (S. M. Smith et al. 2024). This means the availability of scaled CDR will likely be delayed relative to the emissions it must ultimately address. This could again leave the near-term warming associated with these emissions uncompensated; in this case at portfolio or planetary scale.

Whether addressing a delay in CDR at project, portfolio, or planetary scale, shorter-lived GHG mitigation can provide a temporary 'bridge' to ensure the warming response to emissions remains neutralized across the delay. To demonstrate this bridging scenario, we again choose a simple example: an organization uses long-lived CDR to neutralize the warming response to a fossil CO₂ emission, and needs to calculate the quantity of shorter-lived GHG mitigation required per year to neutralize warming across a delay in its delivery. The emission is arbitrarily set as 1 Gt CO₂ in 2030 and the initial shorter-lived mitigation as CH₄. The CDR is arbitrarily assumed to begin 10 years after the emission and then ramp up linearly over the following 10 years to a cumulative mass of 1 Gt.

To design this 'bridge', the method uses the RCCM to calculate the minimum quantity of CH₄ mitigation required in each year across the delay to fully neutralize the warming response to the emission (i.e. the minimum quantity of CH₄ mitigation such that the modeled GMST response to the combination of emission and CH₄ mitigation remains < 0°C, using the 95th percentile estimate to account for uncertainty). CH₄ mitigation continues to be required up until the year in which the cumulative quantity of long-lived CDR becomes equivalent to the CO₂ emission, after which the CDR can itself fully neutralize the warming (see Section A.1.4.3). Through the combination of the CH₄ mitigation and long-lived CDR, the warming is then neutralized across all time periods (Figure 3).

Beyond this simple example, portfolio design under a bridging scenario could use multiple types of shorter-lived mitigation. If shorter-lived CDR is used, a quantity equivalent to the CO₂ emission will be sufficient, provided that its durability claim exceeds the time required for the long-lived CDR to ramp up to an equivalent cumulative quantity. For SLCP mitigation, the quantity required to bridge a delay will depend on the specific properties of the SLCP, the length of the delay period, and the rate of the ramp up of long-lived CDR.

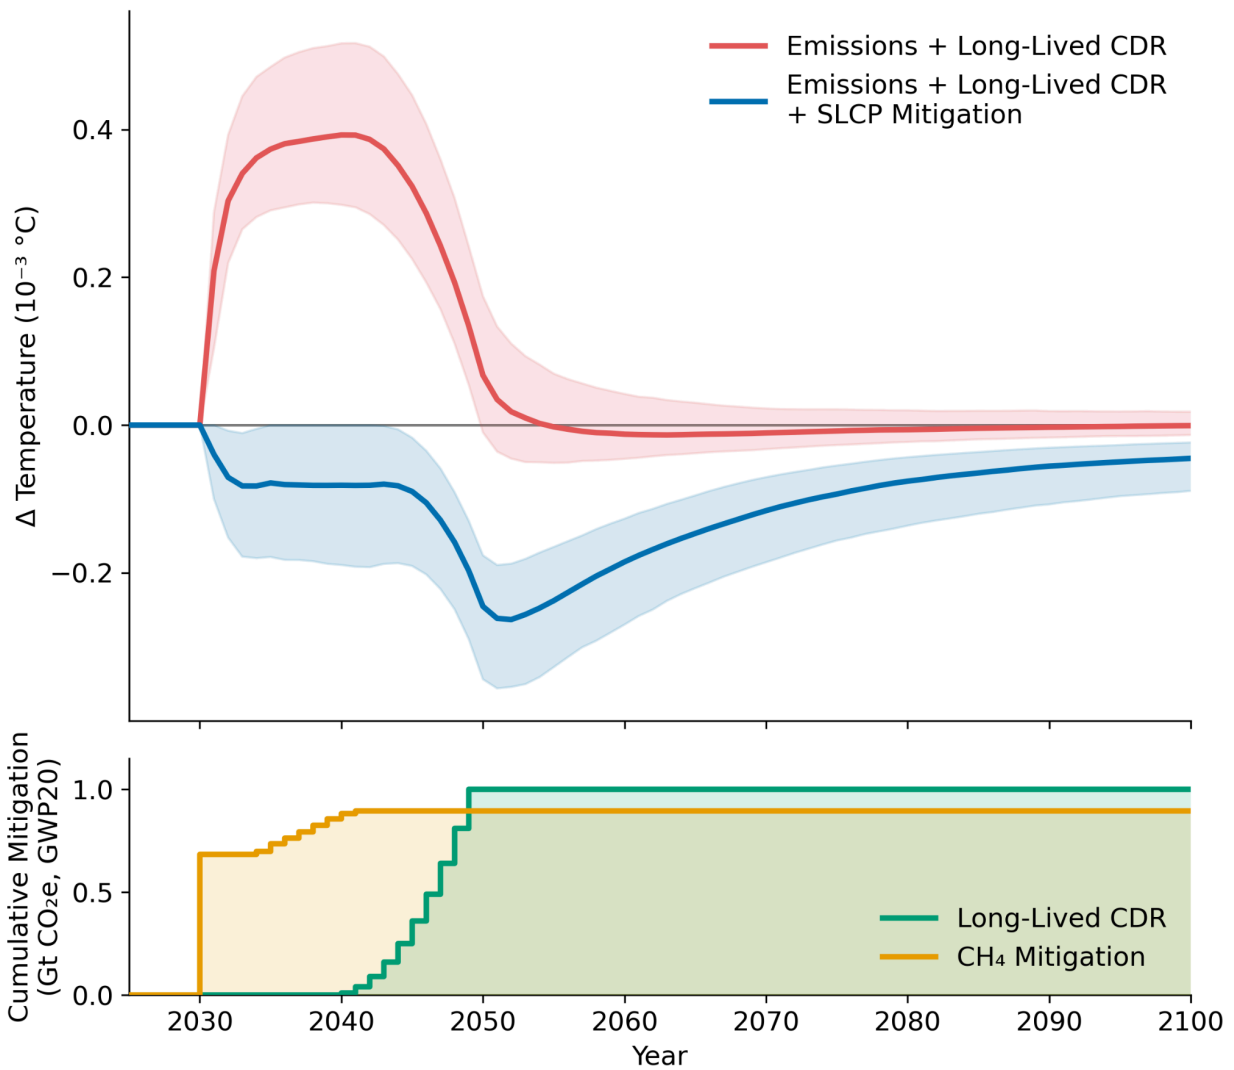


Figure 3. Application of the warming neutralization approach to a bridging scenario. **Top.** Red: modeled global mean surface temperature response to the combination of a 1 Gt CO₂ emission in 2030 and 1 Gt delayed CDR. The CDR is arbitrarily assumed to begin 10 years after the emission and then ramp up linearly over 10 years to a cumulative mass of 1 Gt. Blue: modeled global mean surface temperature response to the combination of a 1 Gt CO₂ emission in 2030, 1 Gt delayed CDR, and the successive additions of CH₄ mitigation required to fully neutralize the warming response to the emission across the CDR delay. The schedule of CH₄ mitigation was calculated using the 95th percentile of the model's ensemble members to account for uncertainty in the GMST response (see Section A.1.4.1). The thick lines are the median (p50) estimates of GMST response, and shaded regions are the p05 to p95 range. **Bottom.** The cumulative schedule of the delayed CDR and the cumulative schedule of CH₄ mitigation required to achieve warming neutralization using the p95 estimate.

3.4 Standardized Relationships for Stacking and Bridging

To avoid the need for new model execution, here we report standardized relationships that enable portfolio design under the warming neutralization approach in a stacking or bridging scenario. These relationships are derived from the method introduced in this paper (see Section A.1.3).

As previously noted, multiple types of shorter-lived GHG mitigation, including SLCP mitigation and shorter-lived CDR, could be used for stacking or bridging. For shorter-lived CDR, standardized relationships are not strictly necessary: the required quantity of shorter-lived CDR equals the quantity of CO₂ emitted, and the replacement time is independently determined by the durability claim. For SLCP mitigation, multiple SLCPs – including CH₄ or HFCs – could be used. The duration for which the SLCP mitigation can fully neutralize an emission depends on the atmospheric lifetime and warming potential of the SLCP (Table 1). For each SLCP, there is therefore a relationship between the quantity of CO₂ emitted, the quantity of SLCP mitigation used to compensate for the emission, and the time at which long-lived CDR must replace the SLCP mitigation to sustain neutralization of the emission’s warming response.

Here we report these relationships for common types of SLCP mitigation. The relationships apply to stacking or bridging under a simple scenario: CO₂ emissions are initially neutralized using a single pulse of SLCP mitigation, and this SLCP mitigation is replaced within, or at, a given time horizon with a quantity of long-lived CDR equal to the quantity of CO₂ emitted. In a stacking scenario, the relationships provide the latest time at which a given quantity of SLCP mitigation can be replaced with long-lived CDR to ensure the warming response to emissions remains fully neutralized. In a bridging scenario, the relationships provide the latest time by which CDR must have ramped up to a quantity equal to the CO₂ emitted to ensure the SLCP mitigation is sufficient to neutralize the warming response in the interim.

Table 2 reports these relationships for a case where the initial quantity of SLCP mitigation used, measured in tonnes of CO₂-equivalent (tCO₂e) at GWP20, is equal to the quantity of CO₂ emitted. GWP20 provides a convenient unit for comparing quantities of mitigation within a portfolio but, because units vary, the quantity of SLCP mitigation is also reported in metric tonnes of gas and in tCO₂e at GWP100. In either the stacking or bridging scenario, the year reported in Table 2 represents the latest possible time at which the SLCP mitigation can be replaced by long-lived CDR, and any replacement that meets or precedes this schedule will be sufficient to sustain neutralization of the emission’s warming response.

Table 2. The latest year following a CO₂ emission in which SLCP mitigation can be replaced with CDR to sustain the neutralization of the warming response to the emission. The initial quantity of SLCP mitigation, when measured in tCO₂e at GWP20, equals the quantity of CO₂ emitted. The quantity of SLCP mitigation is also shown as a mass of gas and in tCO₂e at GWP100. In each case it is shown relative to the initial CO₂ emission, which is itself equivalent to the quantity of long-lived CDR required at the time of replacement. Replacement years were calculated using the p95 estimate and rounded to the nearest year. Equivalent values can be calculated for all major SLCPs; those shown here are for illustration only. GWP values from IPCC AR6 are listed in Table 1.

SLCP mitigation	Latest year of replacement with long-lived CDR	Ton SLCP mitigation per ton CO ₂ emitted	Ton SLCP mitigation as CO ₂ e (GWP20) per ton CO ₂ emitted	Ton SLCP mitigation as CO ₂ e (GWP100) per ton CO ₂ emitted
Methane (CH ₄)	12	0.012	1.00	0.34
HFC-134a (CH ₂ FCHF ₃)	11	0.00024	1.00	0.37

Organizations could use different time horizons for replacement with long-lived CDR in a stacking or bridging scenario: longer replacement times will simply require larger initial quantities of SLCP mitigation to ensure the warming response to the emission remains neutralized prior to replacement. A table reporting the standardized relationships between the quantity of SLCP mitigation and the quantity of CO₂ emitted for replacement times between 1 and 20 years is provided in Section A.1.3. As a useful approximation, a replacement time equal to the atmospheric lifetime of an SLCP minimizes the quantity of SLCP mitigation required. For example, if a universal 20-year replacement time were chosen, then SLCPs with shorter lifetimes would require exponentially larger quantities of mitigation to sustain warming neutralization across those 20 years. This effect is normalized where the atmospheric lifetime of each SLCP aligns with the replacement time, corresponding to the most efficient use of a given type of SLCP mitigation on a mass basis.

The simplifying assumptions used to produce the standardized relationships – namely that SLCP mitigation occurs in a single pulse, and replacement with long-lived CDR occurs on, or before, a single time horizon – mean that the relationships are more conservative than strictly required in stacking or bridging scenarios. In a stacking scenario, the SLCP mitigation will continue to partially neutralize the warming response to an emission following replacement with long-lived CDR. In a bridging scenario, a progressive ramp-up of CDR prior to replacement would provide partial cooling on top of the SLCP mitigation. If organizations prefer to calculate schedules reflecting the minimum quantity of SLCP mitigation or long-lived CDR required in a given year, a method such as that introduced in this paper could be used.

4. Operationalizing the Warming Neutralization Approach

Shorter-lived GHG mitigation strategies are highly effective in addressing acute near-term warming at scale and low cost. However, long-lived CDR also requires near-term support to reach the scale ultimately required to address long-term equilibrium warming. In addition, emission compensation claims will be most defensible if organizations take financial action to replace shorter-lived GHG mitigation with long-lived CDR before, or at, the time of the claim.

At a portfolio level, procurement strategies under the warming neutralization approach will therefore be most impactful when organizations support the scaling of long-lived CDR in parallel to shorter-lived GHG mitigation. Multiple financial structures could achieve this. One possibility is that an organization establishes independent contracts for shorter-lived GHG mitigation and long-lived CDR, and these contracts are then paired to achieve sustained warming neutralization through a method such as that introduced in this paper. The feasibility of this pairing would depend on the contract duration and the expectation that a contracting organization will continue to exist for this duration. As an alternative that addresses this risk, organizations could contribute to an independently-managed third-party fund from which committed dollars appreciate and are used to procure long-lived CDR as the impact of shorter-lived GHG mitigation dissipates. Such a fund could also act as an early off-taker to support the scale-up of long-lived CDR.

This overview does not provide an exhaustive list of financial structures that could be used to operationalize the warming neutralization approach. These structures will require further development and are beyond the scope of this paper. Ultimately, the objective of any financial structure would be for

organizations to take action before the year of claim that ensures the climate impact underlying that claim is as long-lived as the emission. We invite future work to consider the specific design of such structures.

5. Summary

This paper proposes a warming neutralization approach as an alternative to the conventional GWP framework used in carbon accounting. The approach accounts for the combined climate impact of mitigation and emissions through a single metric: the continuous global mean surface temperature (GMST) response. This metric explicitly represents the value of SLCP mitigation in addressing near-term warming, shorter-lived CDR in fully compensating for CO₂ emissions over decades to centuries, and long-lived CDR in addressing millennial-scale warming. While the GWP framework remains simple and convenient, the warming neutralization approach aims to enable emission compensation claims that explicitly consider time-dependent warming impacts and are therefore structurally aligned with global goals for long-term temperature stabilization.

Multiple methods for implementing the warming neutralization approach are possible. A specific method is introduced in this paper to demonstrate how mitigation portfolios can be designed under the approach. The method uses a reduced-complexity climate model to evaluate the continuous GMST response to emissions and mitigation. In a stacking scenario, the method calculates the minimum quantity of long-lived CDR required in a given year to sustain neutralization of the warming response to emissions as the effect of shorter-lived GHG mitigation dissipates. In a bridging scenario, the method calculates the minimum quantity of shorter-lived GHG mitigation required per year to neutralize warming across a delay in the timing of long-lived CDR relative to emissions. For simpler stacking and bridging scenarios, standardized relationships derived from the method are reported so that organizations can design portfolios under the warming neutralization approach without new model execution.

Emission compensation claims made through the warming neutralization approach will be most impactful and defensible where organizations take action to support all mitigation strategies by the time of the claim. The specific financial structures by which this is operationalized will require careful consideration, and we invite future work to focus on this.

Appendix

A.1. A Method for Implementing the Warming Neutralization Approach

Here we describe the method introduced in this paper to demonstrate how portfolios can be designed under the warming neutralization approach. The method will be made accessible through a publicly-available tool, to be released following this paper. Alternative methods and models, with different tradeoffs between computational efficiency, uncertainty, and operational simplicity, could also be used to implement the warming neutralization approach (see Section A.2).

A.1.1. Model Basis

The method uses the calibrated constrained version (1.4.1) of the Finite Amplitude Impulse Response (FaIR) reduced-complexity climate model v2.2 (C. Smith et al. 2024; Leach et al. 2021). FaIR is configured as a generalised impulse response model designed to simulate the Earth system's response to anthropogenic GHG emissions with high computational efficiency. The model is run in an ensemble configuration of 841 calibrated members under an SSP1-2.6 baseline, and again as an 841 member ensemble under the various perturbations to CO₂ emissions and GHG mitigation activities considered throughout the work. These ensemble members incorporate a range of climate sensitivity and carbon cycle feedback parameters that are statistically consistent with the probabilistic assessments provided in the IPCC Sixth Assessment Report (AR6). Global mean surface temperature responses (Δ Temperature) were determined from the difference between the set of ensemble runs under the baseline SSP1-2.6 emissions pathway run and the set of ensemble runs under each perturbation scenario. Projected changes in the CH₄ temperature anomaly are modeled using the 1.4.0 values for unperturbed lifetime, lifetime temperature sensitivity, and lifetime chemical sensitivities to NO_x, VOC, HC, and N₂O.

While FaIR is used for the method introduced in this paper, the warming neutralization approach is model-agnostic and could be implemented using other RCCMs, models with different levels of complexity, or using more simple heuristics (see Section A.2).

A.1.2. Application to Stacking and Bridging Scenarios

The method first uses FaIR to construct a temporal profile representing the GMST response to the combination of GHG emissions and any existing mitigation, relative to a non-perturbed baseline run. The stacking application can evaluate CO₂ emissions and multiple types of mitigation including, but not limited to, CH₄ emissions destruction, F-gas emissions destruction, shorter-lived CDR, and long-lived CDR. The bridging application can evaluate emissions including, but not limited to, CO₂, CH₄, and N₂O, and mitigation including, but not limited to, CH₄ emissions destruction, F-gas emissions destruction, shorter-lived CDR, and long-lived CDR. For each emission and mitigation component, the forcing is assumed to occur in yearly pulses. For mitigation, the forcing year corresponds to the year that the credit for the activity was issued (rather than delivered or retired).

The GMST response to the combination of emissions and any existing mitigation is represented as a series of yearly ΔT ($^{\circ}\text{C}$) values. To ensure conservatism, the ΔT values are evaluated at p95 of the model's constrained parameter ensemble (see Section A.1.4.1). Where the ΔT is positive, the existing mitigation is not sufficient to neutralize the emissions in that year. In the stacking scenario, the method then requires that long-lived CDR is added to the portfolio. The minimum quantity that must be added in a given year is determined by the minimum pulse of CO_2 required to neutralize the ΔT (p95) in that year. In the bridging scenario, the method requires that additional SLCP mitigation is added. The minimum quantity that must be added is determined by the minimum pulse of SLCP that neutralizes the ΔT (p95) in that year.

A binary search algorithm is used to calculate these minimum GHG pulses. The model neutralizes the ΔT for each year in turn, starting from the first year of emissions. This builds a 'schedule' representing the minimum quantity of additional mitigation required in a given year. When moving forward by one year, the new target year's ΔT is re-calculated based on any additions of mitigation in the previous year. This ensures that the non-linear effects of accumulating mitigation are accounted for in the GMST response. The tolerance of the binary search is arbitrarily set as 0.01 Gt CO_2 , 0.01 Mt CH_4 , and 0.01 kt HFC-134a, but these can be updated according to the level of precision required for a given mitigation portfolio.

In the stacking application, for consistency with existing mass-based accounting approaches, the quantity of long-lived CDR additions is capped when the combined total of existing and added long-lived CDR is equivalent to the cumulative mass of CO_2 emitted (see Section A.1.4.3). When the difference between CO_2 emitted and total long-lived CDR becomes less than the tolerance, the difference is allocated to the current year's CDR requirement and the schedule is concluded. Alternatively, if the schedule concludes with unused CDR (for example where shorter-lived mitigation maintains the temperature below 0°C until the final modeled year of 2500 without requiring use of all CDR), this residual mass is retroactively assigned to the final year in which CDR was used in the portfolio.

A.1.3. Standardized Relationships for Stacking and Bridging

Table A1 provides standardized relationships, derived from the method, representing the minimum quantity of SLCP mitigation required to fully neutralize the warming response to a CO_2 emission across time periods of 1 to 20 years. In a stacking scenario, the year reported in Table A1 represents the latest possible time at which the SLCP mitigation must be replaced by long-lived CDR, and any replacement that meets or precedes this schedule will be sufficient to sustain neutralization of the emission's warming response. In a bridging scenario, the year reported in Table A1 reflects the time by which the delayed CDR must have fully 'ramped up' to a quantity equivalent to the CO_2 emission to ensure that the initial SLCP mitigation is sufficient to neutralize the warming response in the interim.

To produce these values, FaIR was used to simulate the combined GMST response to a 1 Gt CO_2 emission pulse in 2030, a 1 Gt CO_2 removal pulse at discrete delay periods ranging from 1 to 20 years, and an SLCP mitigation pulse in 2030. A binary search algorithm with a tolerance of 0.01 Mt CH_4 or 0.01 kt HFC was used to identify the minimum SLCP pulse that ensured the maximum GMST anomaly remained less than or equal to zero from the year of emission to the end of the delay period. To ensure conservatism, the algorithm evaluated the GMST response at p95 of the model's constrained parameter ensemble (Figure A1).

While the SLCP quantities are calculated assuming a 1 Gt CO₂ emission in 2030, these results have minimal scale- and time-dependence, and the p95 values reported in Table A1 remain conservative across the range of probable scenarios for the scale and timing of corporate emissions.

Table A1. The minimum quantity of SLCP mitigation required to fully neutralize the warming response to a CO₂ emission across time periods from 1 to 20 years, shown in mass of SLCP gas and in tCO₂e using GWP20 and GWP100. The quantity of SLCP mitigation is shown relative to the initial CO₂ emission, which is itself equivalent to the quantity of long-lived CDR required at the time of replacement. SLCP quantities were calculated at the p95 estimate of the GMST response. Equivalent values can be calculated for all major SLCPs; those shown here are for illustration only.

Latest year of replacement with long-lived CDR	Ton CH ₄ mitigation per ton CO ₂ emitted	Ton CH ₄ mitigation (CO ₂ e, GWP20) per ton CO ₂ emitted	Ton CH ₄ mitigation (CO ₂ e, GWP100) per ton CO ₂ emitted	Ton HFC-134a mitigation per ton CO ₂ emitted	Ton HFC-134a mitigation (CO ₂ e, GWP20) per ton CO ₂ emitted	Ton HFC-134a mitigation (CO ₂ e, GWP100) per ton CO ₂ emitted
1	0.008	0.68	0.23	0.00022	0.90	0.33
2	0.008	0.68	0.23	0.00022	0.90	0.33
3	0.008	0.68	0.23	0.00022	0.90	0.33
4	0.008	0.68	0.23	0.00022	0.90	0.33
5	0.009	0.69	0.24	0.00022	0.90	0.33
6	0.009	0.72	0.25	0.00022	0.90	0.33
7	0.009	0.76	0.26	0.00022	0.90	0.33
8	0.010	0.79	0.27	0.00022	0.92	0.34
9	0.010	0.84	0.29	0.00023	0.95	0.35
10	0.011	0.89	0.30	0.00024	0.98	0.36
11	0.012	0.94	0.32	0.00025	1.01	0.37
12	0.012	1.00	0.34	0.00025	1.06	0.39
13	0.013	1.06	0.36	0.00026	1.10	0.41
14	0.014	1.13	0.39	0.00028	1.14	0.42
15	0.015	1.21	0.41	0.00029	1.20	0.44
16	0.016	1.29	0.44	0.00030	1.25	0.46
17	0.017	1.37	0.47	0.00032	1.31	0.48
18	0.018	1.46	0.50	0.00033	1.37	0.51
19	0.019	1.55	0.53	0.00034	1.43	0.53
20	0.020	1.66	0.57	0.00036	1.49	0.55

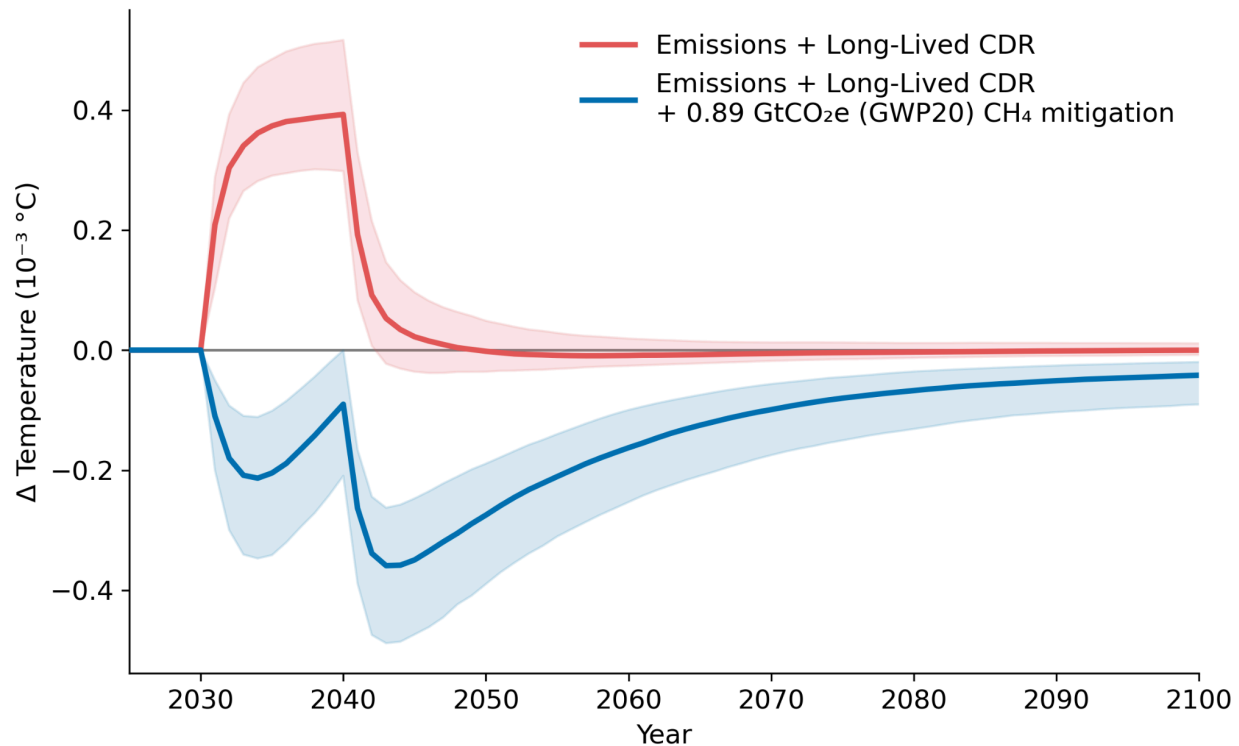


Figure A1. Red: modeled global mean surface temperature response to the combination of a 1 Gt CO₂ emission in 2030 and 1 Gt long-lived CDR after a delay period of 10 years. Blue: modeled global mean surface temperature response to the combination of a 1 Gt CO₂ emission, 0.89 Gt CO₂e (GWP20) of CH₄ mitigation in 2030 (see Table A1), and 1 Gt long-lived CDR after a delay period of 10 years. The thick lines are the median (p50) estimates of GMST response, and shaded regions are the p05 to p95 range.

A.1.4. Method and Model Assumptions

A.1.4.1. Uncertainty in Temperature Projections

Simulation of temperature responses must account for a wide range of possible future climate scenarios and varying Earth system responses. Temperature projections in FalR capture this uncertainty using ensemble configurations that incorporate a broad spectrum of climate sensitivity parameters and carbon cycle feedbacks (Leach et al. 2021; C. Smith et al. 2024).

To ensure that the method suggests additions of GHG mitigation under stacking or bridging scenarios that remain highly conservative in the context of this uncertainty, ΔT values are calculated at the 95th percentile of the 841 member ensembles. This means that the yearly GMST response is over-estimated relative to the median outcome. In turn, the method effectively over-estimates the total quantity of SLCP mitigation required to ‘bridge’ a delay, and accelerates the schedule on which long-lived CDR must be added in a stacking scenario, relative to the median outcome.

A.1.4.2. Selection of SSP Baseline

The choice of SSP baseline can influence the quantity and timing of GHG mitigation required to compensate for emissions. The SSP1-2.6 scenario was selected because it represents a trajectory consistent with deep, near-term mitigation coupled with the large-scale deployment of CDR in the latter half of the 21st century; in keeping with the approach presented in this work, which prioritises GHG

mitigation followed by the larger-scale deployment of long-lived CDR, and ultimately aligns with Paris Agreement goals (Riahi et al. 2023).

Under a higher-emissions baseline more consistent with current policy trajectories (e.g., SSP2-4.5) (Hausfather and Peters 2020), background saturation dampens the marginal radiative forcing – and by extension the GMST response – to both emissions and mitigation. Because this saturation scales non-linearly according to a specific GHG, its effect on the mitigation required to compensate for a given emission is not straightforward. For CO₂, the dampening effect of background saturation is also partially moderated by the reduced efficiency of ocean and terrestrial carbon sinks (Canadell et al. 2021). Together, these effects typically result in marginally more SLCP mitigation being required to compensate for a given mass of CO₂ emitted under a higher-emissions baseline. This is demonstrated in Table A2, which compares the standardized relationships between SLCP mitigation, CO₂ emissions, and replacement timescales under SSP1-2.6 and SSP2-4.5 scenarios.

At the scale of corporate emissions, this overall baseline effect remains small relative to the uncertainty bounds of the temperature projections themselves. At the p50 and p95 likelihoods, the standardized ratios for CH₄ mitigation in stacking or bridging scenarios under SSP1-2.6 and SSP2-4.5 baselines differ by no more than 0.001 ton CH₄ per ton CO₂ emitted (Table A2). While a less mitigated baseline scenario may therefore be more conservative in some cases, most outcomes remain accommodated by the choice to use p95 outputs under SSP1-2.6.

Table A2. The minimum quantity of CH₄ mitigation required to fully neutralize the warming response to a CO₂ emission across various time periods, calculated under SSP1-2.6 and SSP2-4.5 scenarios. The quantity of CH₄ mitigation is shown relative to the initial CO₂ emission, which is itself equivalent to the quantity of long-lived CDR required at the time of replacement. Values are calculated using the p50 and p95 estimates of the GMST response and using a 1 Gt CO₂ emission in 2030.

Latest year of replacement with long-lived CDR	Ton CH ₄ mitigation per ton CO ₂ emitted (SSP1-2.6, p50)	Ton CH ₄ mitigation per ton CO ₂ emitted (SSP1-2.6, p95)	Ton CH ₄ mitigation per ton CO ₂ emitted (SSP2-4.5, p50)	Ton CH ₄ mitigation per ton CO ₂ emitted (SSP2-4.5, p95)
5	0.007	0.009	0.007	0.009
10	0.009	0.011	0.009	0.011
15	0.012	0.015	0.012	0.015
20	0.015	0.020	0.016	0.021

A.1.4.3. Replacement by Long-Lived CDR

Under the stacking scenario, additions of long-lived CDR to a portfolio are capped when the cumulative mass of long-lived CDR equals the cumulative mass of CO₂ emitted. In reality, where the net atmospheric drawdown by a CDR deployment occurs later than the CO₂ emission, the quantity of long-lived CDR required to fully neutralize the warming response may be different to the quantity of CO₂ emitted. This is because the efficiency of natural carbon sinks is state-dependent – responding to varying atmospheric CO₂ concentrations and transient warming levels – meaning that the climate system's response to adding CO₂ and subsequently removing it is not perfectly symmetric (Zickfeld et al. 2021; Koven et al. 2023).

However, assuming a 1:1 mass equivalence between the long-lived CDR and the initial CO₂ emission is an acceptable approximation for corporate GHG accounting. The asymmetry between positive and negative emissions is heavily scale-dependent; it becomes pronounced primarily at large scales (e.g., 100s Gt of CO₂) and over long delay times (Canadell et al. 2021; Zickfeld et al. 2021). At the scale of individual corporate GHG perturbations, the forcing effect of a CO₂ emission is quasi-linear. For example, a 1 Gt CO₂ emission and delayed CDR pulse in years 2030 and 2050 respectively produces a deviation from a symmetric response that is within the p95 likelihood used to estimate quantities of GHG mitigation through the method.

Note that the avoided CO₂ resulting from the atmospheric oxidation of mitigated fossil CH₄ is not explicitly subtracted from the final mass of long-lived CDR, since this mass is typically negligible relative to the scale of the CO₂ emission being compensated (e.g., ≤ 2% of the CO₂ emission for delay periods up to 20 years; Table A1). This simplification is inherently conservative in that it marginally over-estimates the CDR required.

A.1.4.4. Model Stability and Future IPCC Assessments

The modeling of GMST response in FaIR uses physical climate parameters – including radiative efficiencies, atmospheric lifetimes, and climate-carbon feedbacks – primarily derived from the Intergovernmental Panel on Climate Change (IPCC) Assessment Reports. IPCC reports are released on a 5 - 10 year cycle, and it is therefore possible that the standardized relationships in Tables 2 and A1 could change between publication cycles. The relationships will therefore be updated on a cadence reflecting this publication cycle.

Regardless, the physical relationships governing the rapid decay of short-lived climate pollutants and the long-term persistence of CO₂ are generally well established. For example, between the IPCC's Fifth (2013) and Sixth (2021) Assessment Reports, the evaluated atmospheric perturbation lifetime of CH₄ changed from 12.4 to 11.8 years, and its base GWP100 changed from 28 to 27.9 (Myhre et al. 2014; Forster et al. 2023). Barring a major scientific reassessment of the transient climate response or the specific effective radiative forcing of major greenhouse gases following AR6, the values in Tables 2 and A1 are expected to remain stable. If there are changes in these parameters, these will most likely reflect a narrowing of the present-day uncertainty. Because the p95 output used to estimate quantities of mitigation is at the extreme upper bound of present-day uncertainty, future updates will most likely reduce the quantity of mitigation required relative to this output. Irrespective of small changes, the model standardized relationships reported in this work are therefore likely to remain conservative.

A.1.5. Computational Efficiency

Using a binary search algorithm to calculate the minimum quantity of GHG mitigation required in every year can result in long model execution times. To shorten this, it is possible to instead use temperature response functions to approximate the quantity of mitigation required in a given year (Figures A1 and A2). As previously, the GMST response to the combination of emissions and any existing GHG mitigation, relative to a non-perturbed baseline run, is first calculated as a series of yearly ΔT (°C) values at p95. In each year, the minimum mitigation pulse required to neutralize the ΔT is approximated by dividing the ΔT by a temperature response factor for mitigation. This factor is derived from a temperature response function representing the normalized GMST response (in °C mass⁻¹) to a unit-mass pulse of GHG (e.g. 1

Gt for CO₂, 1 Mt for CH₄, 1 kt for F-gases) in that year. Before moving forward by one year, the product of the additional mitigation and its full temperature response function is subtracted from future years' ΔT values, approximating the effect of the mitigation across these years.

To ensure conservatism, the temperature response functions are calculated at the 5th percentile, effectively under-estimating the GMST response to a GHG pulse in each year, and therefore over-estimating the quantity of mitigation required, relative to the median scenario. Because temperature response functions are produced for each year in isolation, they do not account for the effect of accumulating mitigation on marginal radiative efficiency. This tends to further over-estimate the quantity of mitigation required in a given year (noting that this non-linearity is small at pulse magnitudes relevant to corporate emissions and mitigation) (Etminan et al. 2016).

Because the temperature response functions place a conservative upper bound on the GHG mitigation required in a given year, they provide a useful approximation where computational efficiency is prioritised. In a stacking scenario, use of temperature response functions will effectively bring forward the time at which CDR is required relative to the full binary search approach, without modifying the total quantity of CDR required (since this is set as equal to the cumulative quantity of CO₂ emitted) (Figure A2). In a bridging scenario, temperature response functions will marginally increase the quantity of shorter-lived GHG mitigation required in a given year relative to the binary search (Figure A3).

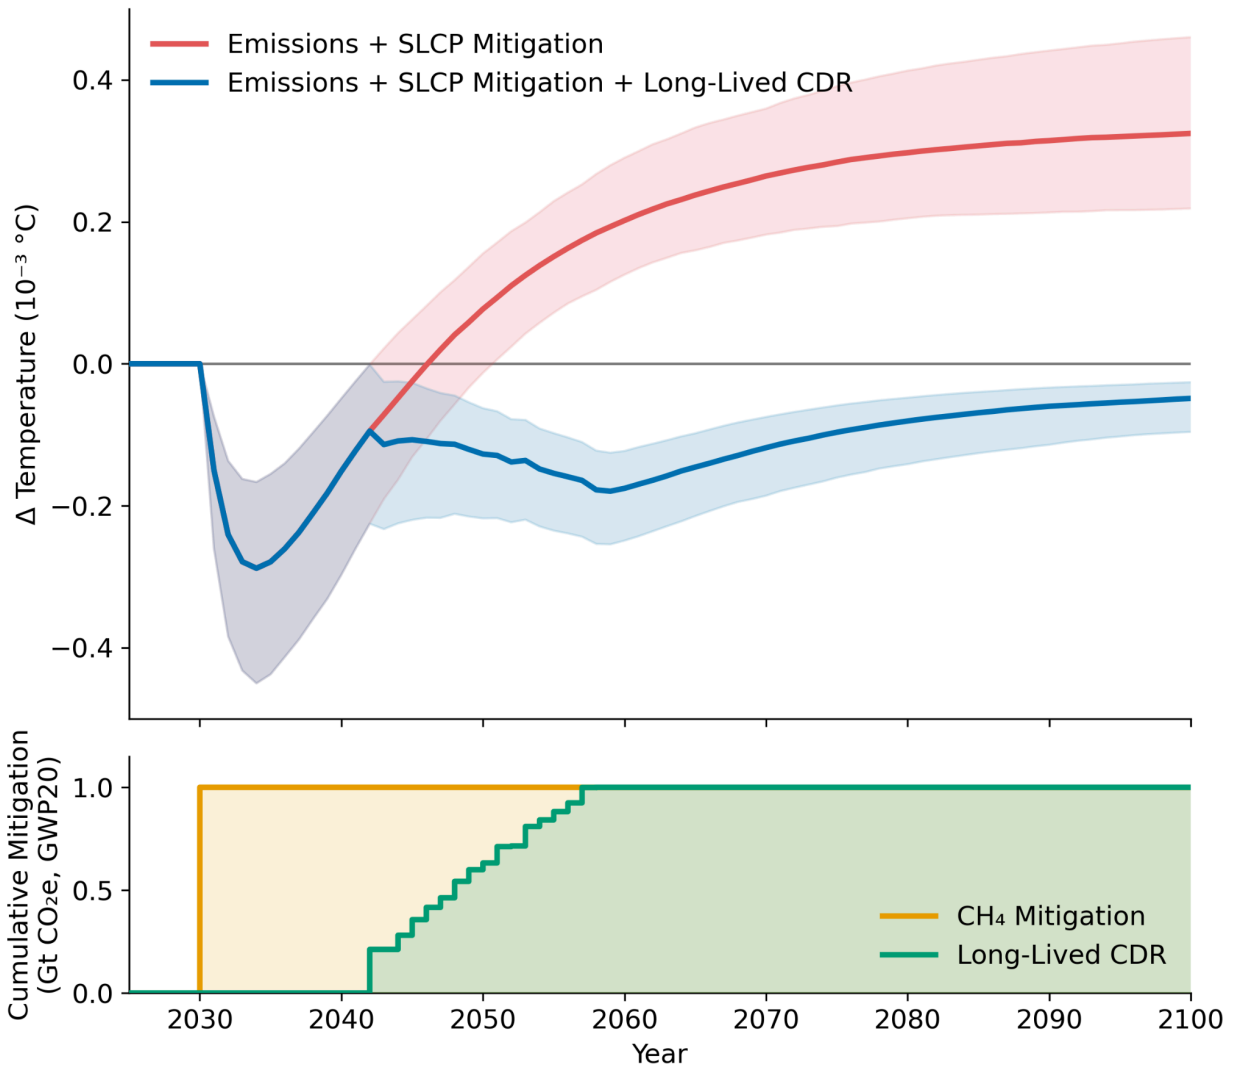


Figure A2. Application of the warming neutralization approach to a stacking scenario, using temperature response functions rather than a binary search algorithm. **Top.** Red: modeled global mean surface temperature response to the combination of a 1 Gt CO₂ emission and 1 Gt CO₂e (GWP20) of CH₄ mitigation, both in 2030. Blue: modeled global mean surface temperature response to the combination of a 1 Gt CO₂ emission in 2030, 1 Gt CO₂e (GWP20) of CH₄ mitigation in 2030, and the successive additions of long-lived CDR required to fully neutralize the warming response to the emission. To account for uncertainty, the schedule of CDR additions was calculated using the 95th percentile of the model's ensemble for the GMST response and the 5th percentile for the temperature response functions. The thick lines are the median (p50) estimates of GMST response, and shaded regions are the p05 to p95 range. **Bottom.** The initial pulse of CH₄ mitigation and the cumulative schedule of CDR additions required to achieve warming neutralization.

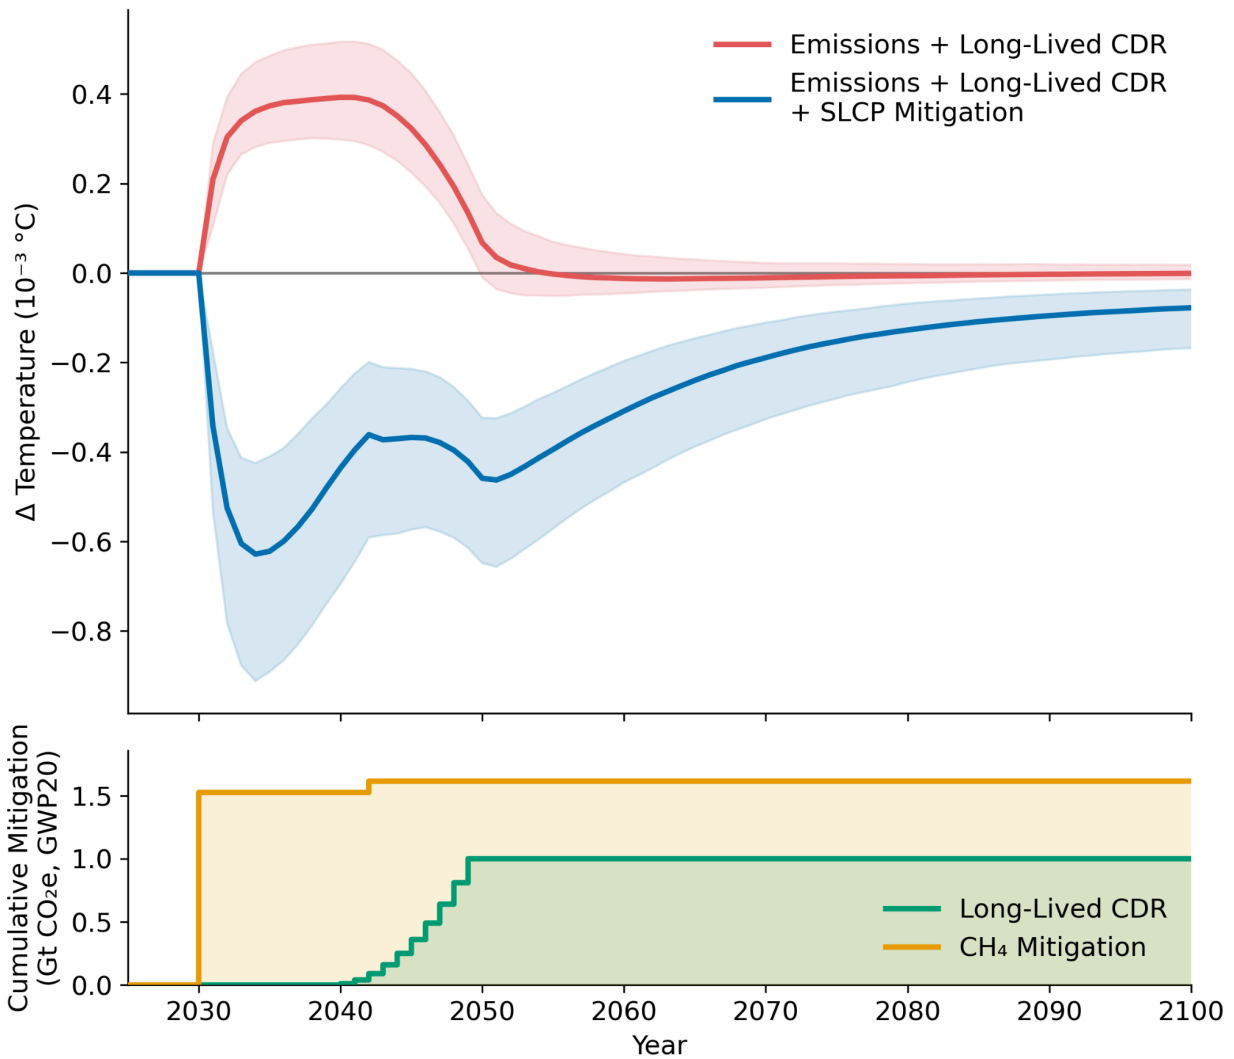


Figure A3. Application of the warming neutralization approach to a bridging scenario, using temperature response functions rather than a binary search algorithm. **Top.** Red: modeled global mean surface temperature response to the combination of a 1 Gt CO₂ emission in 2030 and 1 Gt delayed CDR. The CDR is arbitrarily assumed to begin 10 years after the emission and then ramp up linearly over 10 years to a cumulative mass of 1 Gt. Blue: modeled global mean surface temperature response to the combination of a 1 Gt CO₂ emission in 2030, 1 Gt delayed CDR, and the successive additions of CH₄ mitigation required to fully neutralize the warming response to the emission during the CDR delay. To account for uncertainty, the schedule of CH₄ mitigation was calculated using the 95th percentile of the model's ensemble for the GMST response and the 5th percentile for the temperature response functions. The thick lines are the median (p50) estimates of GMST response, and shaded regions are the p05 to p95 range. **Bottom.** The cumulative schedule of the delayed CDR and the cumulative schedule of CH₄ mitigation required to achieve warming neutralization using the p95 estimate.

A.2. Alternatives to the Warming Neutralization Approach

A.2.1. Measures and Models of Climate Impact

Multiple approaches for accounting for the time-dependent climate impact of GHG emissions and mitigation are possible. The warming neutralization approach proposed in this paper uses GMST response as the measure of impact. Here we consider the primary alternative: use of radiative forcing (RF). A neutralization approach based on RF is directly analogous to the warming neutralization approach, except that RF is used now to represent the climate impact of emissions and mitigation across time. Provided that the forcing response to the combination of emissions and mitigation remains negative across all time periods, the emissions will remain fully compensated.

As for the warming neutralization approach, a ‘RF neutralization approach’ can be implemented through a method that uses an RCM to simulate RF responses relative to a baseline across time. For RF, it is also possible to derive a simpler analytical method to describe the neutralization of emissions by mitigation across time. This is possible because a forcing response can be linked directly to GHG mass through well-constrained physical constants such as specific radiative efficiency. We outline this derivation in Section A.2.2, and then discuss the results of each accounting approach and method of implementation in Section A.2.3.

A.2.2. An Analytical Method to Implement an ‘RF Neutralization Approach’

In a neutralization approach based on RF, the aim is to ultimately ensure that the instantaneous radiative forcing (ΔRF) through a combination of GHG mitigation strategies (e.g. CDR and SLCP mitigation) is equivalent to, or exceeds, the ΔRF of a CO₂ emission across all time horizons. In a single inequality:

$$\Delta RF_{SLCP}(t) + \Delta RF_{CDR}(t) \geq \Delta RF_{CDE}(t)$$

where *SLCP* = SLCP mitigation, *CDR* = long-lived CDR, *CDE* = the CO₂ emission to be compensated, *t* is the time since the CO₂ emission pulse, and all ΔRF are absolute values. ΔRF is a function of many variables, including the mass (*M*) of the emission or mitigation. The first inequality can therefore be rewritten as:

$$M_{SLCP}(t) \cdot i(t, \dots) + M_{CDR}(t) \cdot k(t, \dots) \geq M_{CDE}(t) \cdot k(t, \dots)$$

where *i* and *k* are functions representing the relationships between *M* and *RF* for the SLCP and CO₂ respectively. The decay in atmospheric mass of an SLCP (*M(t)*) can be modeled assuming a pseudo-first order exponential decay from its initial mass *M*₀ using the IPCC AR6 perturbation lifetime (τ):

$$M(t) = M_0 \cdot e^{-t/\tau}$$

The decay in atmospheric mass of CO₂ can be modeled using the standard IPCC AR6 methodology for GWP, which represents the remaining fraction of a CO₂ pulse through an impulse response function, *IRF(t)* (Joos et al. 2013; Forster et al. 2023). Therefore:

$$M_{0,SLCP} \cdot e^{-t/\tau_{SLCP}} \cdot \frac{j(t,...)}{k(t,...)} + M_{0,CDR} \cdot IRF_{CDR}(t - t_{max}) \geq M_{0,CDE} \cdot IRF_{CDE}(t)$$

Here, t_r is the time following the CO₂ emission at which CDR replaces the SLCP mitigation. $\frac{j(t,...)}{k(t,...)}$ expresses the relative effect of a change in mass of the SLCP and CO₂ on RF, and can be approximated by the relative radiative efficiency (RE) and molecular mass of each gas (MW) through the constant α (Table A3). This approximation assumes that the relationship between M and ΔRF for each GHG is linear, which is acceptable for GHG perturbations at the scale relevant to corporate emissions.

$$M_{0,SLCP} \cdot e^{-t/\tau_{SLCP}} \cdot \alpha_{SLCP} + M_{0,CDR} \cdot IRF_{CDR}(t - t_{max}) \geq M_{0,CDE} \cdot IRF_{CDE}(t)$$

$$\text{where } \alpha_{SLCP} \approx \frac{RE_{SLCP}}{RE_{CO_2}} \cdot \frac{MW_{CO_2}}{MW_{SLCP}}$$

Table A3. Table of constants for select SLCPs. All β_{SLCP} values assume the replacement time is equivalent to τ_{SLCP} . Radiative efficiencies from IPCC AR6 (Forster et al. 2023).

SLCP	τ_{SLCP} (yr)	α_{SLCP}	β_{SLCP}
Methane (CH ₄)	11.8	117.6	0.023
HFC-134a (CH ₂ FCHF ₃)	14	5513	0.0005

To further simplify the inequality, it can be assumed that $M_{0,CDE} = M_{0,CDR}$ (since the quantity of CDR is independently set to equal to the quantity of CO₂ emitted), meaning that $M_{0,CDR} \cdot IRF_{CDR}(t - t_{max}) \geq M_{0,CDE} \cdot IRF_{CDE}(t)$. Therefore, until the time at which CDR replaces SLCP mitigation (t_r):

$$M_{0,SLCP} \cdot e^{-t/\tau_{SLCP}} \cdot \alpha_{SLCP} \geq M_{0,CDE} \cdot IRF_{CDE}(t) \text{ for } 0 \leq t < t_r$$

and so the minimum quantities of SLCP mitigation (M_0) required to neutralize the CO₂ emission across the interval between $t = 0$ and $t = t_r$ can be found at t_r :

$$M_{0,SLCP} \cdot e^{-t_r/\tau_{SLCP}} \cdot \alpha_{SLCP} = M_{0,CDE} \cdot IRF_{CDE}(t_r)$$

Finally, neglecting the fact that the forcing effect of a CO₂ emission degrades with time following the IRF (an implicitly conservative assumption):

$$M_{SLCP} / M_{CDR} = e^{t_r/\tau_{SLCP}} / \alpha_{SLCP} \tag{1}$$

Equation (1) provides a simple relationship between the volume of SLCP mitigation, volume of CDR, and time of CDR replacement that ensures that the RF response to the CO₂ emission remains neutralized across all time periods.

An even simpler single ratio representing the relative quantities of SLCP mitigation and CDR required to neutralize the forcing response to a CO₂ emission across time can be derived by setting the timescale over which CDR replacement occurs. If the timescale is assumed to equal the atmospheric lifetime of the SLCP – meaning that CDR replacement must occur within, or at, one lifetime – the simple relationship is reduced to the species-specific ratio β_{SLCP} , where $\beta_{SLCP} = e/\alpha_{SLCP}$ (Equation (2)).

$$M_{SLCP} / M_{CDR} = \beta_{SLCP} \quad (2)$$

A.2.3. Comparison of Accounting Approaches Based on Temperature and Radiative Forcing

Table A4 reports standardized ratios representing the quantity of CH₄ mitigation required to fully neutralize the climate impact of a CO₂ emission across 5- to 20-year time horizons. These ratios are calculated using three distinct methods: (a) the warming neutralization approach, implemented through the RCCM-based method introduced in this work (Section A.1); (b) an RF neutralization approach, implemented through an analogous RCCM-based method (Section A.2.1); and (c) an RF neutralization approach, implemented through the analytical method (Section A.2.2 and Equation (1)).

Table A4. The minimum quantity of CH₄ mitigation required to fully neutralize the climate impact of a CO₂ emission across various time periods, as determined using radiative forcing (RF) and global mean surface temperature (GMST) as the measures of climate impact, and using an analytical or RCCM-based method. The quantity of SLCP mitigation is shown relative to the initial CO₂ emission, which is itself equivalent to the quantity of long-lived CDR required at the time of replacement. Equivalent values can be calculated for all major SLCPs; those shown here are for illustration only. For the RCCM-based methods, values are calculated using the p95 forcing and GMST responses using a 1 Gt CO₂ emission in 2030.

Latest year of replacement with long-lived CDR	Ton CH ₄ mitigation per ton CO ₂ emitted (RCCM-based method for warming neutralization)	Ton CH ₄ mitigation per ton CO ₂ emitted (RCCM-based method for RF neutralization)	Ton CH ₄ mitigation per ton CO ₂ emitted (analytical method for RF neutralization)
5	0.009	0.009	0.013
10	0.011	0.013	0.020
15	0.015	0.019	0.030
20	0.020	0.029	0.046

The RF-based approach, whether implemented through an RCCM or analytical solution, consistently requires larger initial quantities of SLCP mitigation when compared to the GMST-based approach. This is because a pulse of CO₂ immediately exerts a positive radiative forcing, but the Earth’s thermal inertia, driven primarily by the heat capacity of the ocean’s mixed and deep layers, means that the resulting GMST increase is delayed and dampened as the oceans absorb the excess heat (Lynch et al. 2020; Forster et al. 2023). An approach that requires that the instantaneous RF from SLCP mitigation strictly equals or exceeds the RF of the CO₂ emission therefore requires a large upfront volume of SLCP mitigation to match the acute, initial energy flux of a CO₂ emission.

The analytical method for implementing the RF-based approach consistently requires a larger initial quantity of SLCP mitigation compared to the RCCM-based method for RF. This is because a number of simplifying assumptions were made in deriving the analytical method. Ultimately, the RCCM-based method is a more realistic representation of the climate system, and so the higher values produced by the analytical method should be understood as a product of simplification rather than a strict physical requirement. In addition, the mathematical precision of the analytical equations does not necessarily imply greater certainty. There remains inherent uncertainty in radiative efficiencies and atmospheric lifetimes, which the RCCM-based method captures through a probabilistic ensemble approach. A choice to implement the analytical or RCCM-based method would therefore depend on an organization's preference for simple calculations versus RCCM execution.

Because an RF neutralization approach requires larger investment in near-term SLCP mitigation to balance immediate energy fluxes, the choice of RF or GMST to represent climate impact within a neutralization approach will influence portfolio design. When comparing RF and temperature as metrics, the larger upfront volumes of SLCP mitigation required by an RF-based approach could be viewed as more conservative. On the other hand, the physical outcomes of climate change are primarily driven by temperature, rather than a strict neutralization of instantaneous energy fluxes, as highlighted by the development of temperature-aligned approaches like GWP* in recent literature (Lynch et al. 2020).

Ultimately, an approach targeting the sustained neutralization of the impact of emissions could defensibly use either RF or GMST as the accounting metric. In this work, we propose GMST to enable corporate emission neutralization claims that are structurally consistent with global goals to stabilize long-term temperatures. However, neutralization of RF remains a valid and complementary approach. We acknowledge that neither RF nor GMST explicitly represents every biogeochemical consequence of emissions, such as CO₂-induced ocean acidification. Use of a single metric to account for multiple GHGs – whether RF or GMST – inherently requires a compromise between operational utility and representation of all possible climate impacts.

A.3. Data and Code Availability

Source code for the FaIR model is available at <https://github.com/OMS-NetZero/FAIR>. The code required to reproduce the method and analysis in this paper is available upon request.

A.4. Acknowledgments

The authors thank Brad Rochlin and Hara Wang (Cascade Climate); Paulo Quadri (Technical Lead at the Symbiosis Coalition); Kristina Hughes, Calvin Tran, Nathan Truitt and Lynn Riley (American Forest Foundation); Kevin Sutherland (Isometric); Colin McCormick, John Dees and Bodie Cabiyo (Carbon Direct); and Jessica Hinojosa, Gabrielle Dreyfus, Michael Tyka, and Anna M. Michalak, for their thoughtful feedback on this paper.

References

- Allen, Myles R., Pierre Friedlingstein, Cécile A. J. Girardin, et al. 2022. "Net Zero: Science, Origins, and Implications." *Annual Review of Environment and Resources* 47 (1): 849–887.
- Archer, David, Michael Eby, Victor Brovkin, et al. 2009. "Atmospheric Lifetime of Fossil Fuel Carbon Dioxide." *Annual Review of Earth and Planetary Sciences* 37 (1): 117–134.
- Arcusa, Stephanie, and Emily Hagood. 2024. "Definitions and Mechanisms for Managing Durability and Reversals in Standards and Procurers of Carbon Dioxide Removal." *Mitigation and Adaptation Strategies for Global Change* 30. <https://doi.org/10.1007/s11027-024-10184-8>.
- Brunner, Cyril, Zeke Hausfather, and Reto Knutti. 2024. "Durability of Carbon Dioxide Removal Is Critical for Paris Climate Goals." *Communications Earth & Environment* 5 (1): 645.
- Cabiyo, Bodie, Freya Chay, Christopher Field, et al. 2026. "Consistent Temporal Accounting Supports Credible CDR Use." *CDRXIV*, January 20. <https://cdrxiv.org/preprint/302?version=1>.
- Canadell, Josep G., Pedro M. S. Monteiro, Marcos H. Costa, et al. 2021. "Global Carbon and Other Biogeochemical Cycles and Feedbacks." *IPCC AR6 WGI, Final Government Distribution*, chapter 5.
- Cullenward, Danny. 2023. "A Framework for Assessing the Climate Value of Temporary Carbon Storage." *Brussels: Carbon Market Watch*.
- Etminan, Maryam, Gunnar Myhre, Eleanor J. Highwood, and Keith P. Shine. 2016. "Radiative Forcing of Carbon Dioxide, Methane, and Nitrous Oxide: A Significant Revision of the Methane Radiative Forcing." *Geophysical Research Letters* 43 (24): 12,614–12,623.
- Fischer, E. M., S. Sippel, and R. Knutti. 2021. "Increasing Probability of Record-Shattering Climate Extremes." *Nature Climate Change* 11 (8): 689–695.
- Forster, P., T. Storelvmo, K. Armour, et al. 2023. "The Earth's Energy Budget, Climate Feedbacks and Climate Sensitivity." In *Climate Change 2021 – The Physical Science Basis: Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, edited by Masson-Delmotte, V., P. Zhai, A. Pirani, S.L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews, T.K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou. Cambridge University Press.
- Hausfather, Zeke. 2025. "Super Pollutants Are Trendy, but We Should Be Careful How We Use Them." *The Climate Brink*, November 3. <https://www.theclimatebrink.com/p/super-pollutants-are-trendy-but-we>.
- Hausfather, Zeke, and Glen P. Peters. 2020. "Emissions – the 'business as Usual' Story Is Misleading." *Nature* 577 (7792): 618–620.
- Joos, Fortunat, Raphael Roth, Jan S. Fuglestad, et al. 2013. "Carbon Dioxide and Climate Impulse Response Functions for the Computation of Greenhouse Gas Metrics: A Multi-Model Analysis." *Atmospheric Chemistry and Physics* 13 (5): 2793–2825.
- Koven, Charles D., Benjamin M. Sanderson, and Abigail L. S. Swann. 2023. "Much of Zero Emissions Commitment Occurs before Reaching Net Zero Emissions." *Environmental Research Letters* 18 (1): 014017.
- Leach, Nicholas J., Stuart Jenkins, Zebedee Nicholls, et al. 2021. "FaIRv2.0.0: A Generalized Impulse Response Model for Climate Uncertainty and Future Scenario Exploration." *Geoscientific Model Development* 14 (5): 3007–3036.
- Lynch, John, Michelle Cain, Raymond Pierrehumbert, and Myles Allen. 2020. "Demonstrating GWP*: A Means of Reporting Warming-Equivalent Emissions That Captures the Contrasting Impacts of Short- and Long-Lived Climate Pollutants." *Environmental Research Letters* 15 (4): 044023.
- MacDougall, A. H., T. L. Frölicher, C. D. Jones, et al. 2020. "Is There Warming in the Pipeline? A Multi-Model Analysis of the Zero Emissions Commitment from CO₂." *Biogeosciences* 17 (11): 2987–3016.
- Matthews, H. Damon, Kirsten Zickfeld, Mitchell Dickau, et al. 2022. "Temporary Nature-Based Carbon Removal Can Lower Peak Warming in a Well-below 2C Scenario." *Communications Earth & Environment* 3 (1): 65.
- Myhre, Gunnar, Drew Shindell, F-M Bréon, et al. 2014. "Anthropogenic and Natural Radiative Forcing." *Climate Change 2013-The Physical Science Basis*, 659–740.
- Ocko, Ilissa B., Tianyi Sun, Drew Shindell, et al. 2021. "Acting Rapidly to Deploy Readily Available Methane Mitigation Measures by

- Sector Can Immediately Slow Global Warming." *Environmental Research Letters* 16 (5): 054042.
- Pierrehumbert, Ramond T. 2014. "Short-Lived Climate Pollution." *Annual Review of Earth and Planetary Sciences* 42: 341–379.
- Piñeiro-Juncal, Nerea, Miguel Ángel Mateo, Carmen Leiva-Dueñas, et al. 2025. "Soil Organic Carbon Depth Profiles and Centennial and Millennial Decay Rates in Tidal Marsh, Mangrove and Seagrass Blue Carbon Ecosystems." *Communications Earth & Environment* 6 (1): 504.
- Riahi, Keywan, Roberto Schaeffer, Jacobo Arango, et al. 2023. "Mitigation Pathways Compatible with Long-Term Goals (Chapter 3)." *IPCC 2022: Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, 295–408.
- Smith, Chris, Donald P. Cummins, Hege-Beate Fredriksen, et al. 2024. "Fair-Calibrate v1.4.1: Calibration, Constraining, and Validation of the Fair Simple Climate Model for Reliable Future Climate Projections." *Geoscientific Model Development* 17 (23): 8569–8592.
- Smith, S. M., O. Geden, M. J. Gidden, et al. 2024. *The State of Carbon Dioxide Removal - 2nd Edition*. The State of Carbon Dioxide Removal. <https://osf.io/f85qj/>.
- Solomon, Susan, Gian-Kasper Plattner, Reto Knutti, and Pierre Friedlingstein. 2009. "Irreversible Climate Change due to Carbon Dioxide Emissions." *Proceedings of the National Academy of Sciences of the United States of America* 106 (6): 1704–1709.
- Streck, Charlotte, Sara Minoli, Stephanie Roe, et al. 2025. "Considering Durability in Carbon Dioxide Removal Strategies for Climate Change Mitigation." *Climate Policy*, 1–9.
- Szopa, Sophie, Vaishali Naik, Bhupesh Adhikary, et al. 2023. "Short-Lived Climate Forcers (Chapter 6)." *IPCC 2021: Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, 817–922.
- Wang, Seaver, Adrianna Foster, Elizabeth A. Lenz, et al. 2023. "Mechanisms and Impacts of Earth System Tipping Elements." *Reviews of Geophysics (Washington, D.C.: 1985)* 61 (1): e2021RG000757.
- Zickfeld, Kirsten, Deven Azevedo, Sabine Mathesius, and H. Damon Matthews. 2021. "Asymmetry in the Climate–carbon Cycle Response to Positive and Negative CO₂ Emissions." *Nature Climate Change* 11 (7): 613–617.