Carbon Budget Concept and its Deviation Through the Pulse Response Lens

Vito Avakumović [1]

[1]Center for Earth System Research and Sustainability (CEN), University of Hamburg, Grindelberg 5, 20144 Hamburg, Germany

Correspondence: Vito Avakumović (vito.avakumovic@uni-hamburg.de)

Abstract. The carbon budget concept states that the global mean temperature (GMT) increase is roughly linearly dependent on cumulative emissions. The proportionality is measured as the transient climate response to cumulative carbon emissions (TCRE). Since its emergence in natural science, the carbon budget concept has gained prominence as a tool for policymakers and climate economists alike. In this paper, the deviations of the carbon budget and the strict linear relationship implied by the carbon budget are examined. Hereby, two sources of deviations are distinguished: emission scenario and climate state dependence. The former stems from the scenario choice, the emission pathway, under the fixed cumulative emissions, and the latter from the change in TCRE with changing climatic conditions. Previous literature argues for scenario independence using a stylized set of emission scenarios, and offers a way to fit a non-linear carbon budget equation. This paper tests the full portofolio of emissions using an optimization procedure, showing that deviations stemming from emission pathway choices are less than 10% of the overall temperature increase and gradually diminish, even for extreme scenarios. Moreover, introducing Green's formalism using the temperature response to an emission pulse (hereafter, pulse response) as a Green's function, the scenario-dependent effects were replicated to a high degree. Hence, the behaviour of scenario-dependent deviations can be explained and predicted by the shape of the pulse response. Additionally, it is shown that the pulse response changes with climatic conditions, through which the carbon budget state dependency is explained. Using a changing pulse response as an approximation for a state-dependent TCRE, an alternative method to derive a non-linear carbon budget equation is provided. Finally, it is shown how different calibrations of a model can lead to different degrees of carbon budget non-linearities. The analysis is done using FaIRv2.0.0, a simple climate emulator model that includes climate feedback modifying the carbon cycle, along with a one-box model used for comparison purposes. The analysis shows that using the Green's function approach to diagnose a model's carbon budget scenario-dependency, along with the method of deriving the non-linear carbon budget equation, both do not depend on the complexity of the chosen climate model.

1 Introduction

The carbon budget concept, or the carbon budget approach, has gained prominence over the last decade, due to its ability to determine allowable carbon dioxide emissions leading to a specific global mean temperature (GMT) increase. In essence, it assumes a direct link between the total cumulative carbon emissions and the temperature increase without the need to know the preceding emission pathway. Following the concurrent initial discoveries in the late 2000s (Allen et al. (2009), Matthews et al. (2009), Meinshausen et al. (2009), Zickfeld et al. (2009)), the concept received wider recognition after being included in the IPCC AR5 WG1 [1] (Stocker et al., 2013), and after being presented as an explicit policy recommendation tool for limiting future climate change in IPCC AR6 WG1 (Table SPM.2) (Masson-Delmotte et al., 2021), where the 'remaining carbon budgets' indicate how much carbon may be emitted while still reaching low-temperature targets, assuming net-zero emissions afterward. By and large, since its emergence, the carbon budget has become 'a staple of climate policy discourse', having paved the way for various discourses, from policy proposals and international climate justice discussions to financial recommendations and even climate activism arguments for the immediate abandonment of fossil fuels, to name a few (Lahn, 2020).

In addition to its policy implications, the carbon budget approach plays an important part in the field of climate economics. In the analytic climate economy (ACE) models, which combine general production systems with climate dynamics in an analytically tractable way, the carbon budget approach proves to be a convenient tool that simplifies the analytical approach (Dietz and Venmans, 2019). Since ACE models have a similar structure to their numerical counterparts, integrated assessment models (IAMs), the carbon budget approach could potentially make the simple climate models used in integrated assessments redundant – if no other non-CO_2 climate change drivers are examined. In decision-making theory, Held (2019) has shown that it effectively bridges two decision-making analytic frameworks – cost-effectiveness analysis (CEA) and cost-risk analysis (CRA). The former is a dominant paradigm in IPCC ARs 5-6, WGIII (IPCC (2014), Shukla et al. (2022)). However, it is dynamically inconsistent when dealing with decision-making under uncertainty and anticipated future learning, as shown by Blau (1974). Held (2019) reasons that the carbon budget concept, i.e., it being possible to predict temperature solely on the basis of cumulative emissions, is one of the sufficient conditions to retroactively justify CEA-based scenarios as good approximations of CRA-based ones, the latter being dynamically consistent.

Formally, the carbon budget assumes the GMT increases nearly linearly with cumulative emissions, regardless of the preceding carbon emission scenario. Hence, a linear carbon budget equation:

$$T(t) = \Lambda F(t), \tag{1}$$

where $F(t) = \int_0^t E(\tau)\mathrm{d}\tau$ stands for cumulative emissions, and Λ is the proportionality constant, called the transient climate response to cumulative CO_2 emissions (TCRE). The (nearly) linear relationship emerges due to non-linearities cancelling each other out: a concave temperature dependency on the atmospheric carbon content and a convex atmospheric carbon dependency on cumulative emissions (Matthews et al. (2009), Raupach (2013)). The former stems from the radiative efficiency saturation of the atmospheric carbon, the latter from the declining ocean heat uptake and the weakening of natural carbon sinks (MacDougall and Friedlingstein, 2015).

When it comes to explicitly determining the remaining budget to reach a certain temperature target, a segmented framework had been devised by Rogelj et al. (2018). In essence, it determines what amount of cumulative emissions will lead to a given level of peak warming, if historical, non-CO_2 and Zero Emission Commitment (ZEC) warming are subtracted. ZEC is another

[1] It was not labeled explicitly as a budget but rather presented implicitly through the emphasis on temperature dependency on cumulative emissions (see Figure AR5 SPM.10).

metric closely related to TCRE and measures the warming (or cooling) that occurs after emission cessation (Matthews and Weaver, 2010). MacDougall et al. (2020) show that different models perform differently, with an inter-model range of ZEC 50 years following the emission cessation being -0.36 to 0.29 °C. If ZEC were 0, then there would be no time delay in temperature response, and emissions would directly map to temperature according to TCRE. In reality, there is always some time lag between the input and the climate system's response (e.g., Ricke and Caldeira (2014)). Regardless of ZEC, the linear segmented framework concept itself has been challenged by Nicholls et al. (2020), who claim that its assumption of a linear relationship between peak warming and cumulative emissions leads to unrealistically low budgets.

Namely, there is evidence that the relationship between the temperature and cumulative emissions (Eq. (1)) can be non-linear, as either of the two convex or concave mechanisms mentioned above could hypothetically outweigh the other under higher climatic stress (higher T). Indeed, Gillett et al. (2013) show that the linear relationship overestimates temperature response in most Earth System Models (ESMs). Using the FaIR simple climate model (SCM), Leach et al. (2021) quantify the TCRE drop to approximately 10% per 1000 GtC. Additionally, Leduc et al. (2015) have shown that constant TCRE is a good approximation for temperature response under low-emission scenarios, while it overestimates the model's response to high-intensity scenarios; this reaffirms the need for TCRE to decrease in order for the relationship in Eq. (1) to hold true. In this paper, the change in TCRE with the changing climatic conditions is referred to as (climate) *state-dependent carbon budget deviation*. Furthermore, state-dependent deviations of TCRE lead to a non-linear carbon budget equation, as TCRE is no longer a constant. In the extant literature, Nicholls et al. (2020) have derived the non-linear carbon budget equation by positing a logarithmic relationship between cumulative emissions and temperature increase.

One could imagine a second source of deviation from the budget approach that stems only from the choice of emission scenario, and not from the climate conditions of the system. In this paper, this type of deviation is referred to as an *emission scenario-dependent carbon budget deviation*. Previous literature, utilizing the high-complexity climate models (ESMs), tends to argue in favour of scenario independency (Gillett et al., 2013). However, the problem with using ESMs to study the emission scenario effects is that these models are very costly from a computational standpoint, which means only a limited set of emission pathways are examined.Millar et al. (2016) addressed this problem by forcing the simplified, globally aggregated climate model under various emission scenarios. However, to the best of the author's knowledge at the time of writing, the entire portfolio of emission scenarios that would yield the extreme cases of maximum possible scenario-dependent carbon budget deviations has yet to be investigated and scrutinized. This paper does so by using the optimization program introduced in Sect. 3.

There is evidence that state- and scenario-dependent deviations are conditional on the model's complexity (MacDougall, 2017), suggesting that models with low linearity have a higher path dependence and vice versa. In this paper, the two effects are approached as separate entities, as the emission scenario-dependent carbon budget deviation implies the possibility of achieving a different temperature T by following a different emission pathway with the same total cumulative emissions F. On the other hand, it is exactly the change in F (and consequently T) that drives the state dependency of TCRE. As will be shown in this paper, in a simple model, one can have one without the other, depending on the model's parametrization.

At its core, this paper endeavors to define and assess both the scenario- and state-dependent deviations (non-linearities) of the carbon budget approach. It demonstrates that a temperature response to an emission pulse, i.e., the pulse response representation, offers a very convenient tool for doing so.

The maximum possible and realistic carbon budget scenario-dependent deviations are identified by using an optimization program that maximizes and minimizes temperature in a specific year for fixed cumulative emissions. Through the optimization scheme, the full portfolio of emission pathways is tested.

Furthermore, a reinterpretation of the carbon budget equation is suggested using a pulse response in the context of Green's function. It is shown that the linear carbon budget equation is only a special case of the Green's function equation. More importantly, the paper demonstrates that, by utilizing the pulse response as a Green's function in the above-mentioned optimization program and comparing it to the full model results, one can capture scenario-dependent effects. Hence, it is revealed that, merely by assessing the shape of the pulse response, one can directly deduce to which extent the model adheres to carbon budget scenario independency. Ultimately, this means that using the Green's function approach allows us to calculate the maximum possible scenario dependency of the ESM models by using their pulse response and running it through the optimization program, which would be otherwise infeasible due to the computational costs.

Lastly, the changing pulse response under varying climatic conditions is translated into a state-dependent TCRE. The explicitly quantified state-dependent TCRE is then used to derive a non-linear carbon budget equation, one capable of mimicking the temperature dynamics of the full model. Therefore, it is shown that one can deduce the model's degree of carbon budget non-linearity by examining its pulse response. Further, the scenario independency of the carbon budget approach is confirmed, offering an alternative way of deriving the non-linear carbon budget equation to that put forward by Nicholls et al. (2020).

Overall, this paper is intended for both the climate modeling and climate economics audience. For the climate modeling group, it offers a fresh perspective on how to approach the carbon budget and its deviations through the pulse response lens. Additionally, it shows that the Green's function formalism holds considerable potential for predicting the scenario-dependent deviations of complex climate models. For climate economists, it reveals the consequences of using models with incorrect pulse representation, in terms of their inability to adhere to the carbon budget approach.

The paper is arranged as follows. In Sect. 2, we introduce the models and connect Green's framework with the carbon budget equation. In turn, Sect. 3 deals with scenario dependency. The optimization procedure is introduced, followed by the resulting scenario-dependent deviations. Sect. 4 deals with the pulse response representation and its consequences for the carbon budget approach. In Sect. 5, the findings are discussed in a broader context.

2 Models

The numerical optimization procedure used to generate carbon budget deviations requires a substantial number of model runs, so a computationally efficient model has to be chosen. Hence, we restrict this analysis to a class of simple climate models (SCMs), also known as climate emulators. We distinguish between and apply two approaches, the full model approach and the Green's function approach. While the former is sufficient for assessing the carbon budget deviations, the latter mathematically formalizes the carbon budget approach and offers a deeper understanding of the deviations, utilizing the pulse response representation. All of the runs are executed in the GAMS programming language, and the code for all models and runs is available online (https://doi.org/10.5281/zenodo.8314808).

2.1 FaIR model

By FaIR, I am referring to the FaIRv2.0.0 model as provided by Leach et al. (2021). The Cross-Chapter Box 7.1 in IPCC AR6 WG1 argues in favor of FaIR's value as a climate emulator (Forster et al., 2021). Furthermore, Dietz et al. (2021) recommend employing FaIR in the context of climate economics, a field that necessitates the use of SCMs for a previously elucidated reason, namely, computational efficiency. For the purposes of this paper, two features of FaIR are crucial. The first is its ability to correctly capture the temperature response following a single carbon emission pulse, i.e. pulse response (Millar et al., 2017); the second is its ability to incorporate climate feedback on the carbon cycle, with one of the effects being the modification of pulse response with changing climatic conditions.

In essence, the FaIR model is an SCM designed to emulate the gas dynamics of different radiative forcers and their effect on the global mean temperature. Because we are interested only in the deviations from the carbon budget, the non-CO_2 forcers are left out of the analysis, utilizing only the carbon cycle system and its radiative forcing dynamics. The model's description and equations can be found in Leach et al. (2021).

In brief, FaIRv2.0.0 consists of four carbon and three temperature components. Each carbon component has an associated decay timescale which dictates the dissipation of the carbon content into the shared permanent pool that represents the natural global carbon sink. Along with the global temperature increase, the sink's increased content creates a feedback mechanism, resulting in increased decay timescales and, therefore, increased atmospheric CO2 retention time. The atmospheric concentration gives rise to radiative forcing by combining a logarithmic and square root term, which translates into the temperature increase distributed between the components. Throughout Sections 3, 4 and 5, FaIR is implemented with its default parametrization, with the default thermal and carbon cycle feedback parameters provided in (Leach et al., 2021), and with the default carbon cycle parameters presented in (Millar et al., 2017).

In Section 6, uncertainty is addressed via a set of six FaIR calibrations. The parameters can be found in Tables 2 and 3 in (Leach et al., 2021), representing the thermal and carbon cycle feedback parameters tuned to CMIP6 models. Specifically, the sets used in this paper are tuned to the MIROC-ES2L, BCC-CSM2-MR, MPI-ESM1-5, CNRM-ESM2-1, and ACCESS-ESM1-5 models.

2.2 The one-box model

For comparison purposes, another SCM is introduced into the analysis, currently employed[2] as a climate module in the MIND integrated assessment model (Edenhofer et al., 2005). The one-box model consists of only one carbon and one temperature compartment, and it does not include any climate feedbacks. Moreover, Joos et al. (2013) have shown that three to four timescales attributed to individual compartments are necessary to correctly approximate the redistribution of CO_2 in the at-

[2]It is in the process of being replaced by FaIR

mosphere. Hence, the one-box model is not sufficient to fully imitate ESMs. Nevertheless, Khabbazan and Held (2019) have shown that different calibrations can be found with which it can emulate the temperature response of ESMs under RCP scenarios. The model's description and equations can be found in Petschel-Held et al. (1999). In this paper, the thermal parameters were chosen to fit the TCR and ECS values provided by FaIR's default parametrization, with the conversion formulae found in Khabbazan and Held (2019).

Note that the one-box model's pulse response should not be considered a correct representation of climate response, but rather a comparison tool. It is introduced in this article precisely because of its inexact pulse response behavior, in order to underscore how the pulse response is connected to carbon budget deviations. Also, it allows us to explore the effects of structural model uncertainty.

2.3 The Green's function framework

2.3.1 The Green's function formalism

Green's model is one equation motivated by the Green's function formalism. Essentially, a Green's function $f_\mathrm{g}(t-\tau)$ is a specific function unique to a set of linear differential equations $Lx(t) = y(t)$, where $y(t)$ is the input forcing and $x(t)$ is the state variable that changes according to the forcing and the linear operator L. The advantage of Green's function is that it acts as a 'propagator' from the input variable (external forcing) to the output variable (change in state variable), allowing us to replace differential equations with just one equation, which reads as $x(t) = \int_{t_0}^{t} y(\tau) f_\mathrm{g}(t-\tau)\mathrm{d}\tau$.

Using the same formalism, Green's equation is proposed in the context of global mean temperature dynamics with a climate model in lieu of a set of linear differential equations (see Raupach (2013)). Hence, we propose the following equation, imitating the Green's function formalism:

$$T(t) = \int_{t_0}^{t} E(\tau) f_\mathrm{g}(t-\tau)\mathrm{d}\tau. \tag{2}$$

The output variable is the global mean temperature change $T(t)$, and the input (forcing) variable is the emissions $E(t)$. Green's function $f_\mathrm{g}(t-\tau)$ modifies the contribution to a current temperature $T(t)$ stemming from the past emissions $E(\tau)$. According to Eq. (2), the temperature in time t will depend on each emission contributing at time τ prior to t, with the effect modified by Green's function f_g dependent on how far the emission year τ is from t , hence $f_\mathrm{g}(t-\tau)$. Essentially, it is an integration scheme that counts the temporarily modified temperature contributions to each emission pulse, going backwards from moment t, with a temperature being a superposition of modified contributions. Similar approaches can be found in the literature in Shine et al. (2005) and Ricke and Caldeira (2014). The difference is that, in Eq. (6), the temperature is deduced directly from emissions, without the need for quantifying the radiative forcing and/or atmospheric CO_2 response.

2.3.2 The pulse response as Green's function

To make use of Eq. (2), one must opt for a shape matching Green's function f_g. Following the proposed definition, we choose it to be a temperature evolution response following the 1 GtC emission pulse, or simply, the 'pulse response'. Therefore,

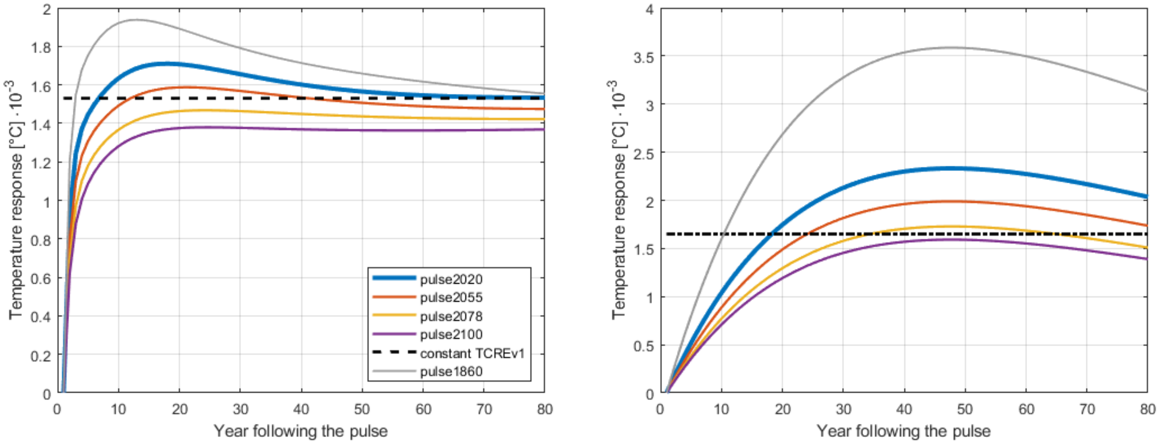

(a) Default parametrization FaIR

(b) One-box model

Figure 1. Temperature evolutions in response to 1 GtC emission pulse for different climatic conditions, i.e., pulse responses (colored lines), and the temperature response implied by Eq. (1) (black dashed line). The numbers correspond to the year of an idealized RCP6.0 scenario in which the pulses were generated. Years 2020, 2055, 2078 and 2100 correspond to the FaIR generated background temperatures of 1, 1.5, 2 and 2.5 °C, respectively, and 1860 to preindustrial climatic conditions. Constants TCREv1 and TCREv2 are equal to 1.53 and $1.6 \cdot 10^{-3}$ °C GtC^{-1} and correspond to the central TCRE estimates in Leach et al (2021) and AR6, respectively.

in this paper, the terms ‘Green’s function’, ‘pulse response’, and ‘temperature evolution following the emission pulse’ are interchangeable. Pulse response experiments are one of the generic experiments applied when evaluating climate models. Following the literature (Joos et al. (2013), Millar et al. (2017)), the pulse response is generated by adding a unit emission pulse on prescribed emissions that keep a constant background atmospheric concentration background, as follows.

The model is forced by the idealized RCP6.0 CO_2-only emission scenario provided by the RCMIP protocol (Nicholls, 2021), starting from the year 1850. In the year of pulse response generation t_p, the emission pathway necessary to keep the level of atmospheric concentration $C_a(t_p)$ constant is generated. Using the derived emissions, two experiments are run: One with the generated emissions only and one with 1 GtC extra added in t_p. Thus, the pulse response (Green’s function) is determined by subtracting the temperature evolution of the two runs.

The pulse response functions generated for different years (and hence, different climatic conditions) can be found in Figs. 1a, 1b, and 9, for the FaIR model standard parametrization, one-box model, and different FaIR parametrizations, respectively.

The Green’s functions f_g utilized in Green’s model (Eq. (2)) are generated at the year $t_p = 2020$ and depicted in blue (pulse2020) in Figs. 1a and 1b.

2.3.3 The carbon budget equation in the context of Green's formalism

Next, the connection between Green's function (Eq. (2)) and the carbon budget suggested by Eq. (1) is examined. As a first test of Green's approach, it is shown that the linear carbon budget equation is merely a special case of the former. Essentially, the linear carbon equation suggests an immediate temperature response to (cumulative) emissions, and that said response does not change in time or with climatic conditions. This implies that the pulse response introduced in the previous subsection should also be a constant function. In Fig. 1, it is plotted as a dashed black line. Formally, a linear budget pulse response can be interpreted as a Heaviside function $\Theta(t)$ multiplied by a constant equal to Λ representing TCRE:

$$f_g^0(t-\tau) = \Lambda\Theta(t-\tau) = \begin{cases} 0 & t < \tau \\ \Lambda & t \geq \tau \end{cases}, \tag{3}$$

where τ is the timing of the emission pulse and is equal to the 0th year in Fig. 1.

Proving that the Green's formalism can be considered an analogue to the carbon budget approach is simple. Inserting the idealized budget Green's function into Eq. (2), one arrives precisely at the linear budget equation (Eq. (1)):

$$T(t) = \int_{t_0}^{t} E(\tau) f_g^0(t-\tau)\mathrm{d}\tau = \int_{t_0}^{t} E(\tau)\Lambda\Theta(t-\tau)\mathrm{d}\tau = \Lambda \int_{t_0}^{t} E(t')\mathrm{d}t' = \Lambda F(t).$$

Therefore, if the temperature response always had the same (constant) shape as the dashed line in Fig. 1, the carbon budget would show deviations – each unit of carbon emission would immediately add to the warming equally and regardless of when it was emitted. However, as shown in Fig. 1, the FaIR-generated pulse responses are not a constant function, a fact that has implications for the carbon budget deviations, as will be shown.

In conclusion, Eq. (2) is to be interpreted as a generalized carbon budget equation that allows the deviations. Sects. 4 and 6 thoroughly scrutinize the connection between pulse response shape and carbon budget deviations. In the following section, the viability of Green's approach (Eq. (2)) is compared to its respective model in the context of scenario-dependent deviations.

3 Scenario-dependent deviation

3.1 Method

3.1.1 Minimization/Maximization scheme

To test the possible scenario-dependent carbon budget deviations, the optimization program is formulated as follows:

$$\underset{\{E(t)\}}{(\mathrm{Max},\mathrm{Min})}[T(t^*)] \quad \text{s.t.} \int_{t_0}^{t^*} E(t)\mathrm{d}t = F_{\mathrm{tot}}, \left|\frac{\mathrm{d}E(t)}{\mathrm{d}t}\right| \leq k,\, E(t) \geq 0,\, E(t_0) = E_0. \tag{4}$$

The program maximizes (minimizes) the temperature variable in a given optimization year t^*. The minimum $T_{\min}(t^*)$ and maximum $T_{\max}(t^*)$ temperatures generated provide the upper and lower bounds for possible temperatures under the given constraints, elaborated on in the following paragraphs. The maximum possible scenario-dependent carbon budget deviation T_{d} is then calculated by subtracting the two boundary temperatures, $T_{\mathrm{d}}(t^*) = T_{\mathrm{max}}(t^*) - T_{\mathrm{min}}(t^*)$.

In the optimization program (Eq. (4)), the emission pathway assumes the role of the free control variable, except in the fixed initial condition $E(t_0) = E_0$. Hence, the novelty of testing scenario independence with the optimization program is that the emission pathway is generated, instead of being assumed as an input by the user. This way, the analysis does not rely on a limited number of emission scenarios but systematically runs through the whole portfolio of possible scenarios under given constraints. To both avoid trivial solutions and keep the generated emission pathways within what is deemed realistic, three boundary conditions are implemented.

The first boundary condition sets the total cumulative emissions at the year of optimization t^* to a fixed value F_{tot}, counting from the initial year t_0, chosen as the year 2020 in RCP6.0. The condition on F_{tot} ensures that the deviation from the carbon budget stems only from the difference between the emission pathways, as it fixes the cumulative emissions to be equal at the end of both the minimization and the maximization run.

The second boundary condition provides the upper bound on the rate of change in emissions per year, effectively setting the allowed absolute slope of the emission pathway to be less than or equal to a prescribed value k. Hence, a trivial solution (e.g., emitting all of the emissions in one year) is avoided. The emission slope k is restricted to the upper bound of 1 $\mathrm{GtC/yr^2}$, roughly corresponding to the emission reduction rate if the annual emissions were linearly reduced to zero between the years 2020 and 2030.

The combination of the restriction on k with the F_{tot} restriction will affect the run's feasibility. The higher the cumulative emissions and the lower the k is, the less feasible the run is. Moreover, the additional requirement that the emissions reach net-zero by t^* further negatively affects the feasibility. The feasibility limiting value of k will correspond to the run where both $T_{\mathrm{max}}(t^*)$ and $T_{\mathrm{min}}(t^*)$ are equal, as they come from the only possible and feasible scenario; hence, the scenario-dependent carbon deviation $T_{\mathrm{d}}(t^*)$ is zero for that specific k. The higher k is, the more the range of possible pathway combinations increases, as does T_{d}.

The last boundary condition excludes negative emissions. This condition is utilized since Green's approach uses a pulse response generated under positive emissions. Nevertheless, for the sake of completeness, negative emissions will be allowed in the last part of the section to see how doing so affects the deviation in FaIR.

3.1.2 Two settings of the scenario-dependent deviations

To examine the carbon budget interpretation, we distinguish between two additional sets of conditions that differ depending on how much we emit after the optimization year (t^*).

The first carbon budget interpretation is addressed as a net-zero budget case, which corresponds to the situation in which all of the carbon has been emitted up until the point in time of interest, and there are no other emissions afterward. This interpretation coincides more with a carbon budget as addressed by the IPCC, which indicates how much more carbon can be emitted while still reaching specific targets. In the corresponding emission scenario set, the emissions are bound to reach zero by the year t^* and stay zero from there onwards ($E(t \geq t^*) = 0$). Note, however, that this is not the case of calculating the ZEC deviations, even though the requirement is emission cessation. ZEC tells us what the temperature evolution will be following emission cessation. In the optimization program, however, one derives two maximally different possible temperatures in a specific year, stemming from different preceding emission choices, and the deviation comes from deducting the two. ZEC affects both boundary temperature cases equally, so when the two are subtracted to get the deviation $T_d(t^*)$, the effect of ZEC is also subtracted.

On the other hand, there is the transient budget case, in which only the momentary relationship between the current cumulative emissions and current temperature increase is of interest; the emissions can evolve freely after t^*. This interpretation can be attributed to the carbon budget approach, as seen through the lens of climate economics, where the direct mapping from cumulative emissions to temperature is of interest, not the 'remaining budgets'.

The additional constraint on the emission pathway negatively affects the feasibility. Therefore, the transient budget case has more possible emission pathway combinations available compared to net zero, which means a higher expected T_{d}.

3.1.3 Deviation time evolution

The optimization procedure (Eq. (4)) calculates the extreme case of scenario-dependent deviations in one specific year t^* only. To see whether these deviations are persistent in time, an additional experiment is designed, one unique to the net-zero approach. For unit of k specified in the setup above, the system is left to evolve for the next 50 years following the optimization year ($t^* = 2070$), without adding new emissions. Hence, $T_{\mathrm{d}}(k)$ is allowed to evolve freely in time, while keeping cumulative emissions at the same level. In this way, one can see how the scenario-dependent deviation obtained in t^* changes in time (again, independent of ZEC, as explained above).

3.1.4 Run configuration

Preceding the initialization of the optimization program, the FaIR model was historically forced from the preindustrial period (the year 1850) until 2020 under the RCP6.0 emission scenario. The quasi-historical run is dynamically separated from the optimization run since, in the former, emissions are prescribed, not generated by the program. The two runs coincide in the year 2020, where the values of the historical run's variables are translated into the initial conditions of the variables of the full-fledged optimization run. Hence, $t_0 = 2020$ in Eq. (4) and the initial emissions value of the optimizer run equals to $E_0 = E_{\mathrm{RCP6.0}}(2020)$. The initial temperature at t_0 is $T_0 = 0.96$ °C, with the associated cumulative emissions counting $F_0 = 584$ GtC.

The Green's model run requires an additional modification to make it comparable with the full model. As we can see in Eq. (2), Green's approach responds only to emissions within the integral. That means that in the optimization run, which starts at t_0, it cannot capture the temperature response stemming from emissions predating t_0. Conversely, this is not a problem for the full model, since that 'leftover' temperature response is fed into the initial conditions of the run. To overcome this in Green's approach, we add the 'temperature leftover' parameter $T_{\mathrm{left}}(t)$ to Eq. (2), so it takes the form of $T(t^*) = \int_{t_0}^{t^*} E(\tau) f_{\mathrm{g}}(t^* - \tau)\mathrm{d}\tau + T_{\mathrm{left}}(t^*)$. The temperature leftover term is generated by feeding the full model with RCP6.0 emissions until the year t_p, and then setting emissions to zero at the moment of pulse response generation. $T_{\mathrm{left}}(t)$ is assessed as the temperature evolution

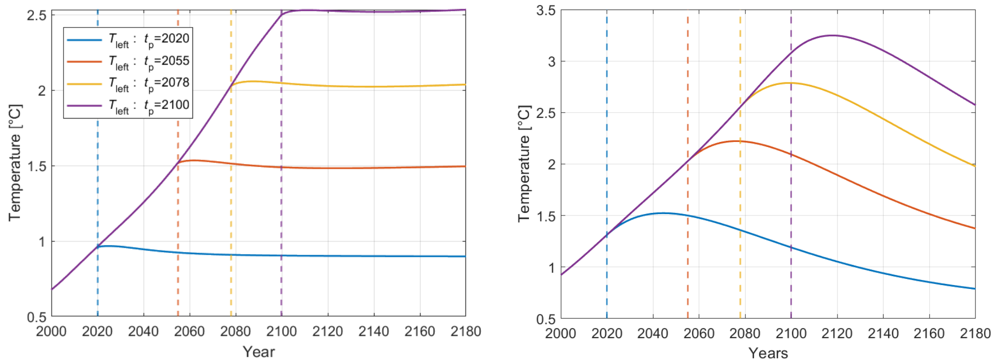

(a) Default parametrization FaIR

(b) One-box model

Figure 2. Temperature evolution run up to (RCP6.0 emission scenario) and following the emission cessation at different years t_p. The blue line represents $T_{\text{left}}(t)$, added to Green's integral to compensate for the temperature evolution leftover from prior to the optimization year $t_0 = 2020$.

after emission cessation. Various temperature leftover values corresponding to different t_p years are shown in Fig. 2. Note that the emission pathways and the years of emission cessation t_p correspond to those of pulse response generation (Fig. 1).

A final note to the reader in this regard: Unlike the net-zero case of scenario-dependent deviations (last subsection), the temperature leftover is de facto ZEC by definition – a temperature evolution following emission cessation.

3.2 Results

3.2.1 Net-zero deviation

In Fig. 3, the generated deviations for the net-zero budget case under one optimization setup are shown in light magenta. The chosen cap on cumulative emissions $F_{\text{tot}} = 416$, in addition to the pre-2020 emitted CO_2, amounts to 1000 GtC, which approximately corresponds to the carbon budget allowed for adhering to 2 °C with 67% probability, as suggested by the IPCC (Masson-Delmotte et al. (2021), Table SPM.2). The lower bound of $k = 0.4$ GtC/yr^2 is close to the feasibility limit, identifiable by the diminishing deviation. As expected, deviation increases with the k. The associated Min$[T(t^*)]$ and Max$[T(t^*)]$ from which the deviations are derived are shown in the figures in the Supplement. For demonstration purposes, the generated emission and temperature pathways for one choice of k are shown in Figs. 4c and 4d, which were generated by the FaIR and one-box model respectively.

In the case of FaIR, the magnitude of net-zero budget deviation is relatively small compared to the associated temperature increase. For the highest slope allowed ($k = 1\,\text{GtC}/\text{yr}^2$), this setup's most significant possible deviation is approximately 0.025 °C, which amounts to roughly 1.5% of the overall temperature increase (roughly 1.58 °C, see Fig 4c).

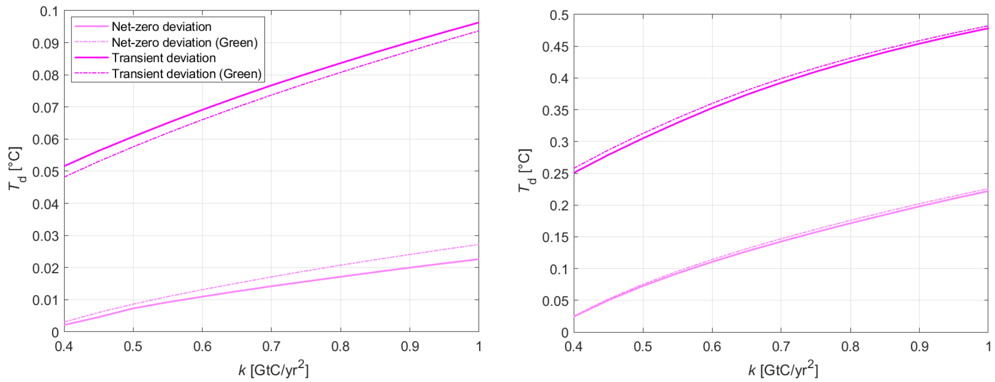

(a) Default parametrization FaIR

(b) One-box model

Figure 3. Maximum scenario-dependent deviations, dependent on the maximum emission slope allowed k, generated by the optimization program described in Sect 3.1, with $F_{\text{tot}} = 416$ GtC and $t^* = 2070$. The solid and dashed lines represent the deviations generated by the complete models (FaIR on the left, one-box on the right) and their associated Green's function models.

Unlike FaIR, the one-box model shows far more pronounced net-zero carbon budget scenario dependency. In the case of one-box model in this setup, the deviation for the highest slope k is 0.225 °C and hence amounts to roughly 10% of the overall one-box generated temperature increase (2.3 °C, see Fig 4d). As one can see, the one-box model generates about ten times larger deviations, compared to its FaIR counterpart.

As for verification of Green's approach in the context of generating scenario dependency, the deviation derived from Green's function (dashed lines, Fig. 3) is the same order of magnitude as the deviation derived from its associated full model (solid lines, Fig 3), and the trend behavior between the two is comparable. The slight shift between the model's and Green's output is due to the fact that we use a constant Green's function, although the pulse response changes. This leads us to use the pulse response and its modification under different climatic conditions (Fig. 1) to explain the source of scenario-dependent deviations in Sect. 4.

3.2.2 Scenario-dependent deviation time evolution

Figs. 4a and 4b show the time persistency of detected $T_{\text{d}}(k)$. The figures represent the time evolutions of $T_{\text{d}}(k)$ (net-zero case) following the optimization year $t^* = 2070$, with different shades of red depicting the k range as given in the abscissae in Figs. 3a and 3b. Therefore, the initial values in the year 2070 correspond to the values of $T_{\text{d}}(k)$ in Figs. 2 and 2b (light pink). Under the FaIR model, the (already small) scenario-dependent deviation ultimately disappears if no additional carbon dioxide is added to the system; hence, the maximum deviations generated by the optimization program are only temporary. In contrast to FaIR, the one-box model's deviations do not 'die out' over time but decrease only to change sign.

The deviations' evolutions for $k = 1$ can be backtracked by examining the max. (red) and min. (blue) generated temperature evolutions shown in Figs. 4c and 4d, as the subtraction of the two yields the $T_{\text{d}}(k)$. The FaIR-generated min. and max. temper-

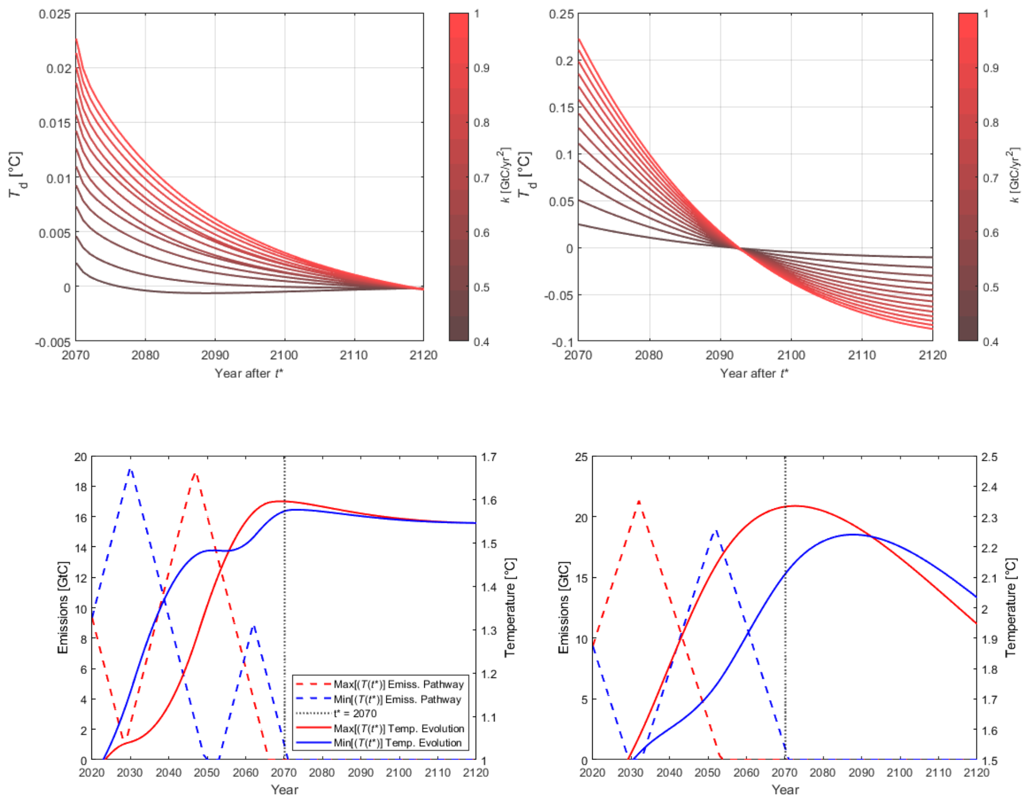

Figure 4. Graphs (a) and (b) show the temporal evolution of the net zero-case carbon budget deviation following the optimization year $t^* = 2070$, generated by FaIR and the one-box model respectively, under the setup discussed in 3.2.1. The colors represent deviations corresponding to the different k allowed, with the darkest red being the lowest allowed (0.4 GtC yr^{-2}) and the brightest red being the highest (1 GtC yr $^{-2}$). The generated emission pathways and absolute temperature evolutions corresponding to the optimization runs (both min. & max.) under the same setup for one value $k = 1$ GtC yr^{-2} are shown in graphs (c) and (d), generated by FaIR and one-box respectively.

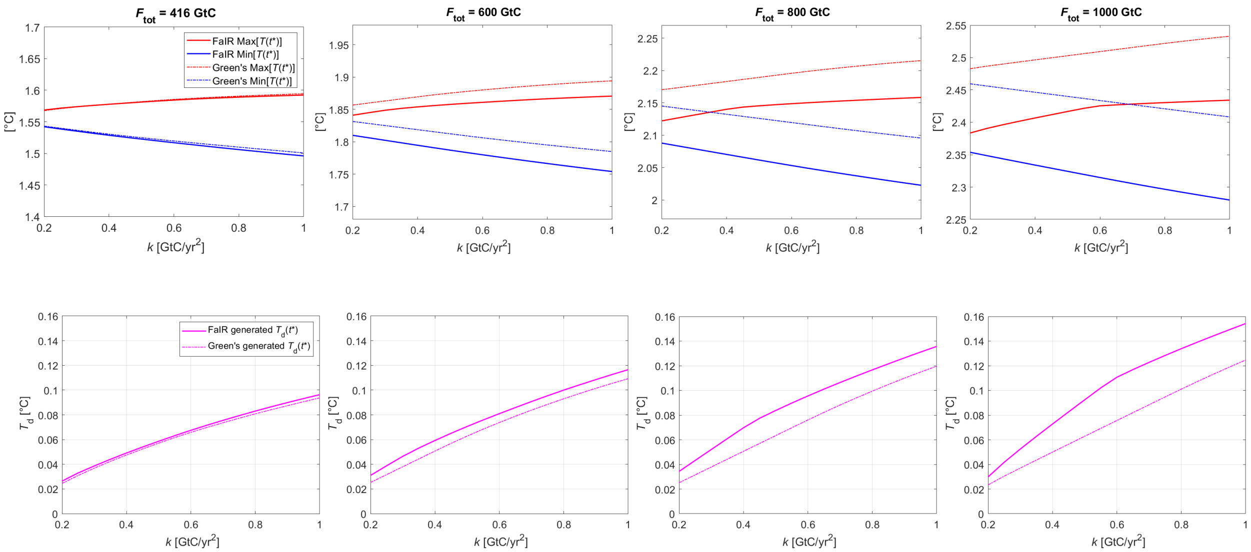

Figure 5. Top row: maximum (red) and minimum (blue) temperatures generated by the optimization program for the transient budget case, dependent on k, set up for different total cumulative emissions levels F_{tot} and $t^* = 2090$, with F_{tot} counted from the initial optimization year $t_0 = 2020$. The graphs are ordered by the magnitude of the associated F_{tot}. Y-axis domains all share the same relative interval of 0.3 °C, but different absolute values. Lower panels: corresponding scenario-dependent deviations T_d plotted against the respective k values. In all graphs, the solid lines represent the FaIR output; the dashed lines represent Green's output.

ature pathways are divided at t^* but eventually coincide, just as the carbon budget approach suggests they should. In contrast, the one-box counterpart's temperature evolutions do not reach the same pathway within the time domain of interest.

3.2.3 Transient budget deviation

In Fig. 3 the transient budget deviations are shown in dark pink lines, under the optimization run with the equivalent setup as introduced in 3.2.1. Both for FaIR- and one-box generated deviations, the transient budget case shows a significantly larger scenario-dependent carbon budget deviation (dark pink) than its net-zero counterpart (light pink). In the transient budget case, the FaIR-produced deviation is around 0.095 °C for the highest allowed k only one-fifth that of the one-box model, which produces a maximum deviation of nearly half a degree (around 0.47 °C).

The difference is due to lower minimum generated temperatures, as a result of a non-constrained $E(t^*)$ and hence allowing emissions to 'stack up' toward the optimization year. This will become clearer in Sect. 4, where we discuss the effects of pulse response shape on the deviations.

The results presented in Figs. 3 and 4 are the last based on the one-box model prior to further analysis, since its simplicity does not allow us to see the effects of the changing climate (e.g. climate feedbacks) and hence the numbers demonstrate its underperformance. They are, however, also crucial to verifying the effects of the model's pulse response (Sect. 4). In the remainder of Sect. 3, the effects of different run setup choices on the FaIR-generated deviations are assessed.

Due to the feasibility issues, we opt for a transient carbon budget approach to show how $T_{\text{d}}(k)$ changes with the increasing total cumulative emissions F_{tot}. In Fig. 5, the results of the optimizer in $t^* = 2090$ for four different F_{tot} choices are presented, explicitly showing the generated min. and max. $T(t^*)$ dependent on k in the top row and and their corresponding $T_d(t^*)$ values in the bottom row. Three main effects can be identified.

First, $T_{\text{d}}(k)$ increases with higher cumulative emissions. A comparison of the top to bottom graphs shows that the deviation increases by roughly 60%, in connection with the F_{tot} increase from 416 GtC to 1000 GtC. In the most extreme case with associated $F_{\text{tot}} = 1000$ GtC, a deviation of ∼0.15 °C is produced.

Second, the choice of the optimization year t^* does not seem to affect the deviation, if infeasibility effects are ignored. One can argue that a difference between two examples in the lowest k choices can be identified. This is where the infeasibility effect manifests: The $t^* = 2090$ case has a slightly higher limiting k close to 0.1 GtC/yr^2. In comparison, the limiting k for $t^* = 2070$ is lower – visually depicted at the intersection of the corresponding blue and red lines on the left. The two are nearly identical from roughly $k = 0.15$ GtC/yr^2 onwards. In the supplement material, various combinations of the same cumulative emissions and different t^*'s show that the deviation not being a function of the optimization year is a robust result.

Third, in the top row, the gap between FaIR and its Green counterpart's maximum and minimum temperatures steadily increases with higher cumulative emissions F_{tot}. Please note that the y-axis domains all share the same relative interval of 0.3 °C, but different absolute values. In this way, the focus is shifted to the changing difference between the Green's model-generated and FaIR-generated temperature with increasing F_{tot}.This increasing gap between the models' (Green's and FaIR) generated temperatures clearly indicates the inability of Green's model to capture non-linearities, as manifested by its use of the same, non-changing pulse response function as Green's function throughout the run. Furthermore, as shown in the bottom row, the difference in $T_{\text{d}}(k)$ between the two models also increases with higher F_{tot}, albeit to a lesser extent. This effect can be attributed to the widening gap between the maximum and minimum temperature of the FaIR approach, which increases its $T_{\text{d}}(k)$ to a larger extent than does Green's model, due to the constancy of Green's function.

As shown in the next section, the last two findings can be interpreted through the lens of the pulse response function, the magnitude and shape of which change under different climatic conditions.

3.2.4 Effect of negative emissions

To round out this section, the effects of negative emissions on the transient budget's scenaro-dependent deviation are shown in Fig 6. The figure shows four different combinations of total allowed cumulative emissions F_{tot}, this time including a choice of $F_{\text{tot}} = 196$ GtC, which, when added to the cumulative pre-optimization emissions, reflects the carbon budget allowed for adhering to 1.5°C with 67% probability, as suggested by the IPCC (Masson-Delmotte et al. (2021), Table SPM.2).

As we can see in Fig. 6 (compared to Fig. 5), including negative emissions increases the generated T_d by roughly 0.04 °C compared to the zero negative emissions scenario, in the highest k case for all F_{tot} combinations.

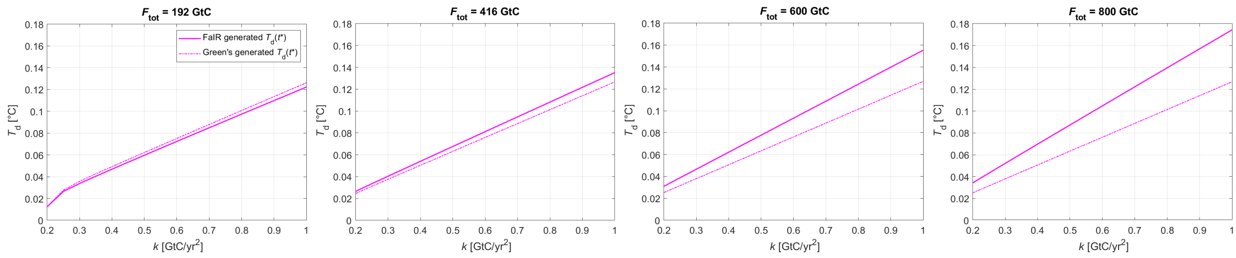

Figure 6. Scenario-dependent deviations, dependent on k, generated by the optimization program for the transient budget case with the allowed negative emissions, dependent on k, set up for different total cumulative emissions levels F_{tot} and $t^* = 2090$, with F_{tot} counted from the initial optimization year $t_0 = 2020$.

4 Pulse response as a deviation indicator

In the previous section, the extreme cases of scenario dependency were demonstrated. Firstly, it was shown that the SCM's generated scenario-dependent deviations can be emulated using a pulse response as Green's function (Eq. (2)). Hence, explaining the sources of scenario-dependent deviations by examining the pulse response's shape is justified. Secondly, the gap between the full model-generated and Green's model-generated temperatures is shown to increase with higher cumulative emissions. This is because a constant pulse response is used in Green's approach. At the same time, the pulse can change shape with changing climatic conditions.

In this section, the carbon budget approach and its deviations are contextualized through the lens of the temperature response to an emission pulse. This pulse response behavior has broader implications for other (simple) climate models and the extent to which they adhere to the carbon budget approach, with the one-box model serving as an example of inaccurate pulse representation and its consequences on deviation.

4.1 Pulse response shape as a scenario dependency indicator

To pinpoint the source of the deviations, a brief review of the discussion from Subsect. 2.3.3 is called for. As previously shown, the linear carbon budget (Eq. (1)) implies that the pulse response is a constant step function (Eq. (3), dashed black line in Fig. 1). However, the pulse response functions used in Green's model, depicted in blue and labelled pulse2020 in Fig. 1, show a dynamic temperature response.

Fig. 1, left graph shows the FaIR-generated Green's function (blue). In contrast to a constant step function, the initial response at the year of the emission pulse is zero. Then it steeply increases until reaching a maximum value of approximately $1.7 \cdot 10^{-3}$ °C GtC^{-1}, 17 years following the pulse. Furthermore, following the peak, there is a slow relaxation of the response, which slowly reaches a constant response later in time. Together, the non-instantaneous response followed by the sudden temperature increase and temperature response peak can help to understand carbon budget scenario deviations in Figs. 3a and 5 (bottom row). In contrast, the relaxation following the temperature peak explains the subsequent diminishing of deviation, as shown in Fig. 4a. The relaxation of the temperature to a certain value is further confirmed by examining Fig. 4c: Even though two temperatures are generated, they eventually reach the same level, just as the pulse relaxation suggests they should. Namely, their cumulative emissions are equal, so their pulse response is the same.

To get a better feel for the deviations and how they are connected to the pulse, one can consider an extreme example. Say that all of the emissions are injected in one year. Total cumulative emissions will then amount to the value of the emissions injection only. Due to the pulse response, the temperature response will depend on what point in time we're at. Tracing the pulse response evolution, we can see a minimum magnitude of temperature in the first year of the pulse and the maximum temperature at the peak of the response, $\sim$17 years after the pulse. Effectively, these are two very different temperatures for the same cumulative emissions. The difference between the two temperatures is the maximum possible scenario-dependent carbon budget deviation. If the cumulative emissions then amount to 100 GtC, the pulse response scales accordingly, and the theoretical deviation between the minimum and maximum response is ~0.17 °C. However, since in the optimization process, the slope restriction is fixed, and the initial emissions in 2020 are set at roughly 10 GtC, the emission pathway is not nearly as steep, resulting in smaller maximum deviations. In essence, the pulse response shows that if one wants to maximize the temperature response in a given year, they should stack the emissions $\sim$17 years before that year; conversely, to minimize the temperature response, they should stack the emissions as close as possible (dictated by k) to that year.

Finally, because of the gradual relaxation of the response, if the year in question is far enough from when we maximized the deviation, the deviation itself diminishes – as shown in Fig. 4a. In the extreme case presented in the previous paragraph, this can be intuitively seen as follows. Although there could have been a considerable difference in temperature stemming from the same cumulative emissions between the 0th (the injection year) and 17th year (the peak year) following the pulse, going forward in time, the temperature response difference between the 80th and 63rd year following the pulse (again, a 17-year difference) is virtually non-existent. Hence, the carbon budget deviation 'fixes' itself as the system enters dynamic relaxation, i.e., the pulse response reaches a nearly constant value. Once it reaches the relaxation phase, the pulse response becomes very similar to the step-function response of the linear budget.

The Green's function used in Sect. 3 and derived from the one-box model is shown in Fig. 1b (blue). Unlike its FaIR counterpart, the one-box model's pulse response never reaches the relaxation phase and peaks much later (roughly 45 years after the pulse). As discussed above, it is these two properties that ensure low-level and non-persistent scenario-dependent deviation. Indeed, Sect. 3 shows that the one-box model produces much larger and persistent carbon budget deviations, which can then be directly attributed to its pulse response shape. The same arguments as for FaIR and its associated small deviations hold true for one-box and its associated larger deviations. For example, one can easily explain why the scenario dependency changes its sign (Fig. 4b) simply by looking at the pulse response of the one-box model. If we examine Fig. 1b, there is a considerable (positive) difference between the temperature at the 0th and the 45th (peak) year. By subtracting the peak temperature from the initial year's temperature, we arrive at a positive deviation. If we go farther forward, specifically 45 years, the observer who was in the initial year now sees their temperature response at the peak, while the observer who was in the peak temperature year now sees a much lower temperature. Subtracting the two now yields a negative value, even though the deviation was previously positive.

Hence, the pulse response shape dictates both the deviation and its evolution, making it critical for the climate model's adherence to the carbon budget approach. The FaIR model shows small, scenario-dependent deviations precisely because its pulse reaches an almost constant regime relatively quickly following a peak. Moreover, if a model cannot emulate reaching the temperature relaxation, it will also show much higher emission scenario-dependent deviations.

4.2 Pulse response alteration as a state-dependency indicator

Until now, only a single pulse response (pulse2020) has been employed (Sect. 3) as Green's function and examined in the previous subsection. However, the experiment shows that this pulse response changes with changing climatic conditions: Following the same procedure described in 2.3.2, pulse responses are generated later in the RCP emission run, accordingly. The generated pulses are depicted in different colors in Fig. 1 for both the FaIR and one-box models. The further analysis considers only the FaIR results, while one-box pulse alteration will be briefly commented on later, when the mechanisms behind the pulse alterations are discussed.

When comparing the pulses (Fig. 1a), a general trend can be recognized. As the system is subjected to higher climatic stress in the form of higher cumulative emissions and higher temperatures, both the shape and the magnitude of the pulse response change. While all the pulse response variations show the aforementioned steep increase in the first few years following the pulse, the magnitude of the peak and the corresponding relaxation temperature level decrease with changing climatic conditions.

This allows us to explain the widening gap between the FaIR-generated and corresponding Green's model-generated max$T[(t^*)]$ and min$T[(t^*)]$, which widens with higher F_{tot} (Fig. 5, top row). Green's model (Eq. 2()) utilizes a non-state -dependent (non-changing) pulse response as Green's function (pulse2020). Therefore, it shows higher temperature anomalies for both maximum and minimum compared with the FaIR model, which by its nature is state-dependent in every sequential timestep (see Leach et al. (2021). The difference between the two models becomes more significant, the more stressed the climatic system is, as the pulse response used as Green's function moves farther away from the actual pulse response under the changed conditions.

In addition to the widening gap between the two models' generated min, and max. temperatures, the widening gap between the corresponding carbon budget deviations $T_{\mathrm{d}}(k)$ is identified. Fig. 5 (bottom row) shows the increasing gap between them, in favor of the FaIR model's lower acquired T_{d}. This can be explained by the flattening of the pulse response curve with higher climatic stress. Green's approach utilizes pulse2020 as a pulse response, which has a distinctly larger 'peak belly' than its counterparts.

One could have also opted for a different Green's function that would reduce the aforementioned gaps between FaIR and Green's model, but then the choice would be case-specific. Green's model could be more precise depending on whether or not we generated Green's function closer to the climatic conditions we were interested in. Hence, Green's model can indeed be seen as a linearized and simplified full model, as the theory suggests. Hypothetically, if the change in pulse response fg with climatic conditions could be incorporated into Eq. (2), the gap between the two approaches would decrease – if not disappear entirely.

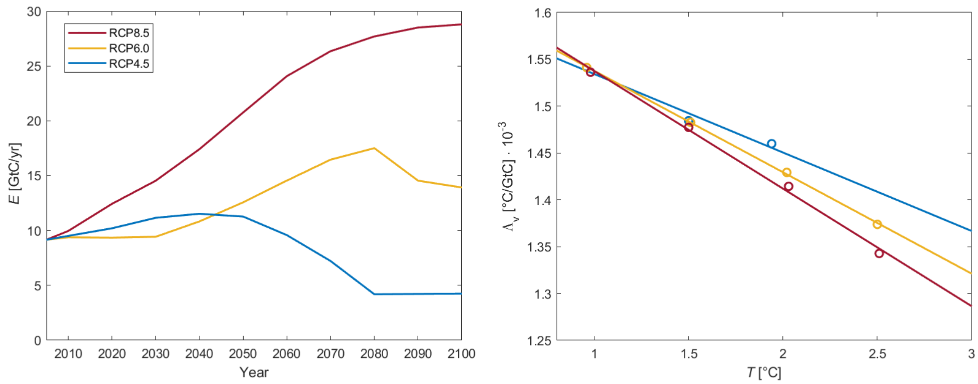

Figure 7. Right graph: TCRE approximations $\Lambda_v(T)$ generated from pulse response functions under different climatic conditions and emission scenarios. Scatter plots are actual values of Λ, while the line is the result of linear regression. The different colors represent the Λ_v dependencies generated from different RCPs, which are plotted in the left graph.

4.3 Changing pulse response as a variable TCRE: nonlinearity of the carbon budget

As discussed in the introduction, the previous literature suggests that TCRE is not a constant value but slowly decreases with cumulative emissions. This can be interpreted as the carbon budget's (=climate's) state dependence, which manifests in the non-linear carbon budget equation (Nicholls et al., 2020). This non-linearity can be identified by the change in pulse response shape with changing background climate conditions.

4.3.1 State-dependent pulse response as a variable TCRE

In Subsect. 2.3.3, it was shown how the step-function pulse response in Green's model translates into TCRE included in Eq. (1). If the TCRE changes with background conditions, the linear carbon budget step-function pulse (black dashed line, Fig. 1) should also change in magnitude following the climatic stress. Indeed, Fig. 1 shows that the SCM-generated pulse response decreases in magnitude with background conditions. If then the changing pulse is approximated with a changing step function, the decrease of the pulse response can be directly linked to the decrease of TCRE. With that approximation, however, the ability to express the time delay and scenario dependency is lost, as the shape of the pulse response function dictates the scenario dependency (Sect. 4.1). As they were shown to be small, this aspect can be safely ignored. Motivated by the findings in Sect. 4.2, in this subsection, a method for using a pulse response representation to explicitly quantify TCRE dependency on climatic conditions is developed, as follows.

To generalize the analysis, the additional pulses are generated under RCP4.5 and RCP8.5 emission scenarios, along with the already generated pulse responses under different climatic conditions with RCP6 (Fig. 1). The first pulse of each run is generated at the benchmark year 2020 and the rest at the same temperature levels (1.5, 2 & 2.5 °C), where possible.

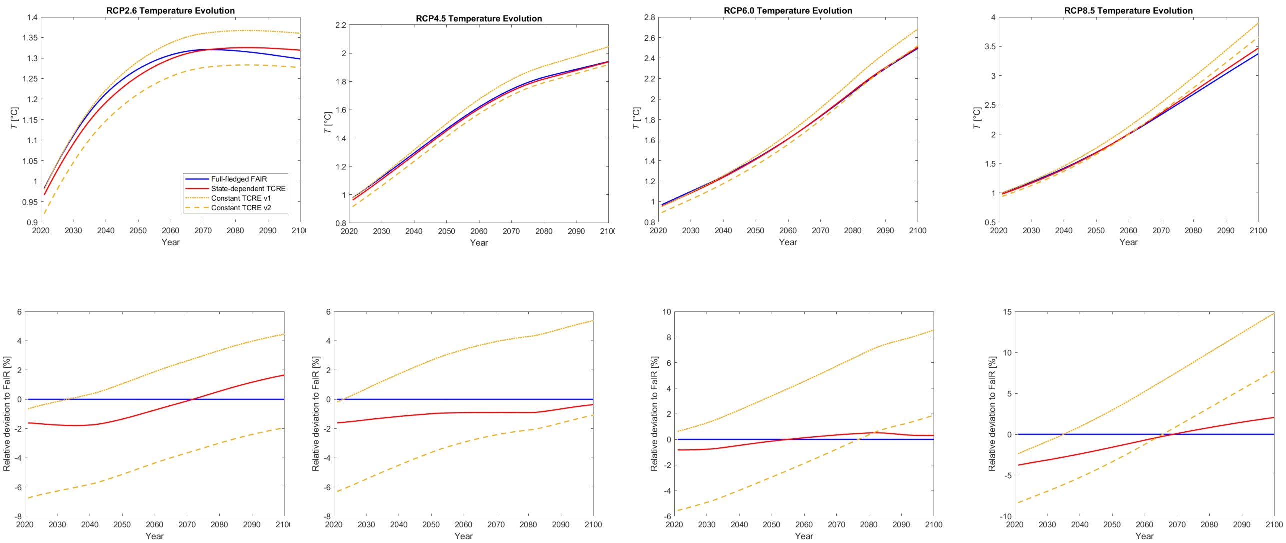

Figure 8. Top row: Temperature evolution under the three RCP emission scenarios, calculated by the full-fledged FaIR model (blue), the derived non-linear carbon budget equation (Eq. (6) (red), and the linear carbon budget equation (Eq. (1) with two different TCRE values (yellow)). Bottom row: Corresponding relative deviations of generated temperatures from FaIR-generated temperature, in percentages.

Next, recalling the perfect budget discussion, the generated pulses are to be approximated with the step function. Ignoring the temperature evolution dynamics in the early years of the pulse response, the pulse is transformed into a constant Λ_v by averaging it between years 70 and 80[3]. As shown in Fig. 1, the pulse dynamics relax by that time, reaching relative constancy. After approximating the pulses, , the corresponding cumulative emissions and temperature values (i.e., the background climatic conditions under which the original pulse was generated) are assigned to each value of generated Λ_v. By doing so, the $\Lambda_v(T, F)$ dependency is mapped, which, when reasoned in line with Eq. (1)[4], can be considered a TCRE dependent on cumulative emissions and temperature increase.

In this way, the carbon budget's state dependency is made explicit: Examining each RCP case separately shows that Λ_v decreases linearly in T and F under the standard FaIR parametrization (Fig. 7b). Moreover, looking at the right figure, one can see that by adding 1000 GtC, $\Lambda_v(F)$ drops by roughly 10%, which is in keeping with the findings of Leach et al. (2021).

4.3.2 From pulse response to carbon budget equation

The RCP6-generated Λ_v (Fig. 7b, yellow dots) is chosen to derive the carbon budget's state dependency from the pulse response representation. The choice of RCP scenario does not constrain the conclusions of this exercise. Fig. 7b suggests a linear

[3] In this way, the approximation for each pulse resembles the black dashed line in relationship to the blue line in Fig. 1.

[4] Note that Λ and Λ_v have the same function in the "perfect budget" equation. The difference is that Λ is a constant, while Λ_v is a function of temperature and cumulative emissions.

relationship $\Lambda_v(T) = -\mathrm{a} \cdot T + \mathrm{b}$, with $\mathrm{a} = 1.083 \cdot 10^{-4}$ GtC^{-1} and $\mathrm{b} = 1.646 \cdot 10^{-3}$ °C GtC^{-1} derived via linear regression. Therefore, TCRE (here Λ_v) is reinterpreted through the lens of T dependency, as temperature is a thermodynamic variable driving the system change. This way, assuming any function for the state dependency is avoided; rather, it is deducted from mapping $\Lambda_v(T)$ (Fig. 7, right graph).

Since Λ_v is, by definition, a temperature response to an emission pulse, the temperature change following the approximated pulse is interpreted as $\Delta T = \Lambda(T) \cdot E_{pulse}$. In words, the temperature change is equal to one unit of pulse emission scaled by temperature response to a pulse Λ_v. Given the fact that the emission pulse brings about a change in cumulative emissions, the aforementioned relation is rewritten in differential form as:

$$\mathrm{d}T = (-\mathrm{a} \cdot T + \mathrm{b})\mathrm{d}F. \tag{5}$$

Solving this differential equation analytically is fairly straightforward. Hence, by integrating Eq. (5), one arrives at:

$$T(F) = \frac{\mathrm{b}}{\mathrm{a}} + (T_0 - \frac{\mathrm{b}}{\mathrm{a}})\mathrm{e}^{-a(F-F_0)}, \tag{6}$$

with T_0 and F_0 being the initial values at the time of the first pulse (pulse2020). Essentially, Eq. (6) represents a non-linear carbon budget equation under a default FaIR parametrization.

When plotted, one can see that $T(F)$ is a closely linear, slightly concave function within the F domain of interest[5].

To check if Eq. (6) yields correct temperature dynamics, it is tested against the full FaIR model under the aforementioned RCP scenarios. The resulting temperature pathways are plotted in the top row of Fig. 8 (red) alongside the FaIR output (blue) and the linear carbon budget Eq. (1) with two values of constant TCRE (yellow), while the bottom row shows the corresponding relative deviations from the FaIR-generated temperature pathway. The two TCRE values are TCRE_{v1}=$1.6 \cdot 10^{-6}$ °C GtC^{-1}, and TCRE_{v1}=$1.53 \cdot 10^{-6}$ °C GtC^{-1}.

Choosing a larger constant TCRE (v1) results in a more accurate temperature diagnosis in the first half of the century under lower cumulative emissions, with deviations increasing in step with rising emissions. The opposite is true for a smaller TCRE. In this sense, Eq. (1) with a constant TCRE is a linearized version of FaIR in a similar way as the Green's function model but without the ability to generate scenario-dependent effects. Additionally, we can see that the state-dependent deviations are not transient like their scenario-dependent counterparts, but ever-increasing with the changing cumulative emissions. The highest detected absolute deviation is around ~ 0.5 °C for the end-of-the-century temperatures in the RCP8.5 run, which amounts to $\sim 15\%$ relative deviation from the FaIR-generated temperature.

Unlike constant TCRE, Eq. (6) replicates the FaIR generated temperatures in RCP2.6, RCP4.5, and RCP6 runs relatively well, with the relative deviation from FaIR being less than $\sim 2\%$ throughout the century. The largest absolute drift from the FaIR-generated temperature is around 0.1 °C at the end of the century under the RCP8.5 scenario. However, this degree of drift is less than 3% in relative terms. Since RCP8.5 is arguably somewhere in the upper bound for possible emission pathways (and RCP2.6 arguably a very optimistic lower bound scenario), one can conclude that Eq. (6) is a good emulator of FaIR

[5]Note that here F represents the total cumulative emissions from the preindustrial era. One could rewrite the equation with $\Delta F = F - F_0$ to derive the temperature increase relative to the initial year $t_0 = 2020$.

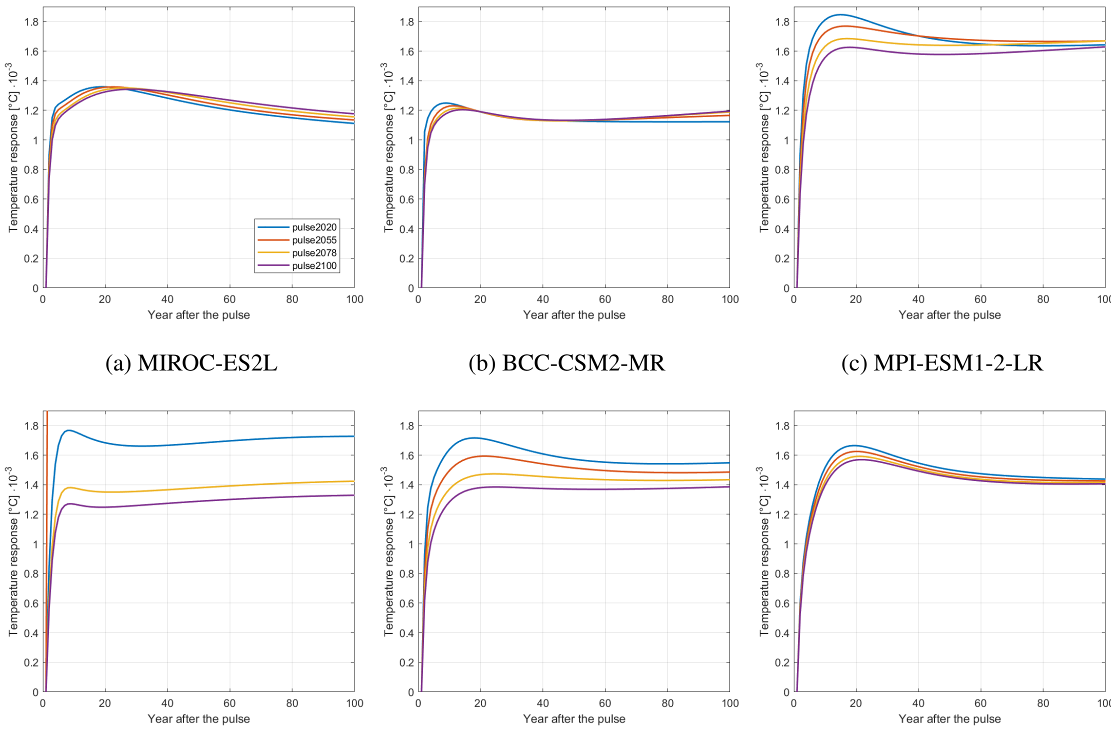

(a) MIROC-ES2L (b) BCC-CSM2-MR (c) MPI-ESM1-2-LR

(d) ACCESS-ESM1-5 (e) Default (f) CNRM-ESM2-1

Figure 9. Pulse responses under different FaIR calibrations. Different parameter sets are each tuned to a specific ESM, with parameter values given in Tables 2 and 3 in Leach et al. (2021). Note that graph (e) matches the left graph in Fig. 1, included here for comparison.

under the single, default parametrization. The incorporation of climate parameters in Eq. (6) and assessing it in view of climate economics lie beyond the scope of this paper.

4.4 Uncertainty in pulse response

By considering the pulse response representation and its implications on the carbon budget framework under one FaIR parametrization, the effects of different model calibrations on pulse response and, thereby the carbon budget are evaluated in the final part of this paper.

Fig. 9 shows pulse responses generated as described in 2.3.2 under six different sets of FaIR parameters, each tuned to a different CMIP6 model, with Fig. 9e being the default parametrization used in the rest of the analyses. We can see that every calibration yields a distinct pattern of behavior. Using the framework from the rest of Sect. 4, one can deduct how each calibration affects FaIR's adherence to the carbon budget approach

To examine scenario dependency, one must examine pulse response shape (Sect. 4.1). Looking at Fig. 9, we can see that all of the parametrizations show a relatively small scenario dependency, as all of them show pulse responses that peak in 10-20 years, followed by some degree of relaxation in the time domain of interest. In other words, one can imagine approximating them with a step function. Two parametrizations that stand out are MIROC-ES2L and ACCESS-ESM1-5. The former reaches a peak and then continually decreases, although at a much slower rate than the one-box model (Fig. 1b). Hence, the scenario-dependent deviations will not fully diminish and are likely to change sign, as per the discussion associated with 3.2.2 and Fig. 4. The same holds true for the latter, although in the other direction as the pulse response of ACCESS-ESM1-5, as the temperature gradually increases following the pulse.

In the context of state-dependent deviations, Fig. 9 reveals an interesting effect of different FaIR parametrizations on the non-linearity type of carbon budget equation. In Sect. 4.3, it was shown that the changing TCRE under different climatic conditions can be reinterpreted as the changing pulse response through $\Lambda_v(T)$. Additionally, it was shown that a decreasing $\Lambda_v(T)$ (Fig 7, right) leads to a concave non-linear carbon budget equation (Eqs. (5) and (6)). The opposite also holds true: If $\Lambda_v(T)$, and hence the pulse response increases in magnitude with higher temperatures, it results in a convex non-linear carbon budget equation. Ultimately, if the pulse response magnitude does not change with changing background conditions, the carbon budget equation is indeed linear[6]. With that in mind, one can easily deduct that not all the combinations of FaIR parameters lead to the concave carbon budget equation, as derived in Eq. (6). For example, MIROC-ES2L tuned to FaIR indicates a slightly convex budget equation, while BCC-CSM2-MR and CNRM-ESM2-1 are closest to the linear carbon budget.

This will be explored further in future work. Due to the constrained set of fully accessible different parameter sets, only six calibrations are presented here. A larger set would provide some insights into which elements of FaIR drive which types of behavior. Additionally, it would be interesting to see to which extent FaIR tuned to a CMIP6 model reproduces the behavior of its corresponding ESM under the same setup. If it were found to do so, one could extend the pulse response framework with FaIR tuned to ESMs to analyze carbon budget deviations as given by the ESM.

5 Discussion

The work shown here utilizes FaIR, the one-box model, and the associated Green's function models. The non-linearities appear in FaIR in both the carbon cycle feedback and in the temperature response saturation. As pointed out in the introduction, the interplay between the changing carbon cycle and temperature response produces the near-linearity of the carbon budget equation, with the former being a convex and the latter a concave driver of the budget equation. equation.

The second model used in the analysis is the one-box model, introduced as an example of a model with a dramatically different pulse response than FaIR, which facilitates comparison in the context of the pulse response behaviour effect on the carbon budget approach deviations. In contrast to FaIR, the one-box model does not include climate feedbacks on the carbon cycle, so non-linearities arise only through the saturation in temperature response.

The inclusion (or lack) of climate feedbacks has an effect on how the pulse response changes with changing climatic conditions. In the one-box model, the carbon cycle response stays the same regardless of background conditions, so the pulse response is modified only by logarithmic temperature response saturation. This manifests in the pulse changing magnitude

[6]Note that a pulse relaxation is still a necessary requirement.

but not shape (Fig. 1b). Conversely, including climate feedbacks changes the shape of the response function and modifies its magnitude. For a more detailed discussion on how the climate feedback changes the carbon cycle in FaIR in the context of decreased atmospheric CO_2 decay, see Millar et al. (2017).

To test whether pulse response behavior offers a trustworthy framework for explaining carbon budget deviations, it is employed as a Green's function in Eq. (2). The methodology is explained in detail in the introduction. However, by proposing Eq. (2) and using a FaIR-generated (or one-box-generated) Green's function, we assume that the climate model is a set of linear differential equations. Hence, although Green's model has been proven to capture scenario-dependent effects, the effects of climate change on the carbon budget approach cannot be explicitly captured with Eq. (2). This effect is visible when comparing the full model and Green's model optimization runs, as the two sets of generated temperatures have an ever-increasing gap with higher cumulative emissions (Fig. 5, top row). One could modify Eq. (2) so as to include a changing pulse response instead of a fixed f_g, but this remains theoretical; the implementation is unclear.

Regardless of Green's model's inability to correctly forecast (or hindcast, for the same reasons) temperature evolution, Sect. 3 shows that it is indeed capable of mimicking the scenario-dependent deviations of both FaIR and the one-box model. Even though there is an ever-increasing gap between temperatures generated by the full model and Green's model, the scenario-dependent deviations are well represented by Green's function even for higher F_{tot}. Hence, one concludes that state and scenario dependencies can arise independently.

In essence, the results show that the changing of the pulse under changing background conditions does not affect Green's model's ability to predict scenario dependency. This implies that one could use any model, of any complexity, generate its pulse response and then plug it into Green's model under the optimization program (Eq. (4)) to arrive at the corresponding model's possible scenario-dependent deviations. In the case of a complex climate model (e.g. an ESM), this would be possible only through the Green's function approach due to the unacceptable computational time [7] required to run an ESM in an optimizing program.

When it comes to purely numerical findings in the context of scenario-dependent deviations, it was shown that how much we emit after the optimization year can dramatically affect the generated deviations. For FaIR, the largest possible deviation we acquire is approximately 0.15 °C for the transient budget case. In the net-zero case, the largest deviation is well below 0.1 °C. From the policy-relevant carbon budget viewpoint, this is good news, as it keeps the carbon budget approach resistant to scenario choice while complying with specific temperature targets and net-zero commitments. Regardless of the interpretation, the carbon budget scenario-dependent deviations identified are not permanent but a result of the optimization in a specific year. The arguably small deviation diminishes relatively quickly if no further emissions are added to the system. Furthermore, scenario-dependent deviations increase with the higher cumulative emissions cap but do not depend on the optimization year. Moreover, in 3.2.4, it was shown that allowing the system to produce negative emissions does not drastically increase scenario-dependent deviations. This shows us that the carbon budget approach is robust to scenario choice under FaIR.

The same conclusion cannot be made for the one-box model. As was shown, the one-box model produces up to 10 times larger scenario-dependent deviations, which evolve in time but do not disappear. The reasons for the dramatically different

[7] On the timescale of a human lifespan.

generated deviations are explained in detail in Sect. 4.1 through the shape of the pulse response function. Essentially, if the model's pulse response shows a large degree of similarity to the step function (dashed line, Fig. 1), i.e., if it peaks quickly, followed by a relaxation phase leading to a nearly constant value, the model will show small scenario-dependent carbon budget deviations. Looking at Fig. 1, one can see that FaIR is close to that behavior, while one-box is far from it. Joos et al. (2013) claim that having four carbon components and two temperature components is the minimum requirement to mimic this kind of pulse response. In climate economics, models often fail to meet this criterion. Besides the one-box model discussed here, the shape comparison of pulse responses presented in Fig. 1 in Dietz et al. (2021) shows that most simple climate models have some potential for carbon budget scenario dependency – adding weight to the argument for replacing climate models with FaIR in integrated assessments if carbon budget adherence is of importance (presumably, it is). Adherence to the carbon budget approach is especially important in the temperature target-based decision-making framework, as it is a crucial difference whether the temperature declines following emission cessation or is kept at the same level, as the carbon budget suggests. In the context of adhering to the temperature target, the declining temperature following emission cessation leads to non-intuitive policy recommendations, namely, to perpetually (albeit at a decreasing rate) continue emitting in order to adhere to the target.

In this spirit, let us consider the connection between the ZEC metric and the pulse response. If ZEC is 0, as the central estimate in MacDougall et al. (2020) suggests, this implies that temperature does not decrease or increase following the cessation of emissions. In the pulse response context, this requires that the pulse response is a step function, or close to it. Plotting the temperature leftover terms (Fig. 2) explicitly shows the two models' generated ZEC's under different climatic conditions (i.e., later in RCP run). Clearly, a model with a pulse response that does not show gradual relaxation (e.g. the one-box model) also shows a negative and declining ZEC. In contrast, FaIR produces a relatively small negative ZEC ($t_p = 2020, 2055, 2078$) that actually increases with changing climatic conditions, becoming slightly positive in $t_p = 2100$. This raises the question as to whether ZEC itself is a state-dependent value, i.e., whether the background climatic conditions dictate ZEC's value. This question is left to be explored in more advanced models.

Concluding that the carbon budget is indeed unaffected by emission scenario choice confirms the carbon budget approach's value as a tool for directly mapping cumulative emissions to temperature increase. However, the question remains as to the functional form of the carbon budget equation. Sect. 4 provides a clue as to how to deduct it from the pulse response representation. Namely, if TCRE is a constant, the carbon budget equation is linear. In Sect. 4, it was shown that the pulse response can be used as a proxy for TCRE, and that the pulse response decreases under changing climatic conditions in the default FaIR parametrization. A method was provided for deriving the non-linear carbon budget deviation from the changing pulse – a general method, which can be used for different models and different model calibrations. This offers an alternative approach to the non-linear carbon budget equation derived in Nicholls et al. (2020), as it does not assume a functional form of the non-linear carbon budget equation in advance, but derives it from TCRE dependency, building on Taylor expansion with respect to temperature, a key thermodynamic variable of the system investigated.

To address the lack of uncertainty in the analysis, Fig. 9 shows different pulse response representations for different FaIR calibrations. Following the methodology explained above, one can deduct that under different parameter sets, FaIR can mimic various levels of carbon budget non-linearity and even full linearity, while keeping scenario-independency robust, as TCRE, which approximates the corresponding pulse responses, can change its magnitude in either direction. This is possible because of the inclusion of both feedbacks on the carbon cycle and the temperature saturation, which counteract each other and can be tuned separately. Deriving the carbon budget equation explicitly for each calibration isn't pursued here, as doing so would not yield any new information and the set is too small to make generalized conclusions on e.g., how each FaIR parameter affects the carbon budget. Among other questions raised, this is an interesting aspect for future research.

6 Conclusion

This article focuses on deviations from the carbon budget approach, seen as a linear mapping from cumulative emissions and temperature increase, and draws a clear distinction between carbon budget emission scenario-dependent and climate state-dependent deviation. Scenario-dependent deviations are the possible differences in resulting temperature that are solely due to the preceding emission choice, while the cumulative emissions of the preceding emission pathway remain fixed. In contrast, state-dependent deviations underline the change in TCRE value, which depends on the change of background climatic conditions – specifically, the cumulative emissions and global mean temperature increase. Importantly, state-dependent TCRE leads to a non-linear carbon budget equation.

Section 2 introduces the reader to the FaIR, one-box and Green's function models. The FaIR model was chosen for the analysis for several reasons: Firstly, it has the ability to capture climate feedback on the carbon cycle; secondly, the model has already been praised in the literature for its efficiency; thirdly, it is relatively easy to implement and is computationally cheap; lastly, and most importantly for this paper, it can accurately capture temperature response to the emission pulse (i.e., pulse response). The one-box model is introduced to study the effects of structural model uncertainty, as the model provides an example of a dramatically different pulse response representation. At its core, this paper shows the implications of pulse response behavior on the carbon budget and its deviations, with the theory not restricted to the type of model under examination.

Section 3 derives maximum scenario-dependent deviations using FaIR in its default parametrization, through an optimization program provided in Eq. (4). The optimization procedure tests the entire portfolio of emission scenarios and diagnoses those that produce extreme potential temperature differences under the same cumulative emissions. FaIR shows that the maximum possible scenario-dependent deviations are small compared to the total temperature increase and gradually diminish, confirming the carbon budget's robustness when it comes to scenario choice. It was also shown that, by using the model's pulse response as a Green's function in Eq. (2), one can calculate the deviations with a correct order of magnitude. Hence, the Green's function approach offers a means of studying maximum possible scenario-dependent deviations in models of higher complexity and in a feasible computational time.

Section 4 shows that the shape of the pulse response dictates scenario dependency. On the other hand, the change of pulse response with background climatic conditions can be reinterpreted as the state-dependent TCRE, which then leads to the non-linear carbon budget equation. The method used to derive the carbon budget equation from pulse response, provided in Section 4, is universal and can be applied under different FaIR calibrations to see how individual climate drivers affect the non-linearity of the carbon budget. This, in combination with employing more complex models' pulse responses as Green's functions, opens a promising avenue for further research.

Code and data availability. The codes and data sets used in this analysis can be found online on https://doi.org/10.5281/zenodo.8314808

Competing interests. The author has declared that there are no competing interests.

Acknowledgements. M. Bekchanov provided a GAMS-coded version of FaIR from Bekchanov et al. (2022). B. Blanz helped fix certain code errors. V. Brovkin pointed out the paper in which the shape of the pulse response function was revealed (Joos et al., 2013). H. Held set this work in motion, as he suggested examining Green's function directly from emissions to temperature and using the minimization and maximization operations to assess scenario dependence. I would also like to thank B. Blanz and H. Held for their helpful discussions. Lastly, I would like to thank H. Held for carefully reading the manuscript. This research was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany‘s Excellence Strategy – EXC 2037 'CLICCS - Climate, Climatic Change, and Society' – Project Number: 390683824.

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, publisher: Nature Publishing Group, 2009.

Bekchanov, M., Stein, L., Mohammadi Khabbazan, M., and Held, H.: Accuracy of subsidiary climate targets (concentration and cumulative emission) as substitutes to a temperature target: trade-offs between overshooting and economic loss, 2022.

Blau, R. A.: Stochastic programming and decision analysis: an apparent dilemma, Management Science, 21, 271–276, 1974.

Dietz, S. and Venmans, F.: Cumulative carbon emissions and economic policy: In search of general principles, Journal of Environmental Economics and Management, 96, 108–129, https://doi.org/10.1016/j.jeem.2019.04.003, 2019.

Dietz, S., van der Ploeg, F., Rezai, A., and Venmans, F.: Are Economists Getting Climate Dynamics Right and Does It Matter?, Journal of the Association of Environmental and Resource Economists, 8, 895–921, https://doi.org/10.1086/713977, publisher: The University of Chicago Press, 2021.

Edenhofer, O., Bauer, N., and Kriegler, E.: The impact of technological change on climate protection and welfare: Insights from the model MIND, Ecological Economics, 54, 277–292, https://doi.org/10.1016/j.ecolecon.2004.12.030, 2005.

Forster, P., Storelvmo, T., Armour, K., Collins, W., Dufresne, J.-L., Frame, D., Lunt, D., Mauritsen, T., Palmer, M., Watanabe, M., Wild, M., and Zhang, H.: The Earth's Energy Budget, Climate Feedbacks, and Climate Sensitivity Supplementary Material, in: Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change, edited by Masson-Delmotte, V., Zhai, P., Pirani, A., Connors, S. L., Péan, C., Berger, S., Caud, N., Chen, Y., Goldfarb, L., Gomis, M. I., Huang, M., Leitzell, K., Lonnoy, E., Matthews, J. B. R., Maycock, T. K., Waterfield, T., Yelekçi, O., Yu, R., and Zhou, B., book section 7, Cambridge University Press, Cambridge, UK and New York, NY, USA, https://www.ipcc.ch/report/ar6/wg1/downloads/report/IPCC_AR6_WGI_Chapter07_SM.pdf, 2021.

Gillett, N. P., Arora, V. K., Matthews, D., and Allen, M. R.: Constraining the Ratio of Global Warming to Cumulative CO2 Emissions Using CMIP5 Simulations, Journal of Climate, 26, 6844–6858, https://doi.org/10.1175/JCLI-D-12-00476.1, publisher: American Meteorological Society Section: Journal of Climate, 2013.

Held, H.: Cost Risk Analysis: Dynamically Consistent Decision-Making under Climate Targets, Environmental and Resource Economics, 72, 247–261, https://doi.org/10.1007/s10640-018-0288-y, 2019.

IPCC: Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, Cambridge University Press, Cambridge and New York, NY, 2014.

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

Khabbazan, M. M. and Held, H.: On the future role of the most parsimonious climate module in integrated assessment, Earth System Dynamics, 10, 135–155, 2019.

Lahn, B.: A history of the global carbon budget, WIREs Climate Change, 11, e636, https://doi.org/10.1002/wcc.636, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/wcc.636, 2020.

Leach, N. J., Jenkins, S., Nicholls, Z., Smith, C. J., Lynch, J., Cain, M., Walsh, T., Wu, B., Tsutsui, J., and Allen, M. R.: FaIRv2.0.0: a generalized impulse response model for climate uncertainty and future scenario exploration, Geoscientific Model Development, 14, 3007–3036, https://doi.org/10.5194/gmd-14-3007-2021, publisher: Copernicus GmbH, 2021.

Leduc, M., Matthews, H. D., and Elía, R. d.: Quantifying the Limits of a Linear Temperature Response to Cumulative CO2 Emissions, Journal of Climate, 28, 9955–9968, https://doi.org/10.1175/JCLI-D-14-00500.1, publisher: American Meteorological Society Section: Journal of Climate, 2015.

MacDougall, A. H.: The oceanic origin of path-independent carbon budgets, Scientific Reports, 7, 10 373, 2017.

MacDougall, A. H. and Friedlingstein, P.: The origin and limits of the near proportionality between climate warming and cumulative CO2 emissions, Journal of Climate, 28, 4217–4230, 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., Burger, F. A., et al.: Is there warming in the pipeline? A multi-model analysis of the Zero Emissions Commitment from CO 2, Biogeosciences, 17, 2987–3016, 2020.

Masson-Delmotte, V., Zhai, P., Pirani, A., Connors, S. L., Péan, C., Berger, S., Caud, N., Chen, Y., Goldfarb, L., Gomis, M. I., Huang, M., Leitzell, K., Lonnoy, E., Matthews, J. B. R., Maycock, T. K., Waterfield, T., Yelekçi, O., Yu, R., and Zhou, B., eds.: Summary for policymakers, pp. 3–32, Cambridge University Press, https://doi.org/10.1017/9781009157896.001, 2021.

Matthews, H. D. and Weaver, A. J.: Committed climate warming, Nature Geoscience, 3, 142–143, 2010.

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, number: 7248 Publisher: Nature Publishing Group, 2009.

Meinshausen, M., Meinshausen, N., Hare, W., Raper, S. C. B., Frieler, K., Knutti, R., Frame, D. J., and Allen, M. R.: Greenhouse-gas emission targets for limiting global warming to 2 °C, Nature, 458, 1158–1162, https://doi.org/10.1038/nature08017, number: 7242 Publisher: Nature Publishing Group, 2009.

Millar, R., Allen, M., Rogelj, J., and Friedlingstein, P.: The cumulative carbon budget and its implications, Oxford Review of Economic Policy, 32, 323–342, https://doi.org/10.1093/oxrep/grw009, 2016.

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, publisher: Copernicus GmbH, 2017.

Nicholls, Z.: Reduced Complexity Model Intercomparison Project (RCMIP), pp. EGU21–3707, https://doi.org/10.5194/egusphere-egu21-3707, conference Name: EGU General Assembly Conference Abstracts ADS Bibcode: 2021EGUGA..23.3707N, 2021.

Nicholls, Z., Gieseke, R., Lewis, J., Nauels, A., and Meinshausen, M.: Implications of non-linearities between cumulative CO2 emissions and CO2-induced warming for assessing the remaining carbon budget, Environmental Research Letters, 15, 074 017, 2020.

Petschel-Held, G., Schellnhuber, H.-J., Bruckner, T., Toth, F. L., and Hasselmann, K.: The tolerable windows approach: theoretical and methodological foundations, Climatic Change, 41, 303–331, 1999.

Raupach, M. R.: The exponential eigenmodes of the carbon-climate system, and their implications for ratios of responses to forcings, Earth System Dynamics, 4, 31–49, https://doi.org/10.5194/esd-4-31-2013, 2013.

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

Rogelj, J., Shindell, D., Jiang, K., Fifita, S., Forster, P., Ginzburg, V., Handa, C., Kheshgi, H., Kobayashi, S., Kriegler, E., et al.: Mitigation pathways compatible with 1.5 C in the context of sustainable development, in: Global warming of 1.5 C, pp. 93–174, Intergovernmental Panel on Climate Change, 2018.

Shine, K. P., Fuglestvedt, J. S., Hailemariam, K., and Stuber, N.: Alternatives to the global warming potential for comparing climate impacts of emissions of greenhouse gases, Climatic change, 68, 281–302, 2005.

Shukla, P. R., Skea, J., Slade, R., Al Khourdajie, A., Van Diemen, R., McCollum, D., Pathak, M., Some, S., Vyas, P., Fradera, R., et al.: Climate change 2022: Mitigation of climate change, Contribution of working group III to the sixth assessment report of the Intergovernmental Panel on Climate Change, 10, 9781009157 926, 2022.

Stocker, T. F., Qin, D., Plattner, G.-K., Tignor, M., Allen, S. K., Doschung, J., Nauels, A., Xia, Y., Bex, V., and Midgley, P. M., eds.: Summary for policymakers, pp. 3–29, Cambridge University Press, Cambridge, UK, https://doi.org/10.1017/CBO9781107415324.004, 2013.

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, publisher: Proceedings of the National Academy of Sciences, 2009.