Optimal Climate Policy with Negative Emissions

RICCARDO REBONATO

DHERMINDER KAINTH

LIONEL MELIN

DOMINIC O'KANE

EDHEC-Risk Climate Impact Institute

Received (Day Month Year)

We can limit the future temperature impact of climate change in two ways: (i) reducing our use of CO2 emitting fuels as an energy source (abatement), and (ii) using negative emission technologies (NETs) to remove existing CO2 from the atmosphere (removal). Using a modification of the DICE model, we analyse the optimal use of these two policy responses to climate change. After calibrating the marginal costs of abatement and CO2 removal to the latest scientific information, we find that carbon removal must play a very important role in an optimal policy. If this policy is followed, we find that the Paris-Agreement 1.5-2 C warming by 2100 target is not just aspirational, but optimal. When an important role is played by NETs to control global warming, the decrease in carbon emissions can be more gradual, reducing transition risk and social dislocations. We examine the impact on the economy of large-scale carbon removal programmes, the potential for moral hazard and the logistic problems associated with the storage of the removed carbon.

Keywords: climate change; global warming; negative emissions technologies; integrated assessment models

1. Introduction

There is a scientific consensus[a] that the physical removal of carbon dioxide from the atmosphere, known as negative emissions, must complement the switch from fossil fuels to renewable sources of energy if global warming is to be limited to a temperature increase of 1.5-2 C above pre-industrial levels by 2100. Kolbert and Pacala (2017) have estimated that to stay below this limit, we need to remove 20% of current emissions by mid-century, and 40% by the end of the century.

[a]According to the 2021 IPCC report 'Carbon Dioxide Removal (CDR) is a necessary element to achieve net zero CO_2 and GHG emissions both globally and nationally, counterbalancing residual emissions from hard-to-transition sectors. It is a key element in scenarios likely to limit warming to 2°C or lower by 2100 (*robust evidence, high agreement*)' IPCCARWGIII (2021), emphasis in the original.

Integrated Assessment Models (IAMs), of which the DICE model by Nordhaus and Sztorc (2013) is arguably the most influential, have been instrumental to shaping climate policies[b]. However, they have either neglected or cursorily dealt with negative emissions. Although the DICE model does implicitly incorporate negative emissions, it incorrectly implies identical marginal costs for abatement and removal, and specifies a link between industrial emissions and abatement that can give rise to perverse results, as we discuss in Section 2.7.

In this paper, we analyse the optimal use of both abatement and negative emissions using a substantially updated version of the DICE model that is in line with the latest research on climate modelling. More precisely,

- we modify the climate-physics modules of the DICE model to reflect the latest scientific findings with regards to the impact over time of CO_2 emissions on the global average temperature;
- we update the damage function which links the change in the global average temperature to economic climate-change damages;
- we allow for the size of global-warming induced economic damages and the rate of economic growth to be stochastic;
- we model the resolution of parameter uncertainty for the damage exponent in a reduced-form Bayesian fashion, following Rudik (2020); and
- we differentiate between the marginal costs of abatement and CO_2 removal using information from the latest IPCC report (IPCC (2022)), and we modify the emission function in the DICE model (see Section 2.7) accordingly.

When we do so, we find that negative emissions must play a very important role not only in attaining the 'aspirational' goal of limiting global warming to 1.5-2 C by the end of the century,[c] but also as part of an *optimal* policy. To give an example of the magnitude of the optimal carbon removal, we note that, at peak, we should remove per annum approximately as much CO_2 as we are currently emitting. This has deep economic and policy implications. In particular, if something close to an optimal policy is followed, this will entail a substantial redirection of investments towards the removal sectors (investments which are at present marginal at best).

The most obvious funding for the necessary subsidies and investment are the current production and consumer fossil-fuel subsidies. Abatement remains essential, but, if negative emissions play the important role our study suggest they should, the abatement schedule can be more gradual than what an abatement-only policy would require, and, given current trends, more realistic. To compensate for this, our analysis shows that we should pursue in parallel aggressive removal policies. If

[b]The DICE model, in particular, is used by the US Environmental Protection Agency to inform government policy.

[c]We refer to the Paris Agreement 1.5-2 C target as 'aspirational' because it was not obtained as the result of a cost benefit-analysis or of an optimization. In the prevalent approach, optimal (e.g. minimum-cost) policies have been determined, after taking the 1.5-2 C target as an exogenous constraint. When seen in this light, our analysis 'endogenizes' the target.

both ‘levers’ (abatement and removal) are used together, the social and economic readjustments associated with the decarbonization of the economy can be more gradual, leaving more time for an orderly transition, and creating smoother and more manageable reallocations of labour and capital resources.

We stress that the importance of the 1.5-2C target is not reduced by our findings. If anything, we find that it becomes not just aspirational, but emerges as the optimal temperature path, achievable through the use of negative emission technologies. And we also emphasize that the importance of negative emissions should not be taken as an excuse for delaying abatement targets, which remain extremely demanding in our optimal solution. We simply obtain from an optimality perspective that abatement by itself is not enough: what scientists have been saying about removal is not just ‘common scientific sense’, but is the best course of action.

2. Our Approach

In order to obtain the optimal removal and abatement policy, we modify in important ways the DICE model. The model and the choices for its parameters have been discussed in Nordhaus and Sztorc (2013) and Nordhaus and Moffat (2017), to which the reader is referred for a detailed description. To make this paper reasonably self-contained, we describe its salient features in Section 2.1, and those features of the model that we have changed in Sections 2.2 to 2.7.

2.1. *The DICE Model*

The Dynamic Integrated Climate-Economy model (DICE) by Nordhaus and Sztorc (2013) is widely used to advise policymakers on the optimal strategies for decarbonization of the global economy. The DICE model combines a model of economic growth, a model for the resulting CO_2 emissions, a model of climate physics to determine the temperature impact of these CO_2 emissions, and a model to determine the economic damage of an increased global temperature on economic consumption. Its output is a calculation of the resulting current economic welfare, which is the sum of the discounted values of the consumption-based utilities between now and the long-term horizon. In the case of the DICE model this is 500 years from 2015. By maximizing welfare over the control variables (the savings rate and the abatement function, defined below), the DICE model determines the optimal CO_2 abatement pathway that maximises current welfare. The DICE model has been revised several times since its initial version, and here we begin by setting out the 2016R version.

DICE’s economic growth model is based on a neoclassical Ramsey model of an economy endowed with an initial stock of capital and labour, and an initial level of technology. The gross economic production $Y_g(t)$ is given by the Cobb-Douglas function

$$Y_g(t) = A(t) \cdot L(t)^{1-\gamma} \cdot K(t)^{\gamma}. \tag{2.1}$$

Gross economic output is a function of the amount of labour $L(t)$ and capital $K(t)$

in the global economy. The exponent γ is the capital elasticity of the production function and in the DICE 2016R model it has the value $\gamma = 0.3$.[d] The pre-factor $A(t)$ is the total factor productivity (TFP), i.e., the component of the production that is not explained by labour and capital, such as technological growth and improvements in efficiency. In the original version of the DICE model, it is assumed to increase deterministically over time. Economic output is then split into the cost of reducing CO_2 emissions (abatement), investment and consumption.

The next DICE module is the *emissions* model which translates gross economic output into CO_2 industrial emissions, E_{ind}, according to the formula

$$E_{ind}(t) = \sigma(t) \cdot Y_g(t) \cdot (1 - \mu(t)). \tag{2.2}$$

The parameter $\sigma(t)$ is the carbon-intensity of the global economy, and measures how much CO_2 is emitted per dollar value of production. Over time $\sigma(t)$ is expected to decline as the world becomes more efficient in its use of fossil-fuel-based energy. We discuss this formula and its shortcomings in detail in Section 2.7. As for the time-dependent function $\mu(t)$, it characterizes abatement, as it is the fraction of output produced using non-fossil-fuel-based energy.

Abatement comes at an economic cost given by

$$\Lambda(t) = Y_g(t) \cdot \Phi(t) \cdot \mu(t)^{\eta}, \tag{2.3}$$

where $\eta = 2.6$. The time-dependent function $\Phi(t)$ has been calibrated to the prices of backstop technologies for removing CO_2 from the atmosphere. See the discussion in Section 2.7 on this point.

The *physics* module tracks the movement of CO_2 between the atmosphere, the lower oceans and the upper oceans. This is required in order to model the long time lags between emission and reabsorption – we know that 35% of CO_2 emitted stays in the atmosphere for over 100 years. Knowing the quantity of CO_2 and other GHGs in the atmosphere allows us to calculate its radiative forcing and hence its impact on the global temperature anomaly, $T(t)$. This is the increase in the global temperature since pre-industrial times. The *damages* module links the temperature increase due to GHG to economic damages such as extreme weather events. The equation for the damages is given by

$$\Omega(t) = a_2 \cdot T(t)^{a_3} \tag{2.4}$$

where a_2 and the exponent a_3 are calibrated to loss data. See the discussion in Section 2.3 in this respect. Given the gross production, $Y_g(t)$, the DICE model

[d]At equilibrium, the exponent γ is equal to the fraction of output that goes to the providers of capital. Until recently, the stability of the labour/capital split (around 70-30%) has been an accepted 'stylized fact' of economic growth. An OECD (2012) report questions this received wisdom, since it finds the capital share to have increased from 33.9% to 3.3% between 1990 and 2009. A more recent survey, OECD (2015), finds that the average adjusted labour share in G20 countries decreased about 0.3 percentage points per year between 1980 and the late 2000s.

calculates the net production $Y_n(t)$ after damages and abatement as

$$Y_n(t) = Y_g(t) \cdot (1 - \Omega(t)) - \Lambda(t). \tag{2.5}$$

The post-abatement, post-damages, residual output $Y_n(t)$ is then split into productive investment,[e] $I(t)$, and consumption, $C(t)$:

$$Y_n(t) = C(t) + I(t). \tag{2.6}$$

The amount of investment is a fraction of net production $I(t) = s(t) \cdot Y_n(t)$, where $s(t)$ is the *savings rate*. For the global economy, capital falls over time due to depreciation, but increases due to investment following $K(t) = (1-\delta)K(t-1) + I(t)$.

Finally we have the *welfare* module. The DICE model assumes a time-separable utility function, and uses the evolution of consumption over time to determine the future per-period utility and the present value of the welfare. This is accomplished by discretizing the 500-year final horizon into 100 5-year time-steps. The time zero welfare, $W(0)$, is calculated according to

$$W(0) = \sum_{t=1}^{100} \frac{U(\tilde{c}(t), L(t))}{(1+\rho)^t} \tag{2.7}$$

where $\tilde{c}$ is the utility per person ($\tilde{c} = C(t)/L(t)$), and $U(\tilde{c}(t), L(t))$ is the power Constant-Relative-Risk-Aversion (CRRA) utility function given by

$$U(\tilde{c}(t), L(t)) = L(t) \cdot \frac{\tilde{c}(t)^{1-\alpha} - 1}{1-\alpha}. \tag{2.8}$$

The utility discount rate, ρ, is the rate of social time preference per year, which Nordhaus sets to 1.5%. The parameter α is the relative risk aversion (RRA) parameter which Nordhaus sets to $\alpha = 1.45$. Within a CRRA framework, it is the reciprocal of the elasticity of intertemporal substitution (EIS). See the discussion of this important point in Section 2.2.

With any equilibrium economic-growth model (of which DICE is a particular case) a zero abatement strategy leads to increased greenhouse gas (GHG) emissions, that lead to increased global warming, that leads to increased economic damage and so to a reduced welfare. At the opposite policy extreme, a sudden 100% decarbonization also has a major negative impact on economic growth and so a negative impact on welfare. The optimal abatement policy lies between these two limit policies and is obtained by solving the following control problem

$$W^*(0) = \max_{\mu^*(t), s^*(t)} \left[\sum_{t=1}^{100} \frac{U(\tilde{c}(t;\mu,s), L(t))}{(1+\rho)^t} \right], \tag{2.9}$$

with the optimal abatement and savings vectors $\mu^*(t)$ and $s^*(t)$ playing the role of control variables. These are determined using standard non-linear maximization algorithms.

[e]The adjective 'productive' refers to investment to produce consumption goods, rather than abatement or removal.

2.2. *RRA and EIS*

The choice of the utility discount rate chosen by Nordhaus for the DICE model, specifically that $\rho = 1.5\%$, has been discussed at great length in the literature (see, e.g., Nordhaus and Moffat (2017), Stern (2007), Pindyck (2013)) and contrasted with the choice made by Stern in his review which is that $\rho = 0.1\%$. Much less attention has been paid to the respective roles of the elasticity of intertemporal substitution (EIS) and the relative risk aversion (RRA). For any time-separable utility function these two quantities must be reciprocal of one another. However, as Ackerman *et al.* (2010) point out, for problems such as climate change, they are 'on a collision course': high aversion to static risk *must* imply high aversion to unequal consumption (and hence a large part of the abatement bill placed on the shoulders of our richer grandchildren). If, to obviate this feature one choses a smaller degree of aversion to uneven consumption, this would imply little static-risk aversion, not a palatable choice given the deep uncertainty surrounding climate outcomes. One way out of the impasse would be to use recursive utility function à la Epstein and Zin (1989) (and, indeed, we are pursuing work in this direction), but, for the level of time-resolution and modelling richness we employ for our study, this entails formidable computational challenges.

Fortunately, Ackerman *et al.* (2010), find that optimal solutions 'are sensitive to the intertemporal elasticity of substitution, but *remarkably insensitive to risk aversion*' (emphasis added). Along similar lines, Crost and Traeger (2014) and Cai *et al.* (2016) find a small role for risk aversion in the social cost of carbon and Belaia *et al.* (2017) also find that risk aversion has less impact than other factors in determining optimal abatement policies. It is therefore more important to capture correctly the elasticity of intertemporal substitution than relative risk aversion.

Unfortunately, there is no consensus in the literature about its value. Since determining the optimal pace of decarbonization is not the goal of this paper, and since a large body of literature takes the DICE choice as reference, we present most of our results for the case of EIS = 0.69. For completeness, we also report the results for the case of EIS = 1.45 (close to the value recommended by Bansal and Yaron (2004)) to show that the importance of negative emissions is a robust feature of our analysis.

2.3. *The Damage Function*

The damage function is one of the key determinants of DICE's optimal policies (see, e.g. Barnett *et al.* (2020)). We therefore look at this aspect in some detail. The damage function in the DICE model takes the form:

$$\Omega(t, T) = a_2 T(t)^{a_3} \tag{2.10}$$

where $T(t)$ is the temperature anomaly at time t. To determine the parameters, a_2 and a_3, Nordhaus (2017) and Nordhaus and Moffat (2017) conduct a meta-analysis of studies of monetary global damage estimates estimated as a function of the temperature anomaly (which is assumed to be less than 4°C). Using a least squares regression, they estimate $a_2 = 0.00236$ and $a_3 = 2$. The problem with this approach is that it parameterizes imprecisely observed damages over a small range of low temperature anomalies; yet, within the model we must extrapolate to levels of warming for which we have no empirical observations. A number of authors have argued that this parametrization is implausible, particularly for large temperature anomalies; to address this, Weitzman (2010) and Botzen and van der Bergh (2012) suggest values of a_3 as high as seven for high temperature anomalies. Howard and Sterner (2017) also emphasise issues with the methodological approach of Nordhaus (2017).

Estimating climate damages as a function of temperature is now a burgeoning field of research. Tol (2022) presents a comprehensive meta-analysis of studies dating back over the past 20 years. We have taken this data presented and plotted it in fig. 1, showing the large variation in estimated damages as a function of temperature anomaly. The estimates are so widespread, in part, because they come from different types of study: i) enumerative methods (where the observed estimates of the physical effects of climate change are obtained one by one from natural science papers, added up and then extrapolated using simple functional forms, see e.g., Tol (2002), Hope (2011)); ii) direct econometric methods (where observed differences in prices, expenditures, self-reported happiness or total output are regressed against contemporaneous or lagged variations in climate and then extrapolated, again using simple functional forms see e.g.,Kalkuhl and Wenz (2020),Kahn *et al.* (2021),Burke *et al.* (2015)); iii) computational general equilibrium models, that can simulate outcomes under counterfactual conditions that haven't occurred in the past and therefore are not restricted in the way econometric models are e.g., Takakura *et al.* (2019)); and iv) expert elicitation (e.g., Howard and Sylvan (2020)). In Figure 1 we show the issues with using a global fit to the data[f]; the importance of capturing the variability is emphasised.

Given the uncertainties associated with the damage function, we have proceeded as follows. First, we require that the observed damages associated with the warming already experienced should be recovered. Next, we note that a global warming of 6C over pre-industrial levels is associated with losses ranging from a few percentage points (econometric methods) to approximately 20% (elicitation methods). We use the parametrization developed by Howard and Sterner (2017) ($a_2 = 0.007438, a_3 = 2.0$) in their meta-analysis of damage studies as our base estimate. We then note that the damages obtained for 6C of warming by the various methods are consistent with the parameter a_3 still centred around a value of 2, but with a distribution whose probability mass is mainly located between 1 and 3.5. We then make a_2 a function of the stochastic a_3 so as to recover the damages that have actually been observed. We have therefore allowed the damage fraction to have the distribution shown in

[f] We show a quadratic fit; this assumes that each study is independent and equally valid, which is not necessarily the case, see e.g., Howard and Sterner (2017)

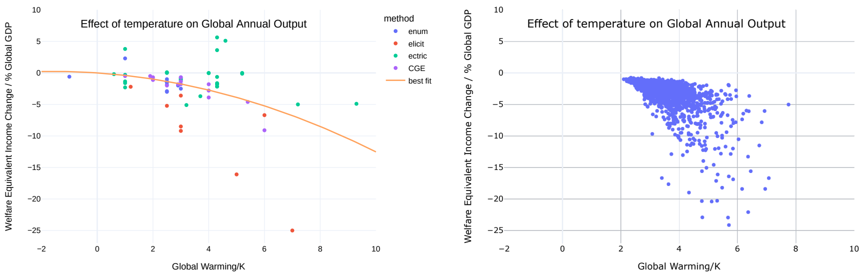

Fig. 1: Economic losses due to climate change temperature increases. The graph on the left shows data points where the colors represent the four methods discussed in the text. We show a quadratic fitted line. The graph on the right shows 2,048 Monte Carlo draws of the simulated losses according to our economic damage model specification.

the simulated draws in Figure 1.

2.4. *Modelling Uncertainty in Economic Growth*

We model uncertainty in economic growth using the Jensen and Traeger (2014) modification of the influential 'long-term growth' Bansal and Yaron (2004), Bansal and Shaliastovich (2012) model. In this approach, the variability in economic outcomes arises from uncertainty in the growth process for the TFP, $A(t)$, denoted by $g_A(t)$. TFP growth is assumed to have both a deterministic and a random component, so that the uncertain technology level one period ahead is given by

$$A(t+1) = A(t) \cdot \exp(g_A(t)) = A(t) \cdot \exp\left(g_A^{\text{det}}(t) + z(t)\right), \tag{2.11}$$

where $g_A^{\text{det}}(t)$ is the deterministic growth trend and $z(t)$ is a stochastic component, i.e., a growth shock. The growth shocks are in turn assumed to consist of two uncorrelated shocks:

$$z(t) = x(t) + w(t) \tag{2.12}$$

where $x(t)$ and $w(t)$ are assumed to be normally distributed and independent. While $x(t)$ follows a Brownian motion, $w(t)$ is assumed to follow a AR(1) process i.e.,

$$\begin{aligned} x(t) &= \mu_x + \sigma_x \cdot dz^x(t) \\ w(t) &= \zeta \cdot w(t-1) + \epsilon(t) \\ \epsilon(t) &= \mu_\epsilon + \sigma_\epsilon \cdot dz^\epsilon(t) \end{aligned} \tag{2.13}$$

where $dz^\epsilon(t)$ and $dz^x(t)$ are independent and $N(0,1)$. The deterministic component of the technology process is assumed to decay with time following the Nordhaus specification:

$$g_A^{\text{det}}(t) = g_A^{\text{det}}(0) \cdot \exp(-\delta_a \cdot t). \tag{2.14}$$

The volatilities, σ_x and σ_ϵ are estimated by Jensen and Traeger (2014) so that $A(t)$ is consistent with empirical long run US TFP data and consistent with Bansal and Yaron (2004). Both are set at 1.9%. The use of an AR(1) process captures the empirically observed strong persistence of TFP. We follow Jensen and Traeger (2014) in setting $\zeta = 0.5$. The drift terms, μ_x, μ_ϵ are developed by requiring that the overall mean of the growth rate of the TFP, $g_{A,t}$ matches the deterministic growth rate component. For this value of ζ, we find $\mu_x = \mu_\epsilon = -3.6107 \times 10^{-4}$.

Bansal and Yaron (2004) show that this process for consumption is consistent with annual consumption data, and that, when coupled in an Epstein and Zin (1989) recursive-utility framework, it accounts well for the equity risk premium, the level of rates, the volatility of the market return, and the price-dividend ratio. We cannot recover such a rich description of the economy in our setting, because we do not use recursive utility functions. However, we inherit the same persistent volatility for the process for consumption.[g] So, the total factor productivity follows the same deterministic decline over time as in DICE model, but displays on top the stochastic behaviour in Equation 2.11-2.13.

2.5. *Updating the Climate Physics of DICE*

The climate physics modules of any IAM should provide a realistic description of the link between net GHG emissions and the associated planetary temperature increase. A 'good' model should be consistent with the scientific consensus in climate research. We take as representative of this consensus the output of the Coupled Model Intercomparison Projects (CMIP) produced by the IPCC, which compare the outputs from a number of academically accepted Global Climate Models (GCMs).

GCMs are sophisticated models of the global climate system attempting to capture atmospheric, oceanic, ice and terrestrial behaviour. A number of distinct models have been built up over the past 30 years, and they are still being improved. The CMIPs make available periodic standardised comparisons of the output of these models, for climate adaptation, mitigation and resilience planning. Their high computational costs, however, render GCMs unsuitable for studying the feedback between the planetary system and human behaviour. The academic community has therefore also developed or repurposed a number of simplified tools (climate emulators) as reduced-form versions of the more complex climate change models.

Unfortunately the output from some economic climate emulators, specifically the one used in DICE, have become divorced from the output of the more recent GCMs. To complicate matters, analysis of the output from the most recent tranche (CMIP6) of GCMs by climate scientists has demonstrated that the output of a

[g]Bansal and Yaron (2004) specify a process for consumption. In the DICE model, consumption is a 'residual', i.e., is obtained from the optimization process. However, the DICE-model dynamics create a strong correlation between consumption and the total factor productivity, and therefore assigning the dynamics in Equation 2.11-2.13 to the total factor productivity generates a very similar process for consumption.

number of these is unsatisfactory and must be treated with caution. To understand the origin of the problem, recall that the CMIP6 collates the results from around 100 different climate models developed by approximately 50 different climate research groups worldwide. This suite of models have been under development since 2013 with final results available in 2021.

Climate modellers define standardised metrics to parametrize the behaviour of a climate model. One is the transient climate response (TCR), or the amount of global warming in the year in which atmospheric CO_2 concentrations have doubled after steadily increasing by 1% every year. A second metric is equilibrium climate sensitivity (ECS), the eventual long-term temperature response to an instantaneous doubling of atmospheric CO_2 concentration. Figure 2 shows the probability density function of the ECS and TCR calculated using the different models underlying CMIP5 and CMIP6. It is apparent that the distribution of the equilibrium climate sensitivity is significantly wider for the CMIP6 suite than for the CMIP5 suite, and a similar picture emerges for the transient climate response.

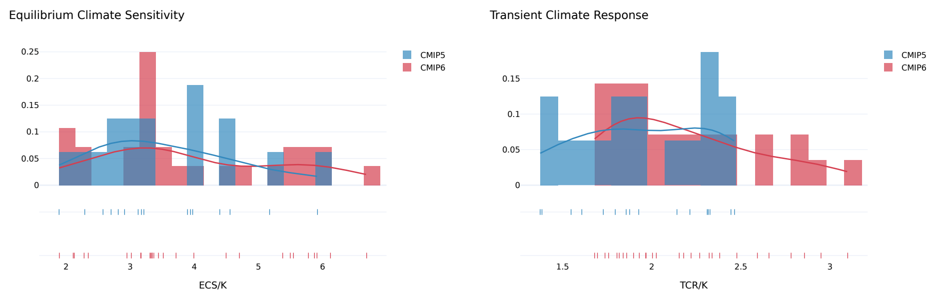

Fig. 2: Probability density function of the Equilibrium Climate Sensitivity (left) and Transient Climate Response (right) from CMIP5 models (red) and CMIP6 models (blue). See text for a discussion.

At first blush, this would be a source of considerable concern, as it would suggest that the Earth is significantly more sensitive to GHG emissions than previously thought. However, the current scientific consensus is that the group of models with high ECS (and TCR) should be down-weighted or rejected, because they have been found to do a poor job of reproducing historical temperatures over the past 150 years[h], or of simulating the climates of the distant past[i]. Hence, a significant proportion of the CMIP6 model suite needs to be treated with some caution.

[h]these models often show no warming over much of the twentieth century and then a sharp warming spike in the past few decades, see Liang *et al.* (2020), which is inconsistent with empirical observations

[i]In particular, they model the last ice age as being much colder than palaeoclimate evidence indicates; see Zhu *et al.* (2020) and Tokarska *et al.* (2020)

Previously (e.g., for CMIP5) the IPCC and indeed economic IAMs referenced the mean and standard deviation of the entire suite of models to provide a 'best' estimate of the impacts of GHG emissions and the associated uncertainty. However, this is only valid if each model is independent and equally valid, assumptions which, as discussed, are questionable for the CMIP6 ensemble. In reality, the large number of poorly performing models introduces a systematic bias, which means that simply using the mean and standard deviation of all models is inappropriate.

A number of approaches have been proposed to address this: for instance the IPCC use a weighted average to produce '*assessed* global warming' projections, with greater weight given to those models that are consistent with historical temperature records in Ribes *et al.* (2021). We follow the simpler prescription of Hausfather *et al.* (2022) who suggest filtering out those models which do not produce ECS and TCRs in line with the assessed warming profiles. For instance they suggest that we choose those models whose ECS (TCR) lies in the range $[2.5°\mathrm{C}, 4°\mathrm{C}]([1.4°\mathrm{C}, 2.2°\mathrm{C}])$ *assessed* as being likely (66% probability). This filters out approximately 55% (40%) of the CMIP6 models. In this spirit, but perhaps more conservatively, we keep those models whose ECS lies in the range $[2°\mathrm{C}, 5°\mathrm{C}]$, which is *assessed* as being very likely.

2.6. *Updating the Physics of the DICE Model*

The climate emulators found in IAMs such as DICE typically consist of two sub-models: a temperature model, which determines how the CO_2 concentration in the atmosphere translates into an (increasing) average temperature (as per the models above) and a 'carbon cycle' model, which predicts how atmospheric CO_2 concentrations evolve as a result of biogeochemical processes (e.g., dissolution into the ocean, weathering, photosynthesis etc).

Unfortunately these economic climate emulators are no longer consistent with the behaviours predicted by more sophisticated models unless appropriately recalibrated. In particular, Folini *et al.* (2021) and Dietz *et al.* (2007) demonstrate that the emulator used in DICE fails standard 'sanity checks' used in the climate-science literature to test a model's reliability.

2.6.1. *Recalibrating the Temperature Model*

Figure 3 illustrates the time evolution of the global mean surface temperature in response to an instantaneous quadrupling of atmospheric CO_2 starting from pre-industrial levels (285 ppm CO_2). Results for 150 years[j] are illustrated for a range of the CMIP6 GCMs and the DICE-2016 model. The latter behaves qualitatively differently: all of the GCMs suggest that the earth will heat up much faster (within 10-to-30 years) than the DICE model predicts.

[j] This represents a comparatively short time horizon, the equilibration time for the Earth is over 1000 years.

We also see the 'hot-tail' issue: the grey lines represent models which are excluded on the basis of a predicted ECS falling outside the temperature range assessed by climate experts on the basis of wide ranging evidence as being 'very likely'. Taking the subset of 'appropriate' CMIP6 models we develop a multi-model mean by parameter averaging, which then defines our revised temperature model within DICE. We highlight that this model is similar to the multi-model mean of the CMIP5 ensemble.

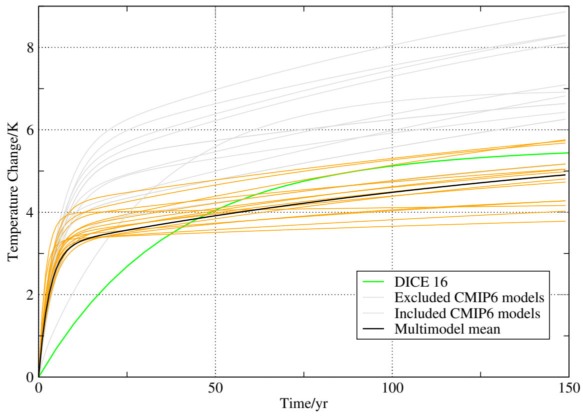

Fig. 3: Temperature response to instantaneous quadrupling of atmospheric CO_2 with respect to pre-industrial values. Each of the orange and grey lines represents a CMIP6 model, parametrized using a two-box EBM. The orange lines represent models of which the black line represents the multi-model mean. The response according to the DICE 2016 model is shown by the green curve.

2.6.2. *Recalibrating the Carbon Cycle Model*

The second test case pertains to the carbon cycle model and follows the evolution of atmospheric CO_2 in the wake of an instantaneous release of 100 GtC into today's atmosphere. Our analysis largely follows the work of Folini *et al.* (2021). Joos *et al.* (2013) have collated the responses from a suite of earth system models, and parametrized them using a sum of exponential decay functions. These are shown in Figure 4 (grey lines), along with the multi-model mean (red line) and the output from DICE-2016 (blue line). Once again, while there is significant variability in the output from the earth science models, the predictions of the DICE model lie well outside this range.

We note that the DICE carbon model is a three-state Markov-chain system, modelling the transfer of carbon from the biosphere into the upper oceans and eventually into the deep ocean. As such, a unique reparametrization of the DICE emulator to recover the model is not possible. However, the flexibility of the carbon

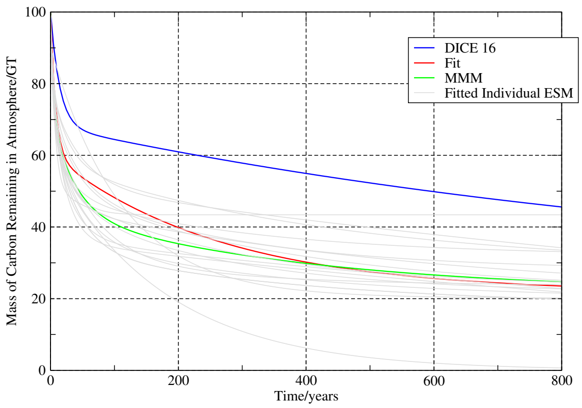

Fig. 4: Removal of a 100GtC emission impulse (47ppm CO_2) on an initial background concentration of 389ppm in climate science models and DICE. The climate science data is drawn from Joos *et al.* (2013). DICE-2016 removes too little of the CO_2 from the atmosphere relative to all of the GCMs. This CO_2 then remains in the atmosphere causing increased warming.

cycle model in the DICE model allows us to fit the outputs of the earth science models (including for example the multi-model mean) to good precision on the timescales of interest. We have given the parameters (as per Folini *et al.* (2021)) in Appendix A.2.

2.7. *Calibration of Marginal Costs of Abatement and Removal*

DICE allows for 'net negative emissions'. However, the approach is somewhat rudimentary and a key focus of our work is to improve upon this. By way of background we introduce negative emission technologies (NETs) in section 2.7.1; we then explore how abatement (and NET) are currently modelled in DICE through 'generalised' marginal cost curves. These cost curves specify the marginal cost of abatement both as a function of abatement fraction and their evolution with time – describing the cheapening of abatement as e.g., technology improves. We have then extended this approach to CO_2 removal, basing our cost curves on the latest IPCC survey. This is discussed in 2.7.3.

2.7.1. *Background to Negative Emission Technologies (NET)*

The IPCC defines NET as anthropogenic activities that remove CO_2 from the atmosphere and durably store it. NETs are in reality a range of technologies from nature-based practices, such as forestation, soil carbon sequestration and wetland restoration, to technological alternatives such as enhanced weathering, bioenergy with carbon capture and storage, and direct air capture and storage. Table 1 provides a very brief overview of these technologies; for further details we refer to the comprehensive reviews in Fuss *et al.* (2018). Broadly speaking, nature-based solutions are already-existing ('good-to-go') technologies. However, they are limited in the amount of carbon they can remove as there is competition for land use. Conversely, technological approaches, while in principle less limited in scope, frequently require large amounts of energy and are therefore somewhat self-defeating unless energy supplies are significantly decarbonized.

Table 1: A summary of Negative Emission Technologies.

Technology	Brief Description
A/R	Afforestation and Reforestation • Planting new forests and better land management can sequester CO_2. • Effectiveness is latitude dependent (most effective in tropics); death and decay of forests leads to release of CO_2, so forest sinks saturate.
Wetlands	• Better wetland management can also sequester CO_2 in temperate latitudes
SCS	Soil Carbon Sequestration • Better soil management aided by techniques such as bio-char can capture atmospheric CO_2 as biomass and store it in the soil, enhancing fertility
BECCS	Bio-Energy with Carbon Capture and Storage • CO_2 extracted from atmosphere by biomass when it grows • Energy is extracted from biomass through, e.g., combustion, fermentation etc
Biochar	• Capture the CO_2 either as a gas to be stored geologically or as biochar. Biochar is the solid residue left over after the pyrolysis of biomass.
EW	Enhanced Weathering • Weathering is a natural geological process whereby particular rock types (e.g., basalt) are broken down and atmospheric CO_2 is sequestered as minerals; EW speeds this up by grinding rocks down to dust • Energy costs (grinding rock) are high; but potential fertilisation of soil.
DACCS	Direct Air CO_2 Capture and Storage • Large fans draw in air from the atmosphere; CO_2 is filtered out by amine / hydroxide solvents. These are then heated and the released concentrated CO_2 stored geologically. • Energy requirements (e.g., for releasing CO_2 from the sorbent, pumps) are high
OA/OF	Ocean Alkalinisation and Ocean Fertilisation Very speculative technologies, whereby oceans' acidity is reduced (allowing greater dissolution of CO_2) or are fertilised to allow for greater biomass. Huge amounts of iron/fertilisers required; more akin to geo-engineering.

2.7.2. *Incorporating Negative Emission Technologies in DICE*

In the DICE model, the emissions of carbon, $E_{\text{ind}}(t)$, due to industrial production and abatement are given by Equation 2.2, where $\sigma(t)$ is the carbon intensity of the economy (i.e., the extent to which a unit of production leads to GHG emission) and $\mu(t)$ is the abatement fraction. The carbon intensity is assumed to be a decreasing function of time, reflecting the fact that as economies evolve, becoming e.g., more service-oriented, a smaller quantity of GHG emissions are required per unit of production.

The original specification of DICE allows for net negative emissions from the year 2165 onwards. While not explicitly stated, it is likely that 'near term' negative emission technologies such as reforestation are accounted as part of the abatement function, and that DICE models as NETs 'industrial' technologies such as direct air capture only. Negative emissions are then simply treated as abatement rates larger than one; furthermore they are constrained not to exceed 20% of the business-as-usual industrial emissions in each time step.

Before examining how to incorporate negative emissions, we first examine how the marginal cost of abatement is determined within the standard DICE model. Following Equation 2.2, the quantity of carbon abated is given by

$$E_{\text{abate}}(t) = \sigma(t) \cdot Y_g(t) \cdot \mu(t). \tag{2.15}$$

The abatement $\mu(t)$ is a control variable which is optimised over. The amount of abatement chosen is determined by its cost, $\Lambda_{\text{abate}}(\mu(t), t)$, and the DICE model parametrizes the marginal cost of abatement as:

$$\frac{\partial \Lambda_{\text{abate}}(\mu(t), t)}{\partial E_{\text{abate}}} = p_{\text{back}}(t) \cdot \mu(t)^{\lambda - 1}. \tag{2.16}$$

where $\lambda = 2.6$. This power law captures the fact that while the first tonne of carbon costs nothing to abate, the marginal cost of abatement increases rapidly with each successive tonne of CO_2 abated. Once the economy is fully decarbonized, removing one more tonne of CO_2 at time t costs $p_{\text{back}}(t)$ real US dollars, which is a decaying and deterministic function of time. Both λ and $p_{\text{back}}(t)$ were estimated by Nordhaus (2017) and Nordhaus and Moffat (2017) on the basis of observed and projected abatement costs.

Given the potential importance of NET, there are at least two areas where it is important to improve upon this approach. First, the same marginal costs (post 2165) are assumed in the original DICE model for abatement and for negative emissions. While this might be appropriate for the distant future, earlier adoption of NETs is a distinct possibility. The marginal costs of abatement and removal and how they evolve in the near term, while linked, are quite different. This is because the technologies are at different stages of development with industrial NET at an earlier stage. Furthermore, there is likely a sequencing, with industrialised NET requiring significant decarbonization of energy supply before they become truly viable.

The second significant problem with Equation 2.2 is that when μ is greater than one, the greater the output and the higher the carbon intensity, $\sigma(t)$, of the economy, the more negative the emissions. As a consequence, when the decarbonization of the economy is low (σ is high), the optimizer perversely seeks to increase both economic output and abatement to produce negative emissions.

To model NETs in a more satisfactory way, we therefore introduce an additional set of control variables allowing the central planner to divide post-consumption economic output among savings, abatement and *carbon removal*. Our only constraints on the degree and timing of CO_2 removal are (i) to require that the atmospheric carbon concentration cannot be reduced below pre-industrial levels, and (ii) to ensure that industrial emissions in Equation 2.2 remain non-negative. The resulting optimal policy can then capture the interplay between abatement and removal costs. As in the DICE model, the key to modelling this behaviour lies in the marginal cost curve of carbon dioxide removal both as a function of amount of CO_2 removed, and as a function of time, which we now model and calibrate.

2.7.3. *Marginal Cost Curves for NETs*

To develop the marginal cost curves for NETs, we make use of the comprehensive review in IPCC (2022) and the work of Fuss *et al.* (2018). Appendix A.3 contains the raw data extracted from these works detailing expected technological readiness as well as the marginal cost of removal per tonne of CO_2. We use $P_{\min}$ and $P_{\max}$ to represent the minimum and maximum US dollar price for CO_2 removal at rates $Q_{\min}$ and $Q_{\max}$ respectively. The CO_2 removal rates are given in units of $GtCO_2$ per annum. The IPCC gives estimates of the range of expected removal rates by 2050 for these two prices. We also define a midpoint estimate $Q_{\text{mean}} = (Q_{\min} + Q_{\max})/2$.

Table 2: Proposed scenarios for the take-up of Negative Emission technology showing which NETs are effective in 2050 and 2100.

	2050	2100
Pessimistic Case	Biochar and DACCS remove at rate $Q_{\min}$	Biochar and DACCS remove at rate Q_{mean} BECCS removing $Q_{\min}$
Optimistic Case	Biochar and DACCS removed at rate Q_{mean}	Biochar and DACCS removed at rate $Q_{\max}$ BECCS remove Q_{mean} EW allowed remove $Q_{\min}$

We focus our attention on the next 100 years, for which the IPCC provides a forecast. While the IPCC report gives a good indication of the removal costs, their time evolution is less explicitly detailed. We therefore infer this by considering technological readiness[k]. Specifically, we develop two scenarios – an optimistic (high uptake/development of NET) and a pessimistic scenario for 2050 and 2100 based around the readiness of the various NET. The scenario assumptions are detailed in Table 2. This can then be translated into optimistic and pessimistic marginal cost curves for CO_2 removal, which we have plotted in Figure 5. We are therefore able to develop a 'term structure' for the marginal cost curve for CO_2 removal. For 2050 and 2100, we take the mid-points of the optimistic and pessimistic scenarios. For the longer term horizon, we assume that the marginal costs for abatement and removal are approximately consistent with one another.

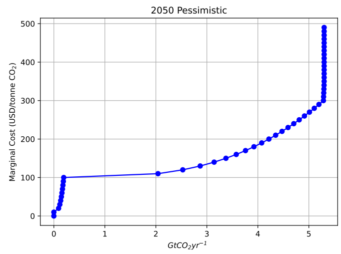

(a) 2050 Pessimistic scenario

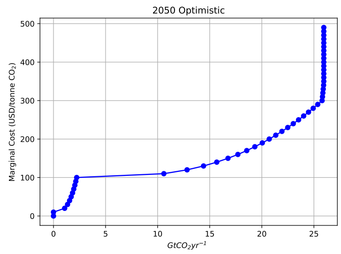

(b) 2050 Optimistic scenario

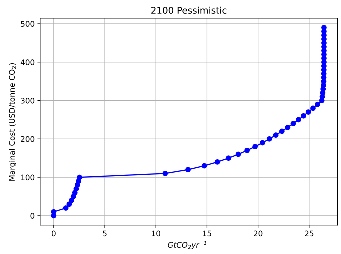

(c) 2100 Pessimistic scenario

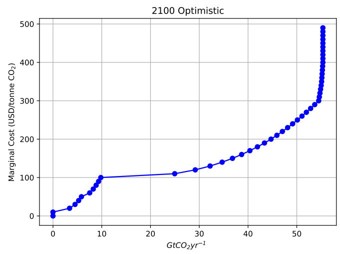

(d) 2100 Optimistic scenario

Fig. 5: Marginal costs in 2050 and 2100 of CO_2 removal (in real US dollars per tonne) for the amount of CO_2 removed ($GtCO_2$ /year) displayed on the x axis. This is shown in the Pessimistic and Optimistic scenarios which are detailed in the text.

We note from Figure 5 that the marginal cost of carbon removal has two distinctive features. First, contrary to abatement, the first tonne of carbon to be removed and stored has non-zero cost. We would expect this on purely thermodynamic grounds, as removing CO_2 from the atmosphere and then storing it will require energy. Second, the marginal cost becomes effectively 'infinite' beyond a

[k]Abatement and removal marginal costs for very long horizons are extremely speculative, but they have a small effect on the optimal solution because they are heavily discounted.

time-dependent level of removal. This is in contrast to the abatement cost which is modelled as a power law. While one might argue the merits of these choices,[l] the cost structure again reflects the fact that, at any given time, only some removal technologies are available on the scale required, and so the cost should therefore be more rapidly increasing. We reflect this by modelling the cost of removal, $\Lambda_{\rm rem}(Q_{\rm rem}(t), t)$, with an exponential marginal cost structure given by

$$\frac{\partial \Lambda_{\rm rem}(Q_{\rm rem}(t), t)}{\partial Q_{\rm rem}} = p_{\rm back}(t) \cdot \nu(t) \cdot e^{\xi(t) \cdot Q_{\rm rem}(t)}, \tag{2.17}$$

where $Q_{\rm rem}(t)$ is the rate at which CO_2 is removed at time t. The product $p_{\rm back}(t)$ times $\nu(t)$ characterizes the level at a given time t of the marginal cost of removal (i.e., $\nu(t)$ represents the cost differential between removal and abatement technologies) while $\xi(t)$ captures the curvature, i.e. the dependence of cost on the amount of CO_2 removed. We assume the following behaviour in time:

$$\begin{aligned} \nu(t) &= \nu_\infty + (\nu_0 - \nu_\infty) \cdot e^{-\nu_{decay} \cdot t} \\ \xi(t) &= \xi_\infty + (\xi_0 - \xi_\infty) \cdot e^{-\xi_{decay} \cdot t}. \end{aligned} \tag{2.18}$$

More precisely, we have fitted the cross sectional cost of removal (developed from the data in Figure 5). The results are shown in Figure 6. This shows the marginal costs (in USD per tonne of CO_2) for abating or removing the quantity of CO_2 (in $GtCO_2$ /year) on the x axis for horizons from 5 to 250 years. We give in Appendix A.4 the parameters that we obtain following this procedure.

3. Results

In this section we present our results as follows: we first discuss the importance of the EIS coefficient in determining the optimal abatement policy (Section 3.1); then we present the optimal policies when negative emissions are allowed (Section 3.2).

3.1. *The Impact of EIS on the Abatement Schedule*

Before presenting the results showing the effect of negative emissions on the optimal policy, we present in Figure 7 the optimal solutions in a deterministic setting and without negative emissions. Assuming an EIS of 0.69, full decarbonization of the economy is reached close to the end of the century; the temperature anomaly is well above 2 C (2.2C by 2100) and the amount of CO_2 in the atmosphere peaks at well over 1,000 Gt CO_2 (corresponding to an atmospheric concentration of CO_2 of 470 ppm, a 20% increase over 2016 levels) in fifty years[m]. When the EIS is increased to the levels suggested in Bansal and Yaron (2004), i.e. the EIS is 1.45, the economy

[l]For instance, it is likely that beyond a certain level also the marginal cost of abatement will become very high, far in excess of that specified by the power law.

[m]These results differ from those in Nordhaus (2017) because we employ the more 'severe' damage function from Howard and Sterner (2017). The damage exponent is the same as in the original DICE model. See the discussion of this point in Section 2.3.

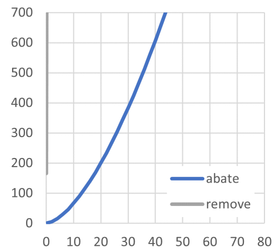

(a) 5-year horizon

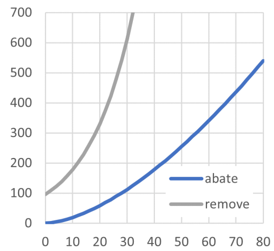

(b) 30-year horizon

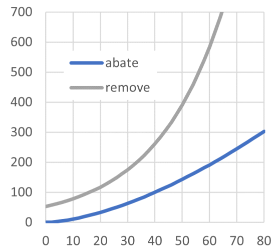

(c) 80-year horizon

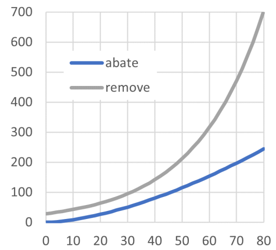

(d) 150-year horizon

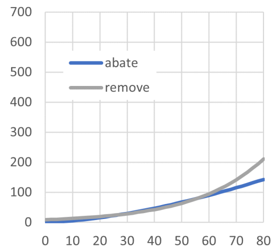

(e) 250-year horizon

Fig. 6: The marginal costs (in USD dollars per tonne of CO_2) for abating or removing CO_2 at a rate shown on the x-axis (in units of $GtCO_2$ /year) for horizons of 5, 30, 80, 150 and 250 years.

is fully decarbonized in 40 years; the temperature anomaly never reaches 2 C and is 1.8 C by the end of the century; and the CO_2 concentration peaks below 920 Gt CO_2 (an atmospheric concentration of 425ppm) in approximately 30 years' time (an 8% increase over 2016 levels).

The reason for this marked change in the abatement schedule is clear: by increasing the EIS, the aversion to uneven consumption is reduced. As a consequence, despite the fact that, on average, future generations are still expected to be richer than the present, the 'abatement tax' is no longer pushed as much into the future, and a greater abatement burden is accepted today. How big is this effect? The results in Figure 7, and comparisons with results obtained using an EIS of 0.69 and a discount rate $\rho = 0.10\%$ (as used in Stern (2006)) suggest that the choice of value for EIS is at least as important as the more widely discussed value for the rate of utility discounting.

All the results we present in the following section were obtained with an EIS of 0.69 (hence an RRA of 1.45), and allowing both the damage exponent and the total factor productivity to be stochastic, as described in Sections 2.3 and 2.4, respectively. As explained, we present our results for this value of EIS not because we think it is more defensible, but for ease of comparison with the DICE results. In any case, as far as carbon removal is concerned, the two choices for the EIS produced qualitatively similar results, as they both point to a large role played by negative emissions in the optimal decarbonization policy.

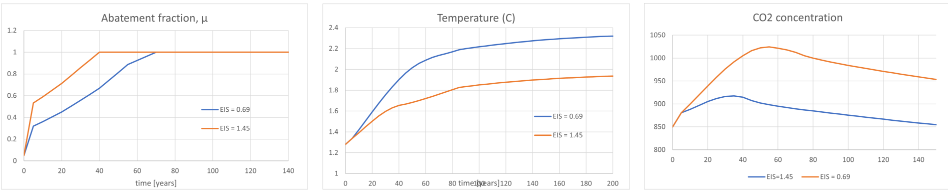

(a) Abatement fraction, μ (b) Temperature anomaly (c) CO_2 concentration

Fig. 7: The optimal abatement fraction, μ, (left panel), the optimal temperature anomaly (middle panel) and the optimal CO_2 concentration in the deterministic case (no negative emissions) for EIS = 0.69 and EIS = 1.45.

3.2. *The Impact of Negative Emissions on Optimal Policies*

We can now move to the analysis of the comparison of the optimal abatement and removal policies when negative emissions are available or not. The first observation is that, when negative emissions are available, the pace of decarbonization can be much slower (see Figure 8a) than if abatement is the only decarbonization lever, and yet the asymptotic value of the temperature anomaly is much lower (approximately 1.1C instead of close to 2.3C: see Figure 8b). This is obtained with a modest overshoot in temperature (by less than 0.1C) in 50 years' time, and with very-similar end-of-century temperatures obtained with the two policies.[n] For longer horizons, the temperature obtained with negative emissions reaches an asymptotic value similar to today's temperature anomaly, while without negative emissions it keeps on rising (albeit at a declining rate).

The rate of emission abatement necessary to obtain this asymptotically lower temperature path is much slower, and 80% decarbonization of the full economy is reached in approximately 60 years' time (instead of 45 years without negative emissions). This can greatly facilitate an orderly decarbonization of the economy, reducing transition risk. It is also important to note that the temperature pattern obtained when carbon removal is allowed is more robust to a misspecification of the much-debated rate of utility discounting: in the distant future – heavily discounted by the DICE impatience coefficient $\rho = 1.5\%$ as to be almost 'irrelevant' today – the temperature is lower with carbon removal by more than a full degree centigrade. If we think that we should care more about future generations than the DICE impatience coefficient implies, the temperature schedule associated with

[n]We note that the overshoot is relatively higher in the case of EIS = 1.45.

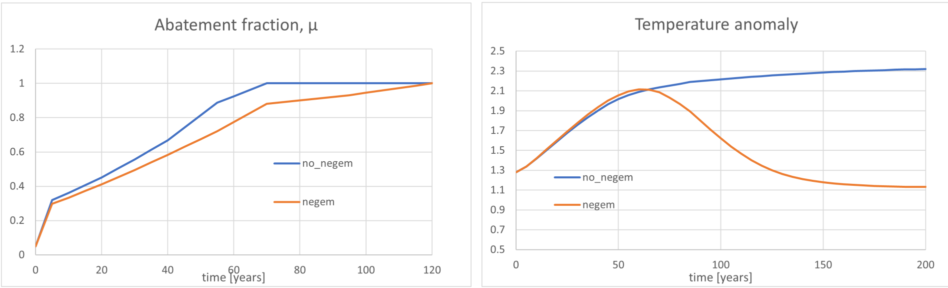

(a) Abatement fraction, $\mu(t)$. (b) Temperature anomaly, $T(t)$.

Fig. 8: The optimal abatement fraction, $\mu(t)$, (left) and the resulting temperature anomaly $T(t)$ (right). We show this with and without negative emissions (curves labelled 'negem' and 'no negem', respectively).

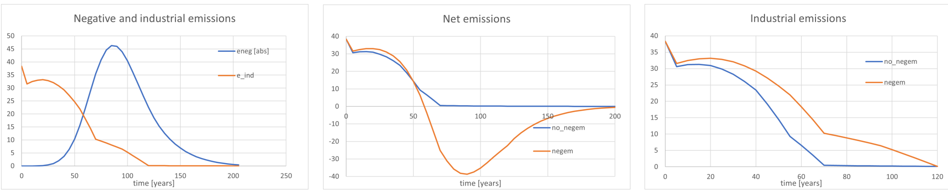

(a) Negative emissions (b) Net emissions (c) Industrial emissions

Fig. 9: Negative and industrial emissions (left panel), net emissions (middle panel), and industrial emissions with or without negative emissions.

carbon removal would certainly be more desirable: it provides, in a sense, a form of preference-misspecification insurance.

How are these results obtained ? The three panels of Figure 9 hold the key to the explanation. Note first (see Figure 9a) that, when negative emissions are available, industrial emissions can reach zero as late as in 120 years because negative emission pick up pace rapidly after the first 25 years, reaching a maximum CO_2 removal by the end of the century close to today's emissions[o]. The stark differences in the paths of industrial emissions with and without carbon removal are clearly shown in Figure 9c. The *net* emissions, shown in Figure 9b, with and without carbon removal then explain the slight temperature overshoot associated with the negative-emission policy: note, in fact, that with carbon removal emissions don't have to decline as rapidly because the policymaker knows that she will be able to remove CO_2 relatively soon.

[o]The attending logistical CO_2 storage problems, not considered in this study, but briefly discussed in the concluding section, deserve careful analysis. See National Academies of Sciences, Engineering, and Medicine 2019 (2019).

4. Conclusions and Suggestions for Future Extensions

In this study we have carried out an extensive and original review of the climate modules and of the marginal cost of abatement. When these important changes are made to the DICE model, we obtain results that can be pithily summarized by saying that that an optimal policy should target net-negative, not net-zero, emissions.

We find that an optimal solution to the global warming problem points unambiguously to the important role played by the active removal of CO_2 from the atmosphere, and to fact that such large carbon removals must entail substantial *negative* net emissions. Indeed, given projected marginal costs, the maximum amounts removed by the end of the century should be of the same order of magnitude as (and probably larger than) today's emissions. This not only confirms and strengthens the conclusions drawn in the latest IPCC report, it also adds the important element of optimality: substantial negative emissions are needed not just to reach an aspirational temperature target, but are an essential component of an economically optimal policy.

A few words of caution are in order. In a world populated by rational, time-consistent and politically unencumbered policymakers the case for a somewhat slower pace of abatement now (compensated by significant carbon removal in the near future) is compelling. In the real world, there is, however, the clear moral hazard risk that the pace of abatement will be slackened, but the necessary removal policies will then not be implemented. What one can reply to this valid objection is that the optimal cum-carbon-removal pace of abatement is not that different from the actual 'net-zero' targets that have recently been pledged. So, a pragmatic reading of our results is that the current emission targets should be maintained, but an important additional carbon-removal programme should be undertaken.

We also note that, given the current high marginal costs of removal, the optimal removal strategy takes off rather late (only picking up speed in 20-to-30 years). These marginal removal costs are, however, not a physical datum, but strongly depend on the direction towards which the scarce subsidies are channelled between abatement and removal. Experience with renewable energy shows that the best form of research is the process of 'learning by doing'. (The dramatic fall in energy cost for solar and wind installations is a testament to how effectively subsidies can be.)[p] Since current subsidies mainly target renewables and fossil fuels, the centrality of negative emissions in meeting the Paris Accord target suggests that the present

[p]This is echoed in a recent report by the National Academy of Science, that concludes that "[the US] nation should launch a substantial research initiative to advance negative emissions technologies (NETs) as soon as practicable. A substantial investment would (1) improve existing NETs [...]; (2) make rapid progress on direct air capture and carbon mineralization technologies, which are underexplored but would have essentially unlimited capacity if the high costs and many unknowns could be overcome; and (3) advance NET-enabling research on biofuels and carbon sequestration that should be undertaken anyway as part of an emissions mitigation research portfolio." See National Academies of Sciences, Engineering, and Medicine 2019 (2019), page 20.

subsidy policy could be re-examined.

An important conclusion from our study is therefore that the current allocation of subsidies, strongly skewed in favour of abatement, is far from optimal, and needs re-thinking. A second conclusion is that, if anything close to the optimal policy we obtain will indeed be followed, this will entail major sectoral shifts within the economy, with the size of the removal (and storage) sectors having to increase by orders of magnitude. To give an idea of the scale of the task, MacDowell *et al.* (2017) point out that, in order to store 25 years of current emissions 'in 2050 the [carbon sequestration and storage] industry will need to be larger by a factor of 2-4 in volume terms than the current global oil industry. In other words, we have 35 years to deploy an industry that is substantially larger than one that has developed approximately over the last century'.

Our work can be extended in several directions. First, the abatement and removal policies that we have presented are the policies to which a policymaker would commit and not modify in the future. They are, in other words, 'irrevocable optimal policies'. This way of presenting results is common (both the results from the DICE model and the Stern review are presented in this manner), but hardly satisfactory. Yes, in a world of perfect certainty the irrevocable policy is *the* best policy; however, in a world of uncertainty, one should specify state-dependent policies. We are extending our analysis in this direction.

Second, Jensen and Traeger (2014) have shown that the measure of relative prudence, which dictates whether in the presence of uncertainty one should invest in abatement or not, depends on both EIS and RRA. Recursive utility allows the decoupling of the two channels, but this comes with heavy modelling costs. Early results suggest that coupling high values of EIS and RRA make investment to control climate change even more desirable.

Another area that needs further exploration is the impact on optimal policies of the existence, severity and location of tipping points. Howard and Sterner (2017) provide estimates of damage functions that include 'catastrophic' events (presumably a place-holder for tipping-point climate cascades). Their formulation, however, retains a simple dependence of damages on temperature, while tipping points tend to display features of irreversibility or hysteresis that bring path-dependence into play. Clearly, adding the effect of tipping points would make abatement and removal schedules more 'aggressive', but precise quantification requires additional modelling work.

References

Ackerman, F., Stanton, E. A. and Bueno, R. (2010). Fat tails, exponents, extreme uncertainty: Simulating catastrophe in dice. *Ecological Economics*, **69** (8), 1657–1665.

Bansal, R. and Shaliastovich, I. (2012). A long-run risks explanation of predictability puzzles in bond and currency markets. *Review of Financial Studies*, **26** (1), 1–33.

— and Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, **59** (4), 1481–1509.

Barnett, M., Brock, W. and Hansen, L. P. (2020). Pricing Uncertainty Induced by Climate Change. *The Review of Financial Studies*, **33** (3), 1024–1066.

Belaia, M., Funke, M. and Glanemann, N. (2017). Global Warming and a Potential Tipping Point in the Atlantic Thermohaline Circulation: The Role of Risk Aversion. *Environmental and Resource Economics*, **67** (1), 93–125.

Botzen, W. J. W. and van der Bergh, J. C. J. M. (2012). How sensitive is nordhaus to weitzman? climate policy in dice with an alternative damage function. *Economics Letters*, **117**, 372–374.

Burke, M., Hsiang, S. M. and Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature*, **527** (7577), 235–239.

Cai, Y., Lenton, T. M. and Lontzek, T. S. (2016). Risk of multiple interacting tipping points should encourage rapid CO2 emission reduction. *Nature Climate Change*, **6** (5), 520–525, number: 5 Publisher: Nature Publishing Group.

Crost, B. and Traeger, C. P. (2014). Optimal CO2 Mitigation under Damage Risk Valuation. *Nature Climate Change*, **4** (7), 631–636.

Dietz, S., Hope, C., Stern, N. and Zenghelis, D. (2007). Reflections on the stern review: A robust case for strong action to reduce the risks of climate change. *World Economics*, **8** (1), 121–168.

Epstein, L. G. and Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, **57** (4), 937–969.

Folini, D., Kübler, F., Malova, A. and Scheidegger, S. (2021). The climate in climate economics.

Fuss, S., Lamb, W. F., Callaghan, M. W., Hilaire, J., Creutzig, F., Amann, T., Beringer, T., de Garcia, W. O., Hartmann, J., Khanna, T., Luderer, G., Nemet, G. F., Rogelj, J., Smith, P., Vicente, J. L. V., Wilcox, J., del Dominguez, M. M. Z. and Minx, J. C. (2018). Negative emissions—part 2: Costs, potentials and side effects. *Environmental Research Letters*, **13** (6), 063002.

Hausfather, Z., Marvel, K., Schmidt, G. A., Nielsen-Gammon, J. W. and Zelinka, M. (2022). Climate simulations: Recognize the 'hot model' problem. *Nature*, **605** (7908), 26–29.

Hope, C. (2011). The social cost of co2 from the page09 model. *Economics discussion paper*, (2011-39).

Howard, H. P. and Sterner, T. (2017). Few and not so far between: A meta-analysis of climate damage estimates. *Environmental and Resource Economics*, **68** (1), 197–225.

Howard, P. H. and Sylvan, D. (2020). Wisdom of the experts: Using survey responses to address positive and normative uncertainties in climate-economic models. *Climatic Change*, **162** (2), 213–232.

IPCC (2022). Climate change 2022: Impacts, adaptation and vulnerability. *Working Group II Contribution to the Sixth Assessment Report of the IPCC*, pp. 1–3676.

IPCCARWGIII (2021). *Climate Change 2022 – Mitigation of Climate Change*. Tech. Rep. Chapter 12, WGIII WMO UNEP.

Jensen, S. and Traeger, C. (2014). Optimal climate change mitigation under long-term growth uncertainty: Stochastic integrated assessment and analytic findings. *European Economic Review*, **69** (C), 104–125.

Joos, F., Roth, R., Fuglestvedt, J. and et al (2013). Carbon dioxide and climate impulse response functions for the computation of greenhouse gas metrics: A multi-model analysis. *Atmospheric Chemistry and Physics*, **13**, 2793–2825.

Kahn, M. E., Mohaddes, K., Ng, R. N., Pesaran, M. H., Raissi, M. and Yang, J.-C. (2021). Long-term macroeconomic effects of climate change: A cross-country analysis. *Energy Economics*, **104**, 105624.

Kalkuhl, M. and Wenz, L. (2020). The impact of climate conditions on economic production. evidence from a global panel of regions. *Journal of Environmental Economics and Management*, **103**, 102360.

Kolbert, E. and Pacala, S. (2017). *Climate Solutions: Is It Feasible to Remove Enough CO2 from the Air?* Yale environment 360, Yale School of the Environment.

Liang, Y., Gillett, N. P. and Monahan, A. H. (2020). Climate model projections of 21st century global warming constrained using the observed warming trend. *Geophysical Research Letters*, **47** (12).

MacDowell, N., Fennel, P. S., Shah, N. and Maitland, G. (2017). The role of co2 capture and utilization in mitigating climate change. *Nature Climate Chnage*, **7**, 243–249.

National Academies of Sciences, Engineering, and Medicine 2019 (2019). *Negative Emissions Technologies and Reliable Sequestration: A Research Agenda*. Washington, DC, USA: The National Academies Press.

Nordhaus, W. and Moffat, A. (2017). *A Survey of Global Impacts of Climate Change: Replication, Survey Methods, and a Statistical Analysis*. NBER Working Papers 23646, National Bureau of Economic Research, Inc.

Nordhaus, W. D. (2017). Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences of the United States of America*, **114** (7), 1518–1523.

— and Sztorc, P. (2013). *DICE 2013R: Introduction and User's Manual*. New Haven, Connecticut USA: Yale University.

OECD (2012). *OECD Employment Outlook 2012*. Organisation for Economic Co-operation and Development.

OECD (2015). *The Labour Share in G20 Economies*. Tech. rep., Organisation for Economic Co-operation and Development.

Pindyck, R. S. (2013). Climate Change Policy: What Do the Models Tell Us? *Journal of Economic Literature*, **51** (3), 860–872.

Ribes, A., Qasmi, S. and Gillett, N. P. (2021). Making climate projections conditional on historical observations. *Science Advances*, **7** (4).

Rudik, I. (2020). Optimal climate policy when damages are unknown. *American Economic Journal: Economic Policy*, **12** (2), 340–373.

Stern, N. (2006). *The Economics of Climate Change, The Stern Review*. Tech. rep., UK HM Treasury.

— (2007). *The Economics of Climate Change: The Stern Review*. Cambridge, UK: Cambridge University Press.

Takakura, J., Fujimori, S., Hanasaki, N., Hasegawa, T., Hirabayashi, Y., Honda, Y., Iizumi, T., Kumano, N., Park, C., Shen, Z., Takahashi, K., Tamura, M., Tanoue, M., Tsuchida, K., Yokoki, H., Zhou, Q., Oki, T. and Hijioka, Y. (2019). Dependence of economic impacts of climate change on anthropogenically directed pathways. *Nature Climate Change*, **9** (10), 737–741.

Tokarska, K. B., Stolpe, M. B., Sippel, S., Fischer, E. M., Smith, C. J., Lehner, F. and Knutti, R. (2020). Past warming trend constrains future warming in CMIP6 models. *Science Advances*, **6** (12).

Tol, R. (2002). Estimates of the damage costs of climate change. part 1: Benchmark estimates. *Environmental & Resource Economics*, **21** (1), 47–73.

Tol, R. S. J. (2022). A meta-analysis of the total economic impact of climate change.

Weitzman, M. L. (2010). Ghg targets as insurance against catastrophic climate damages. *Mimeo, Department of Economics, Harvard University*.

Zhu, J., Poulsen, C. J. and Otto-Bliesner, B. L. (2020). High climate sensitivity in CMIP6 model not supported by paleoclimate. *Nature Climate Change*, **10** (5), 378–379.

Appendix A. Appendix

A.1. *Parameters of the Temperature Model*

The two-layer energy balance model in DICE is parametrized by the heat capacities of the biosphere (C_{AT}), the lower oceans (C_{LO}), the heat exchange coefficient γ between these layers and the radiative feedback parameter $\lambda = \frac{F_{2\times CO2}}{T_{2\times CO2}}$ which is the ratio of the forcing from a doubling of CO_2 ($F_{2\times CO2}$) to the associated temperature change ($T_{2\times CO2}$). We tabulate the parameters for DICE-16, CMIP5 and our new calibration for CMIP6 in Table 3.

Table 3: Parametrization of the DICE two-layer EBM specifying temperature dynamics. The coefficients used in the actual model are $c_1 = \Delta/C_{AT}, c_3 = \gamma$ and $c_4 = \frac{\gamma\Delta}{C_{LO}}$, where Δ is the time-step used for the simulation. While a Δ of 5 years is appropriate for the original DICE-16 parametrization, it is necessary to use a smaller time step ($\Delta = 1$) when using the updated parametrization highlighted here.

Model	C_{AT}	C_{LO}	γ	$F_{2\times CO2}$	$T_{2\times CO2}$
DICE 16	9.9502	88.0	0.44	3.6813	3.1
CMIP5	7.3	106.0	0.73	3.45	3.25
CMIP6	6.843	121.97	0.9838	4.113	3.184

A.2. *Parameters of the Carbon Model*

The carbon model in DICE is parametrized by the equilibrium and current masses of carbon in the atmosphere, lower oceans and upper oceans, the mass transfer rates between the different reservoirs and the values of the initial temperature anomaly. We give the results in Table 4 below.

A.3. *Marginal Cost of Carbon Removal*

To develop an estimate of the marginal cost curve for carbon dioxide removal, we have used data provided by the IPCC in IPCC (2022). In particular the IPCC reviews key technologies and develops consensus estimates of the status of technological readiness (ranked from one to ten, where one is in effect theoretically proposed whereas eight to ten suggests technologies which are well understood and implemented). It also proposes a range in USD of the cost per tonne of CO_2 removed, as well as the minimum, $Q_{\min}$ and maximum, $Q_{\max}$, estimated capacities of CO_2 removal in Gigatonnes per year. We have summarised the data in Table 5.

Table 4: Calibration of the DICE-2016 carbon model and the revised calibration as per Folini *et al.* (2021). The DICE mass transfer coefficients b_{12}, b_{23}, specifying rate of carbon transfer from the atmosphere (ATM) to the upper ocean (UO) and then from the upper oceans to the lower oceans (LO) are given assuming a time step of 1 year; to convert to the similar quantity seen in Nordhaus (2017) we need to multiply by Δ, the length of a time step. To generate the other transfer rates we use the same ratios as Nordhaus (2017).

Model	Current Masses/GtC (ATM,UO,LO)	Equilibrium Masses/GtC (ATM,UO,LO)	Mass Transfer coefficients (b_{12}, b_{23})	Initial Temp Anomaly (K) (Land,Ocean)
DICE-2016	851, 460, 1740	588, 360, 1720	0.024, 0.0014	0.85, 0.0068
Revised	850, 765, 1799	607, 600, 1772	0.053, 0.0042	1.28, 0.31

Table 5: Technologies for removal of CO_2 , along with assessments of technological readiness (Status), $P_{\min}$, $P_{\max}$ are proposed minimum and maximum prices in USD for removal of one tonne of CO_2 ; also shown are $Q_{\min}$ and $Q_{\max}$, the minimum and maximum 'capacities' for CO_2 removal in $GtCO_2$ per yr.

Negative Emission Technology	Status	$P_{\min}$	$P_{\max}$	$Q_{\min}$	$Q_{\max}$
A/R	8	5	240	0.5	10.0
SCS	8	45	100	0.6	9.3
Peatland	8	5	100	0.5	2.1
Agroforestry	8	5	240	0.3	9.4
Forest Management	8	5	240	0.1	2.1
Biochar	6	10	345	0.3	6.6
DACCS	6	100	300	5.0	40.0
BECCS	5	15	400	0.5	11.0
EW	3	50	200	2.0	4.0
Ocean Alkalinity	1	40	260	1.0	100.0
Ocean Fertilisation	1	50	500	1.0	3.0

A.4. *Parameters of the Negative Emission Cost Function*

We calibrate to the removal technology marginal cost functions in Equation 2.18. The calibrated parameter values are given in Table 6.

Table 6: Parametrization of the carbon removal technology cost functions in Equation 2.18.

Coefficient	Value	Units
ν_0	3×10^{-4}	-
ν_∞	1×10^{-6}	-
ν_{decay}	0.7%	Yrs^{-1}
ξ_0	2.0	Yrs/GtCO_2
ξ_∞	0.04	Yrs/GtCO_2
ξ_{decay}	10%	Yrs^{-1}