

# Allowable carbon emissions lowered by multiple climate targets

Marco Steinacher, Fortunat Joos & Thomas F. Stocker

Corresponding author Climate and Environmental Physics, University of Bern, 3012 Bern, Switzerland Oeschger Centre for Climate Change Research, University of Bern, 3012 Bern, Switzerland Contact Marco Steinacher

Nature (2013) doi:10.1038/nature12269 Published online 03 July

Climate targets are designed to inform policies that would limit the magnitude and impacts of climate change caused by anthropogenic emissions of greenhouse gases and other substances. The target that is currently recognized by most world governments<sup>1</sup> places a limit of two degrees Celsius on the global mean warming since preindustrial times. This would require large sustained reductions in carbon dioxide emissions during the twenty-first century and beyond<sup>2, 3, 4</sup>. Such a global temperature target, however, is not sufficient to control many other quantities, such as transient sea level rise<sup>5</sup>, ocean acidification<sup>6, 7</sup> and net primary production on land<sup>8, 9</sup>. Here, using an Earth system model of intermediate complexity (EMIC) in an observation-informed Bayesian approach, we show that allowable carbon emissions are substantially reduced when multiple climate targets are set. We take into account uncertainties in physical and carbon cycle model parameters, radiative efficiencies<sup>10</sup>, climate sensitivity<sup>11</sup> and carbon cycle feedbacks<sup>12, 13</sup> along with a large set of observational constraints. Within this framework, we explore a broad range of economically feasible greenhouse gas scenarios from the integrated assessment community<sup>14, 15, 16, 17</sup> to determine the likelihood of meeting a combination of specific global and regional targets under various assumptions. For any given likelihood of meeting a set of such targets, the allowable cumulative emissions are greatly reduced from those inferred from the temperature target alone. Therefore, temperature targets alone are unable to comprehensively limit the risks from anthropogenic emissions.

The ultimate objective of the United Nations Framework Convention on Climate Change (UNFCCC) is the “stabilization of greenhouse gas concentrations in the atmosphere at a level that would prevent dangerous anthropogenic interference with the climate system”<sup>18</sup>. This goal is commonly expressed as a global mean temperature target, most notably the 2 °C temperature limit<sup>1</sup>. Yet the “climate system” within the UNFCCC refers to “the totality of the atmosphere, hydrosphere, biosphere and geosphere and their interactions”, and the broad objective specified in Article 2 of ref. 18 also covers the sustainability of ecosystems and food production. This objective thus cannot be encapsulated in one single target but may require multiple global and regional targets. Various variables essential to the habitability of Earth are discussed<sup>8, 19, 20</sup>, including climate change, sea level rise, ocean acidification, biodiversity loss, land-use change, and terrestrial net primary production (NPP). For policy-makers it is crucial to link these targets quantitatively to anthropogenic greenhouse gas emissions. Probabilistic methods<sup>21, 22</sup> can be used to account for uncertainties along the cause-and-effect chain from targets to emissions (Methods) and to provide results in terms of probability distribution functions.

For this study we define six target variables and four limits for each target that attempt to reflect levels of comparable stringency (Methods, Table 1). We stress their illustrative nature and that these choices may be refined in a dialogue with stakeholders. Two variables quantify physical changes in the climate system: the traditional global mean surface air temperature increase above preindustrial levels ( $\Delta\text{SAT}$ ) and steric sea level rise (SSLR). Two ocean acidification targets are defined in terms of area fractions. The first, ASO, is the fraction of the Southern Ocean surface area that undergoes a transition from supersaturation to undersaturation with respect to aragonitic calcium carbonate ( $\Omega_{\text{arag}} < 1$ , Methods), where sea water becomes corrosive to aragonitic shells of marine organisms<sup>6, 7, 23, 24</sup>. The second,  $A\Omega > 3$ , represents the loss of the global ocean surface area with at least threefold supersaturation ( $\Omega_{\text{arag}} > 3$ ), commonly associated with

coral reef habitats<sup>25, 26, 27</sup>. The third pair of targets addresses impacts on the terrestrial biosphere that could potentially affect food production and ecosystem services<sup>9, 28</sup>: CNPP >10% is the fraction of the global cropland area that suffers from substantial local NPP reductions (>10% relative to 2005 ad), and C<sub>carbon loss</sub> is the percentage of carbon lost from cropland soils since 2005. The response of the selected target variables and their associated uncertainties are illustrated with emission-driven simulations under the lowest (representative concentration pathway RCP2.6) and highest (RCP8.5) scenarios<sup>14</sup> used in the IPCC’s Fifth Assessment Report (Supplementary Fig. 1, Methods).

Table 1: Target variables and limits

Full table


**Table 1: Target variables and limits**

Target variable (annual mean)		Target set number				Units
		1	2	3	4	
$\Delta$ SAT	Global mean SAT increase since 1800	1.5	2	3	4	°C
SSLR	Steric sea level rise since 1800	20	40	60	80	cm
$A_{SO}$	Aragonite undersaturation of Southern Ocean surface	5	10	25	50	Percentage of area south of 50 °S
$A_{\Omega > 3}$	Global loss of surface waters with $\Omega_{arag} > 3$	60	75	90	100	Percentage of area in 1800
$C_{NPP > 10\%}$	Cropland area with NPP losses >10%	5	10	20	30	Percentage of crop area in 2005
$C_{carbon loss}$	Global soil carbon loss on croplands	5	10	20	30	Percentage of soil carbon in 2005

The targets are applied either for the time horizon of the twenty-first century or for years 2000–2300.

To quantify the allowable emissions compatible with the defined targets we ran the observationally constrained model ensemble (Methods, Supplementary Figs 2 and 3) for a set of 55 greenhouse-gas concentration pathways that represent a wide range of economically plausible scenarios<sup>14, 15, 16, 17</sup> (Fig. 1, Methods, Supplementary Table 3, Supplementary Fig. 4). We characterize the scenarios by the atmospheric CO<sub>2</sub> concentration, [CO<sub>2</sub>]<sub>2100</sub>, and the radiative forcing from non-CO<sub>2</sub> agents in the year 2100 (RFNC<sub>2100</sub>) and interpolate the target variables between the individual scenarios to sample the full two-dimensional scenario space ([CO<sub>2</sub>]<sub>2100</sub>, RFNC<sub>2100</sub>) spanned by the 55 scenarios for each of the 1,069 model configurations (Supplementary Fig. 5, Methods).

Figure 1: Flowchart illustrating the applied methodology.




First, an ensemble of model configurations is generated from prior distributions of model parameters (Supplementary Table 1, Supplementary Fig. 2). Then the ensemble is constrained by 26 observational data sets (Supplementary Table 2, Supplementary Fig. 3) by calculating a skill score ( $S_m$ ) for each ensemble member. In the next step, the constrained model ensemble is run into the future under multiple greenhouse gas scenarios (Supplementary Table 3, Supplementary Fig. 4). Finally, probability distributions of allowable CO<sub>2</sub> emissions are calculated from the simulation results for the defined targets (Table 1).

We then calculate the probabilities of not exceeding the defined limits for the scenario space (Fig. 2), considering uncertainties in physical and carbon-cycle parameters (Methods). Here we focus on target set 3 (results for all target sets are shown in Supplementary Figs 6–13).  $\Delta$ SAT and SSLR increase with both CO<sub>2</sub> and with non-CO<sub>2</sub> radiative forcing (RFNC), resulting in slanted isolines of equal probability. Depending on the concurrent non-CO<sub>2</sub> radiative forcing, CO<sub>2</sub> must not exceed 550–870 p.p.m. to be considered ‘likely’ (>66%) to stay below the  $\Delta$ SAT limit of 3 °C by 2100 (Fig. 2a). In contrast, it is extremely likely (>95%) that SSLR will not exceed 60 cm in any of the considered scenarios by 2100 (Fig. 2b). On longer timescales, however, the probability of exceeding the SSLR limits increases significantly (Supplementary Fig. 10). Ocean acidification is mainly driven by the CO<sub>2</sub> increase (vertical isolines in Figs 2c, d and Supplementary Fig. 5). It is likely that aragonite undersaturation is limited to 25% of the Southern Ocean surface by 2100 if CO<sub>2</sub> stays below 625 p.p.m. (Fig. 2c). The goal to preserve surface waters with  $\Omega_{arag} > 3$  proves harder to achieve. It is unlikely (<33%) that less than 90% of these waters are lost during this century in scenarios with [CO<sub>2</sub>]<sub>2100</sub> > 550 p.p.m. (Fig. 2d). The two cropland targets are less directly connected to [CO<sub>2</sub>]<sub>2100</sub> and RFNC<sub>2100</sub> (Supplementary Fig. 5). For CNPP >10%,

we find higher values in scenarios with very low CO<sub>2</sub> than in scenarios with higher CO<sub>2</sub> levels but relatively low RFNC. This is explained by the partially opposed effects of CO<sub>2</sub>-fertilization and climate change on NPP. Similar to SSLR, it is unlikely that the limits of set 3 are exceeded during this century for these variables (Fig. 2e, f), but the probabilities of exceeding the limits increase beyond 2100 (Supplementary Figs 1 and 12).

Figure 2: Probabilities of staying below the targets defined in set 3 up to year 2100.




(Results for all target sets are provided in the Supplementary Information.) Dark (light) brown shadings indicate low (high) probability of meeting the listed target for a given point in the scenario space defined by [CO<sub>2</sub>]<sub>2100</sub> and RFNC<sub>2100</sub>. The symbols indicate the ensemble average of the target variables (scale bars in each panel; maximum in the twenty-first century; see Supplementary Fig. 5). The representative concentration pathway (RCP, stars), Energy

Modeling Forum (EMF-21, circles), and Greenhouse Gas Initiative (GGI, diamonds) scenarios include all major anthropogenic forcings, whereas the Asia Modelling Exercise (AME, squares) scenarios are less complete and we make conservative assumptions for aerosol emissions, which results in very low RFNC.

Allowable cumulative twenty-first-century fossil-fuel CO<sub>2</sub> emissions (Eff) are diagnosed by closing the carbon budget for each concentration pathway and model ensemble member (Methods). We first examine the allowable emissions that are likely (>66%) to be compatible with the limits defined in set 3 (Fig. 3). The criterion of not exceeding the limits is applied to the time horizons 2000–2100 and 2000–2300 under the assumption of stabilizing CO<sub>2</sub> and RFNC by 2150 (Supplementary Fig. 4 and Methods). The  $A_{\Omega} > 3$  target is the most restrictive in this set, and the corresponding ensemble mean Eff are around 625 gigatons of carbon (GtC), independent of the time horizon. Up to 2100 and for moderate to high RFNC2100, the 3 °C temperature and ASO limits yield similar emissions of 750–1,200 GtC and 975 GtC, respectively. Which limit is the more restrictive depends on RFNC2100 in this case (Fig. 3a). For the very low RFNC2100 assumed in the Asia Modelling Exercise (AME) scenarios (see Methods), Eff are significantly higher for the 3 °C target than for the other targets and range up to 1,600 GtC. On the longer timescale,  $\Delta$ SAT becomes more important and approaches the  $A_{\Omega} > 3$  limit (Fig. 3b). The SSLR and Ccarbon loss limits are only relevant on the longer timescale, and CNPP >10% is insignificant for determining Eff for target set 3.

Figure 3: Allowable cumulative fossil-fuel CO<sub>2</sub> emissions for target set 3 selected for illustrative purposes.




The shading shows the ensemble average of Eff in the ([CO<sub>2</sub>]<sub>2100</sub>, RFNC<sub>2100</sub>) space given by 55 scenarios (stars for RCP, circles for EMF-21, diamonds for GGI and squares for AME). Contour lines indicate the 66% probability of not exceeding the limits of set 3 within the twenty-first century (a, compare shading in Fig. 2) and years 2000–2300 (b). The red line represents the multi-target 3, that is, the requirement of meeting all targets simultaneously, which requires smaller cumulative emissions than any of the individual targets. The aberration in Eff around RCP4.5 (at [CO<sub>2</sub>]<sub>2100</sub> = 575 p.p.m.) is due to different land-use assumptions (Methods).

A crucial question is what implications arise if we require that multiple limits must not be exceeded at the same time. Generally, Eff are lower for the combined multi-targets than for the most restrictive single limit, particularly in the long term (Fig. 3b). Therefore, if CO<sub>2</sub> were stabilized at about 500 p.p.m. by 2100, each target in set 3 would—by itself—be likely to be met, even up to 2300. Meeting all targets simultaneously, however, is less probable and is only achieved for [CO<sub>2</sub>]<sub>2100</sub> < 490 ± 20 p.p.m. when considering the 2000–2100 period, and for [CO<sub>2</sub>]<sub>2100</sub> < 460 ± 20 p.p.m. in the long term. This is related to the interdependence of target variables. If, for example, a certain model configuration simulates a weak oceanic CO<sub>2</sub> uptake and a low climate sensitivity, it is likely that surface ocean acidification is enhanced owing to the relatively high CO<sub>2</sub> in that model, whereas the temperature increase remains relatively small due to the low climate sensitivity. Hence, this specific model contributes below-average ΔSAT and above-average ocean acidification to the corresponding probability distribution functions of Eff for the two targets. Therefore, it will contribute to higher Eff for the ΔSAT target and to lower Eff for the ocean acidification target, if the probability distribution functions are evaluated independently. Likewise, another model with relatively high ΔSAT might be at the high end of the Eff probability distribution function for the ocean acidification target. If, however, the probability distribution function of Eff for meeting all targets simultaneously is considered, it is likely that the contribution from each of these individual models will be the respective lower value, that is, the Eff given by the ocean acidification and ΔSAT targets for the first and second model, respectively.

All four multi-target sets yield significantly lower Eff than the corresponding temperature target or any of the other targets in the set alone (Fig. 4, Supplementary Table 4). ΔSAT is the most limiting target only at the low end of the emission ranges and for low targets or high probabilities (Fig. 4 and Supplementary Fig. 14). For the most part, other targets (most notably OAQ > 3; Supplementary Figs 15 and 16) are more limiting and Eff inferred from the temperature target alone would be too optimistic. The implied limits on the other target variables given by the temperature targets alone are listed in Supplementary Table 5. The requirement to meet all targets simultaneously further reduces Eff considerably as explained above. For target set 3, the average Eff values at the 66% (90%) likelihood level are 40% (26%) lower for the multi-target than for the 3 °C temperature target when excluding the AME scenarios with very low RFNC<sub>2100</sub> (Fig. 4). Eff for the multi-target sets depend on the specific combination of the individual targets. When combining the temperature targets with additional targets from more (or less) stringent sets, the resulting reduction of Eff is bigger (or smaller). Nevertheless, we still find a considerable reduction for most combinations, except when combining either of the low 1.5 °C or 2 °C temperature targets with the least ambitious additional targets from set 4

Figure 4: Allowable cumulative twenty-first century fossil-fuel carbon emissions for multiple targets.



Allowable emissions are given for a likely (66%; a) and very likely (90%; b) chance of staying below the targets up to the year 2100. Results for the full probability space are provided in Supplementary Fig. 14. Blue symbols indicate the temperature-only targets and red symbols represent the corresponding multi-target sets. Green symbols show the results when considering the most limiting target with respect to the whole ensemble, but allowing the other limits to be exceeded in individual model realizations. Results are given for the case when the AME scenarios with very low RFNC are excluded (coloured bars) and also for the full scenario set (small symbols). The ranges indicate the RFNC-scenario uncertainty. No ranges are given where the targets lay outside of the examined scenario range in a significant number of model configurations, indicating that these results are upper-limit estimates (Methods). Historical emissions<sup>29</sup> and simulated emissions for the four RCP scenarios without climate targets (median, 66% and 90% confidence intervals) are shown for comparison in c.

Meeting the multi-target 1 is very unlikely (<10%) if Eff exceed  $360 \pm 40$  GtC (mean and minimum–maximum range from RFNC-scenario uncertainty; Supplementary Fig. 14), although it becomes likely to meet the 1.5 °C target (which is part of set 1) at this range of emissions (Fig. 4 and Supplementary Fig. 14). Similarly, it is unlikely that multi-target 2 can be met if Eff exceed  $470 \pm 80$  GtC, while it is still likely to meet the 2 °C target in 2100 if they stay below 470 GtC. That means that for emissions on the order of RCP2.6 it becomes likely that global warming can be limited to 2 °C, but at the same time

there is a considerable risk that at least one of the other limits of target set 2 is exceeded. To be likely to meet multi-targets 1 and 2, we estimate (Methods) that Eff must remain below 180–270 GtC and 290–350 GtC, respectively. Multi-target 3 is likely to be met for Eff below 350 GtC, which is at the high end of the emission-range for RCP2.6 (Fig. 4). Multi-target 4 is likely to be met if Eff stay below 350 GtC, a range that covers the high and low ends of RCP4.5 and RCP6.0, respectively.

Our results show that including additional targets along with the conventional global temperature limits can considerably reduce the allowable CO<sub>2</sub> emissions. In particular, ocean acidification limits pose strong constraints on CO<sub>2</sub> emissions and reduce the scenario uncertainty with respect to RFNC (Methods), which suggests that CO<sub>2</sub> targets should be treated separately from other greenhouse gases in policy frameworks. In probabilistic assessments it is not sufficient to choose only the most limiting target from a set. Instead, all targets should be taken into account simultaneously. Clearly, multiple socio-economically relevant, global and region-specific targets need to be considered in combination when the risks associated with anthropogenic emissions of CO<sub>2</sub> and other climate agents are to be assessed correctly on global to regional scales. Our results are based on ensemble simulations with an EMIC, a limited number of emission scenarios, and an illustrative set of targets. For future assessments, stakeholders should define relevant target variables and agree on limits for acceptable risks associated with environmental changes caused by anthropogenic emissions. We have shown that including additional targets would probably lead to even more stringent emission reductions than reported here. Similar studies with more comprehensive Earth system models should be carried out to include more regional and impact-related targets, such as extreme events like flooding, heat waves, or droughts.

### **Target selection**

The conventional global mean temperature increase is a straightforward metric for climate change because it comes relatively early in the causal chain from emissions to impacts, just after translating emissions to concentrations and radiative forcing. As such, this metric sometimes also stands for impacts that are associated with global warming but are more difficult to assess directly. Yet it represents other anthropogenic impacts only to a limited degree. SSLR, for example, is strongly connected to global warming but shows a delayed response owing to the relatively slow vertical mixing of heat into the ocean interior. Sea level continues to rise even after stabilization of surface temperatures<sup>5</sup>, and thus global mean temperature is not a suitable metric for the committed sea level change before equilibrium in SSLR is reached. An even more obvious example for the limited validity of global mean temperature as a metric for anthropogenic disturbance is ocean acidification from the uptake of CO<sub>2</sub>, which is a direct geochemical effect of increased atmospheric CO<sub>2</sub> concentrations and is largely independent of climate change in most regions<sup>6, 7</sup>. As a consequence it has been suggested to incorporate indicators of both climate change and ocean acidification in a common policy framework such as the UNFCCC<sup>30</sup>. Various other variables essential to the habitability of Earth have also been proposed<sup>8, 20</sup>, including biodiversity loss, land-use change and terrestrial NPP. In the light of these considerations, we define six illustrative global change target variables and four limits for each target (Table 1), which are described below.

### **Physical targets**

Two variables quantify physical changes in the climate system, that is, the traditional global mean surface air temperature increase above preindustrial (1800 ad) levels ( $\Delta\text{SAT}$ ) of 1.5–4 °C and steric sea level rise (SSLR) of 20–80 cm. We note that SSLR does not include contributions from other sources such as melting glaciers and ice sheets because this is not simulated by the EMIC applied here. SSLR is estimated to contribute about 40% of the observed total sea level rise from 1972 to 2008 with a decreasing proportion as the ice contributions increase<sup>31</sup>. We illustrate the response of the selected target variables and their associated uncertainties with emission-driven ensemble simulations under the RCP2.6 and RCP8.5 scenarios<sup>14</sup> and their extensions<sup>32</sup> to 2300 (Supplementary Fig. 1). These scenarios are the lowest and the highest of the four representative concentration pathways (RCP) defined in preparation of the IPCC's Fifth Assessment Report. Uncertainties in the response of the carbon cycle, most notably the CO<sub>2</sub> absorption of the oceans and the release of carbon from soils, introduce uncertainties in simulated atmospheric CO<sub>2</sub> concentrations that increase considerably with higher emissions and in the long term (Supplementary Fig. 1a). The uncertainties in CO<sub>2</sub> add up with the weakly constrained climate sensitivity and produce a relatively large range in  $\Delta\text{SAT}$  by 2100. Somewhat more than half of the distribution exceeds the 4 °C limit by 2100 under RCP8.5, while a small fraction projects a  $\Delta\text{SAT}$  of 2–3 °C. In the RCP2.6 scenario, more than half of the distribution exceeds 1.5 °C but not 2 °C (Supplementary Fig. 1b). SSLR shows a similar but delayed response due to the thermal inertia of the oceans. Recent estimates of  $\Delta\text{SAT}$  (ref. 11) and SSLR (ref. 5) are mostly compatible with our results but are somewhat higher for RCP8.5, particularly in the long term (Supplementary Fig. 1b, c).

### **Ocean acidification targets**

A common metric for ocean acidification is the saturation state of sea water with respect to aragonite ( $\Omega_{\text{arag}}$ ; ref. 23), a mineral form of calcium carbonate. We define two ocean acidification targets in terms of area fractions. The first, ASO, is the fraction of the Southern Ocean surface area that undergoes a transition from supersaturation to undersaturation ( $\Omega_{\text{arag}} < 1$ ; annual mean), which means that sea water becomes corrosive to aragonitic shells of marine organisms<sup>23, 24</sup>. The selected limits for this target variable range from 5% to 50%. High-latitude waters have a naturally low saturation state and thus are generally most prone to undersaturation<sup>6, 7, 33</sup>. The second ocean acidification target, OA $\Omega > 3$ , addresses areas with high saturation states ( $\Omega_{\text{arag}} > 3$ ) that are mainly found in the tropics and subtropics, and are commonly associated with coral reef habitats<sup>25, 26</sup>. Following this broad characterization, we define this variable as the percentage of the global ocean surface area with  $\Omega_{\text{arag}} > 3$  that has been lost since preindustrial times, and select limits from 60% to 100%. Many corals show a reduction in calcification rates with decreasing  $\Omega_{\text{arag}}$  over the range  $2 < \Omega_{\text{arag}} < 4$  (ref. 34), and laboratory experiments with one species have found negative net calcification for  $\Omega_{\text{arag}} < 2.8$  (ref. 35). The calcification response among species, however, is highly variable and probably depends on the interactive effects of ocean acidification and other environmental factors<sup>36, 37</sup>. Ocean acidification and warming are concurrent stressors to corals, which motivates a combination of ocean acidification and temperature targets<sup>27, 38</sup>. The simulations under the RCP8.5 and RCP2.6 scenarios illustrate that the responses of the selected surface ocean acidification variables depend mostly on CO<sub>2</sub> and the rate of ocean CO<sub>2</sub> uptake. They can be characterized as relatively fast transitions that are reversible to some extent when anthropogenic emissions remain low and CO<sub>2</sub> decreases, as is the case in RCP2.6 (Supplementary Fig. 1d, e). In RCP2.6, the Southern Ocean surface remains supersaturated in most simulations and the median  $\Omega_{\text{arag}} > 3$  area loss peaks at 60% with a considerable uncertainty. Under RCP8.5, half of the ensemble distribution projects that the entire surface of the Southern Ocean becomes undersaturated by 2100 and that virtually no surface waters with  $\Omega_{\text{arag}} > 3$  exist after 2050 and until the end of the simulation. As shown earlier<sup>39</sup>, ocean acidification changes in the deep ocean and in the surface ocean from business-as-usual carbon emissions during the twenty-first century remain irreversible on human time scales.

## Cropland targets

The third pair of targets addresses impacts on the terrestrial biosphere that could potentially affect food production and ecosystem services. The first is the fraction of the global cropland area that suffers from substantial local net primary production (NPP) reductions ( $>10\%$  relative to 2005 ad), denoted CNPP  $>10\%$ . We note that our model generally projects an increase in crop NPP on the global average for most scenarios. NPP changes, however, are spatially very heterogeneous, and our metric is chosen to capture potential negative impacts on regional food production<sup>9</sup>, although the global productivity might increase. The second terrestrial target variable is the percentage of carbon lost from cropland soils since the year 2005 (Ccarbon loss). In contrast to NPP changes, simulated changes are approximately homogeneous (in relative terms) and can be used as a global metric. Changes in soil carbon content can have large impacts on soil properties that are relevant to ecosystem functioning and crop growth<sup>28</sup>. Land that is converted from natural vegetation to cropland after 2005 is not included in these metrics. The selected limits range from 5% to 30% for both cropland targets (Table 1). The cropland targets are affected by a series of processes which introduce considerable uncertainties (Supplementary Fig. 1f, g). Changes in NPP depend on the interplay of changes in temperature, precipitation and CO<sub>2</sub> fertilization. The large climatic changes in RCP8.5 are accompanied by CO<sub>2</sub> fertilization, which explains the fact that the median area with NPP losses is smaller in RCP8.5 than in RCP2.6. Owing to the large uncertainties in RCP8.5, there is, however, still a substantial probability of high losses. The amount of carbon lost from cropland soils generally increases with higher temperatures but is also associated with considerable uncertainties (Supplementary Fig. 1g).

## Probabilistic approach

Connecting climate targets to allowable emissions is challenging because it involves several steps along the cause-and-effect chain which all include uncertainties. First, the translation of carbon emissions to atmospheric concentrations is complicated by uncertainties in the response of the carbon cycle such as the release of carbon from mineral, peat and permafrost soils in a warmer climate<sup>12</sup>, CO<sub>2</sub>-fertilization of plants<sup>40</sup>, anthropogenic land-use interactions<sup>41</sup> or the evolution of the oceanic carbon sink<sup>13</sup>. In the next step, the weakly constrained climate sensitivity<sup>11</sup> and radiative forcing from aerosols<sup>10</sup> likewise hamper the robust prediction of global temperature changes for a given atmospheric composition. Other processes further down the chain, such as agricultural productivity, typically depend on multiple environmental variables and are accordingly associated with larger uncertainties. Probabilistic methods can be used to account for these uncertainties and to provide results in terms of probability distribution functions<sup>21, 22</sup>. Here we apply our EMIC—the University of Bern three-dimensional Earth system model with Lund-Potsdam-Jena dynamic global vegetation (Bern3D-LPJ)—in a Bayesian framework to quantify allowable carbon emissions for multiple targets as depicted in Fig. 1 and described below.

## Bern3D-LPJ model parameter sampling

The Bern3D-LPJ model features a three-dimensional dynamic ocean<sup>42, 43</sup> including sea-ice<sup>44</sup>, a single-layer energy and moisture balance model of the atmosphere<sup>44, 45</sup>, and a comprehensive terrestrial biosphere component with dynamic vegetation<sup>46</sup>, permafrost, peatland<sup>47</sup> and land-use<sup>41</sup> modules (Supplementary Information). We generate a 5,000-member ensemble from the prior distributions of 19 key model parameters (Supplementary Fig. 2, Supplementary Table 1) using the Latin hypercube sampling method<sup>48</sup>. The prior distributions are selected such that the median matches the standard model configuration and the standard deviation is a quarter of the plausible parameter range based on literature and/or expert judgement (Supplementary Information). The perturbed model parameters affect terrestrial photosynthesis, hydrology, vegetation dynamics, soil organic matter decomposition and turnover, diffusivities in atmosphere and ocean, atmosphere–ocean gas transfer, the radiative forcing from greenhouse gases and aerosols, as well as the nominal climate sensitivity of the model.

## Observational constraints

To reduce uncertainties, we exploit a broad set of observation-based data to constrain the model ensemble to realizations that are compatible with observations. The data set combines information from satellite, ship-based, ice-core and in situ measurements and includes estimates of surface air temperature change, ocean heat uptake, seasonal and decadal atmospheric CO<sub>2</sub> change and ocean and land carbon uptake rates, seven physical and biogeochemical three-dimensional ocean tracer fields, as well as land carbon stocks, fluxes and fraction of absorbed radiation (Fig. 1, Supplementary Table 2, Supplementary Fig. 3). Thus, both the mean state and transient responses in space and time are probed. The model ensemble is run over the historical period (1800–2010) driven by reconstructed historical CO<sub>2</sub> emissions, the radiative forcing from additional greenhouse gases, anthropogenic and volcanic aerosols, maps of anthropogenic land cover changes, as well as changes in solar irradiance and orbital forcing. From the simulation results ('mod') and the large set of observational ('obs') constraints we assign a score to each ensemble member,  $1 \leq m \leq 5,000$ :

This likelihood-type function basically corresponds to a Gaussian distribution of the data-model discrepancy ( $\epsilon$ ) with zero mean and variance  $\sigma^2$ , which represents the combined model and observational error (Supplementary Information). The overbar indicates that the error-weighted data-model discrepancy is first averaged over all data points of each observational variable (volume- or area-weighted) and then aggregated in a hierarchical structure by averaging variables belonging to the same group (Supplementary Information, Supplementary Fig. 3, Supplementary Table 2). Cross-correlation of errors is not considered owing to computational and methodological limitations. Finally, the total score  $\sum_m \epsilon_m^2$  is normalized to one. Ensemble members with very low scores are excluded from the scenario simulations to reduce the computational cost. The reduced ensemble with 1,069 members fully represents the 5,000-member ensemble within an error of <1% (Supplementary Information).

## Greenhouse-gas scenarios

The constrained model ensemble is run for a set of 55 greenhouse-gas scenarios from the integrated assessment community. The resulting set of about 59,000 simulations permits us to quantify the allowable CO<sub>2</sub> emissions

compatible with the targets defined in this study. Thus we focus on economically feasible multi-gas emission trajectories spanning a large range from high business-as-usual pathways to low mitigation pathways that require negative CO<sub>2</sub> emissions by the end of the century (Supplementary Table 3 and Supplementary Fig. 4). These scenarios include the four RCPs (ref. 14) and 22 scenarios from the EMF-21 project<sup>15</sup>, which served as a basis for the RCP selection. In addition, the scenario set comprises 29 ‘post-RCP’ scenarios from the GGI16 of IIASA, and 23 scenarios from the AME17 (Supplementary Table 3). For these simulations, we prescribe CO<sub>2</sub> and RFNC derived from the emission scenarios (Supplementary Information). Fossil-fuel CO<sub>2</sub> emissions are translated to concentration pathways in a simulation with prescribed CO<sub>2</sub> emissions and standard model parameters. RFNC is modelled following ref. 49 with radiative efficiencies and lifetimes updated according to ref. 10. The AME scenarios, however, are less complete because they do not provide emission paths for aerosols and some minor greenhouse gases. To include these scenarios in our framework, we chose the most conservative approach by assuming constant aerosol emissions at the level of the year 2005 (radiative forcing of  $-1.17 \text{ W m}^{-2}$ ) and neglecting the forcing from the missing additional greenhouse gases, which implies a significant cooling effect continued into the future (Supplementary Fig. 4f). Following the approach of ref. 32 for RCP4.5 and RCP6.0, we extend the scenarios from 2100 to 2300 by stabilizing CO<sub>2</sub> and RFNC by 2150 (Supplementary Fig. 4).

### **Allowable emissions**

Fossil-fuel CO<sub>2</sub> emissions are diagnosed in the Bern3D-LPJ model by closing the global carbon budget for each concentration pathway and ensemble member. These emissions do not include emissions from land-use changes which are simulated internally by the model<sup>41</sup>. To derive the allowable carbon emissions for the defined targets, we first interpolate the results for each ensemble member in the two-dimensional space ( $[\text{CO}_2]_{2100}$ , RFNC<sub>2100</sub>) between the 55 scenarios using ordinary kriging<sup>50</sup>. This method is appropriate owing to the relatively simple relation between ( $[\text{CO}_2]_{2100}$ , RFNC<sub>2100</sub>) and the target variables for an individual ensemble member (Supplementary Fig. 5). Then we determine the contour lines in the interpolated fields that correspond to the defined target values. From the maximum, minimum and average emissions along these contour lines we obtain the allowable emissions (mean and RFNC-scenario uncertainty range) for each ensemble member. Finally, we calculate the probability distribution of the allowable carbon emissions from the ensemble and the weights  $S_m$  (Supplementary Information). We note that the range of considered scenarios is limited at the low end, implying that allowable emissions cannot be determined adequately for low targets and high confidence levels that require very low emissions that are hardly covered even by the most stringent mitigation scenarios included in our large set. This is the case for the multi-target sets 1 and 2 (dashed lines in Supplementary Fig. 14). In those cases only upper-limit estimates for the average allowable emissions can be given, as indicated by symbols without uncertainty ranges in Fig. 4.

### **Scenario uncertainties**

Sampling the scenario space in two dimensions, that is,  $[\text{CO}_2]_{2100}$  and RFNC<sub>2100</sub>, which varies by about  $1.6\text{--}2.9 \text{ W m}^{-2}$  for a given  $[\text{CO}_2]_{2100}$ , adds considerable scenario uncertainty to the diagnosed allowable emissions (Fig. 4). This uncertainty is generally lower for the multi-target sets than for the temperature targets because the ocean acidification metrics are largely independent of the radiative forcing. It is important to note that this uncertainty is only related to the choice of the emission scenario and neither to the parameter uncertainty of the model nor to the uncertainty of translating emissions to radiative forcing, which are both included in the probability distribution function of the allowable emissions. Another scenario uncertainty arises from the choice of the land-use scenario for the non-RCP simulations (Supplementary Information). The presented results are based on the assumption that the total land-use area increases in the non-RCP scenarios as in the RCP8.5 and RCP2.6 scenarios. If the land-use area decreases during the

twenty-first century, as assumed in RCP4.5 and RCP6.0, allowable cumulative fossil-fuel CO<sub>2</sub> emissions are 50–100 GtC higher (~5–10%, Supplementary Fig. 19).