

Long-term climate implications of twenty-first century options for carbon dioxide emission mitigation

P. Friedlingstein^{1*}, S. Solomon², G-K. Plattner³, R. Knutti⁴, P. Ciais⁵ and M. R. Raupach⁶

Long-term future warming is primarily constrained by cumulative emissions of carbon dioxide^{1–4}. Previous studies have estimated that humankind has already emitted about 50% of the total amount allowed if warming, relative to pre-industrial, is to stay below 2 °C (refs 1,2). Carbon dioxide emissions will thus need to decrease substantially in the future if this target is to be met. Here we show how links between near-term decisions, long-term behaviour and climate sensitivity uncertainties constrain options for emissions mitigation. Using a model of intermediate complexity^{5,6}, we explore the implications of non-zero long-term global emissions, combined with various near-term mitigation rates or delays in action. For a median climate sensitivity, a long-term 90% emission reduction relative to the present-day level is incompatible with a 2 °C target within the coming millennium. Zero or negative emissions can be compatible with the target if medium to high emission-reduction rates begin within the next two decades. For a high climate sensitivity, however, even negative emissions would require a global mitigation rate at least as great as the highest rate considered feasible by economic models^{7,8} to be implemented within the coming decade. Only a low climate sensitivity would allow for a longer delay in mitigation action and a more conservative mitigation rate, and would still require at least 90% phase-out of emissions thereafter.

The world is currently debating options for future greenhouse-gas emissions. Ambitions for emission reductions per country vary widely across the globe⁹. Meanwhile, global anthropogenic emissions of CO₂ are still growing, reaching about 9 Pg C yr⁻¹ since the year 2000 (ref. 10). The atmospheric CO₂ growth rate for the past decade corresponds to ~4 Pg C yr⁻¹, implying that about 45% of anthropogenic emissions stay in the atmosphere, with the rest being absorbed by the world's oceans and lands¹⁰. This atmospheric CO₂ increase, combined with increases in other greenhouse gases and aerosols, very likely led to most of the global temperature change of the past 50 years¹¹.

In this paper, we explore the consequences of three specific issues that will need to be dealt with if the world chooses to mitigate CO₂ emissions.

The first issue is how large the global emission reduction should be on human timescales to limit the global warming to 2 °C or less. On one hand, one could argue that complete decarbonization of the world economy is extremely challenging and could be beyond reach¹², so there is a risk of non-zero

emissions continuing for many decades. On the other hand, some recent mitigation scenarios include deployment of bioenergy combined with carbon-capture-and-storage (BECCS) technologies to achieve negative emissions^{13,14}. One of the four representative concentration pathways (RCP) being used for the fifth assessment of the Intergovernmental Panel on Climate Change includes BECCS leading to negative emissions by the end of the twenty-first century¹⁵. The second issue is how steep emission reduction rates should be to keep warming below 2 °C. Emission mitigation is achieved mainly through decreases in carbon intensity (use of non-fossil-fuel-based energy, such as nuclear, renewable and biomass) and increases in energy efficiency¹⁶. As mitigation costs increase with mitigation rate, the highest rate of emission reductions over the next four to five decades found in the integrated assessment model (IAM) literature is around 3.5% yr⁻¹ (refs 7,8). IAMs have also been used to investigate the risk that some technology options will fail or their potential be overestimated, making it harder to reach the mitigation target¹⁶, which would translate into lower than anticipated mitigation rates. The third issue is how soon global emissions reductions should begin. Not all countries are likely to start mitigation at the same time; countries will take on differentiated responsibilities, with some taking on more mitigation initially and others needing several decades before being able to decrease emissions effectively^{17–19}. Also, existing infrastructures commit the world to large CO₂ emissions over their lifetimes²⁰. Again, the cost of mitigation increases dramatically for shorter delays in action²¹.

Clearly, near-term rates of mitigation and potential delay will have an impact on the possibility of allowing positive emissions or requiring negative emissions in the long term if one aims to limit the long-term change in climate. Although a few studies have explored either the implications of non-zero emissions in the long term²², or the effect of different mitigation routes on the twenty-first-century-climate^{7,13,16,17}, no studies so far have highlighted how long- (coming millennium) and near-term (coming decades) decisions are intimately connected through the combination of three key factors: (1) how much, (2) how steep, and (3) how soon mitigation will take place. These three factors form an envelope of future emission choices that will jointly determine the potential for achieving climate stabilization within 2 °C above pre-industrial, or any other chosen target. The cumulative emission approach^{1–4} implicitly combines these factors and reflects to a good approximation how they will interact over the next ~100 yr. But in practice, the implementation of emissions mitigation will inevitably

¹College of Engineering, Mathematics and Physical Sciences, University of Exeter, Exeter EX4 4QF, UK, ²Department of Atmospheric and Oceanic Sciences, University of Colorado, Boulder, Colorado 80309, USA, ³Climate and Environmental Physics, Physics Institute, University of Bern, Sidlerstrasse 5, 3012 Bern, Switzerland, ⁴Institute for Atmospheric and Climate Science, ETH Zurich, Universitätstrasse 16, CH-8092 Zürich, Switzerland, ⁵Laboratoire des Sciences du Climat et de l'Environnement, UMR CEA-CNRS-UVSQ, Bat. 709, CE, L'Orme des Merisiers, 91191 Gif-sur-Yvette, France, ⁶CSIRO Marine and Atmospheric Research, Clunies Ross Street, Black Mountain, Acton, Australian Capital Territory 2601, Australia. *e-mail: p.friedlingstein@exeter.ac.uk.

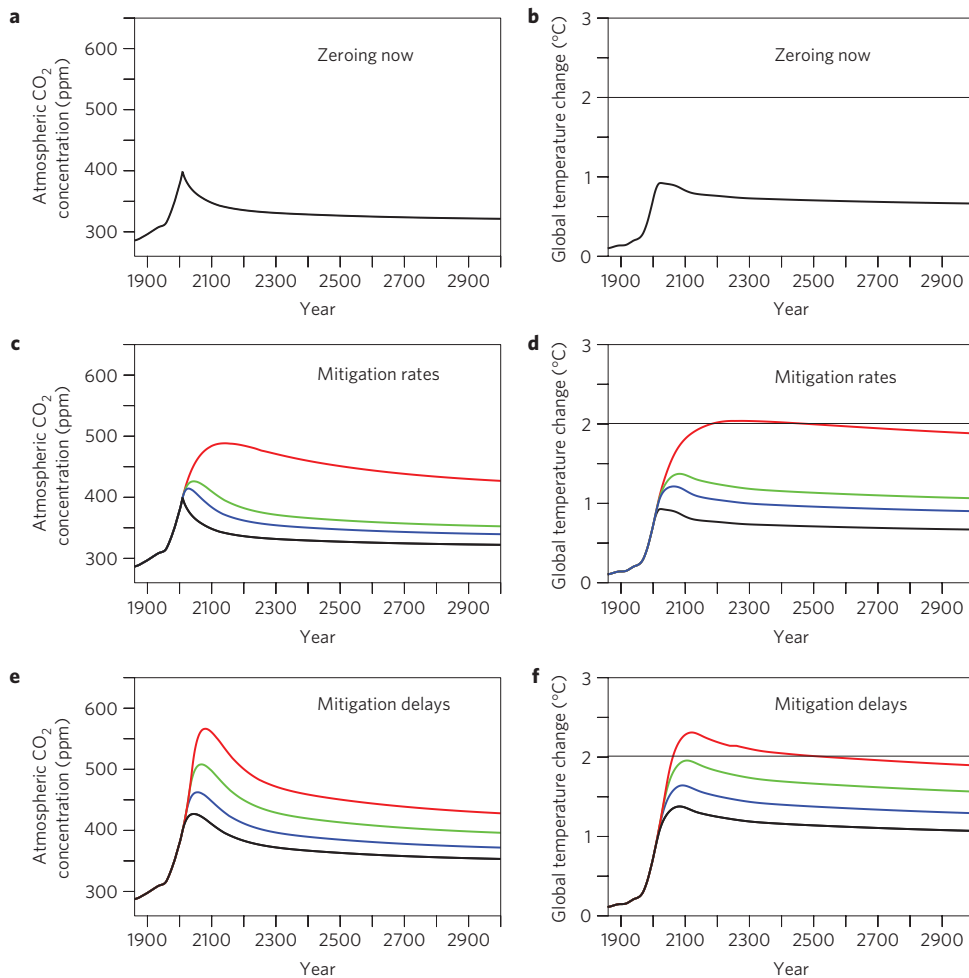


Figure 1 | Long-term effect of zero emissions. Illustrative scenarios where emissions are instantaneously cut by 100% (**a** and **b**), compared with scenarios where emissions decrease today following a mitigation rate of 5%, 3% and 1% (**c** and **d**, blue, green and red lines respectively) and where emissions start decreasing globally in 10, 20 or 30 years at a 3% yr⁻¹ mitigation rate (**e** and **f**, blue, green and red lines respectively). The left panels show atmospheric CO₂ concentration (ppmv) and the right panels show global surface temperature change relative to the pre-industrial (°C). The instantaneous mitigation case of **a,b** is reproduced on **c,d** (black lines) as the 'baseline' case (infinite mitigation rate). Similarly, the 3% mitigation rate with no delay case (**c,d**, green lines) is reproduced on **e,f** (black lines) as the 'baseline' case (no delay in mitigation). Also shown on each temperature panel is the 2 °C isotherm (thin black line).

have to deal with each of these three factors. Obviously, a more complex mitigation route could occur with, for example, time varying mitigation rates, so our results apply strictly in the context of these assumptions.

Starting from present-day anthropogenic CO₂ emissions, we explore different mitigation routes in the near-(this century) and long-term (this millennium) future and evaluate the resulting long-term consequences for the climate system. We use the BERN2.5D EMIC (refs 6–8; see Methods), versions of which have been used extensively for long-term climate projections^{6,23,24}. The climate sensitivity of the BERN2.5D model version is set to a median value of 3.0 °C for the base set of calculations presented here; the assessed uncertainty in climate sensitivity is probed below. Our study focuses on long-term climate change, so we consider CO₂ only.

Figure 1 shows the BERN2.5D model response of the carbon cycle and climate for cases where future emissions are reduced completely (100% mitigation). In the 'zeroing now' case, emission reduction happens now (no delay) and instantaneously (infinite mitigation rate). The model shows a continuous decrease of atmospheric CO₂ concentration (Fig. 1a) owing to uptake by the land and ocean. Changes in atmospheric CO₂ induce an imbalance between modelled photosynthesis and respiration over land, with respiration lagging photosynthesis owing to the long residence

time of the increasing amount of carbon in the biomass and soil carbon pools. Oceanic CO₂ exchange is similarly out of balance while atmospheric CO₂ is changing: increases in atmospheric concentrations of CO₂ and the timescales of centuries for oceanic mixing imply that the in-gassing fluxes of carbon to the ocean will be larger than the out-gassing flux. As emissions suddenly drop to zero, the CO₂ growth rate becomes negative and the strength of these land and ocean imbalances in turn declines. On multi-centennial timescales, most of the sink occurs in the ocean, whereas the land returns to almost neutral within a century. As shown in previous studies^{24,25}, the inertia in the climate response to the sudden zeroing of CO₂ emissions leads to a near-stabilization of global temperature above pre-industrial levels for many centuries (Fig. 1b).

We first evaluate the effect of less drastic near-term mitigation trajectories using cases with finite mitigation rates: a high (5% reduction per year), a middle (3% yr⁻¹) and a low (1% yr⁻¹) rate (see Methods). The highest of these rates is above what is currently considered feasible by IAMs (refs 7,8). Compared with the 'zeroing now' case, these illustrative mitigation rates obviously lead to additional emissions to the atmosphere before the ultimate reduction level is reached (180, 310 and 950 PgC, respectively), leading to a larger peak warming. With the high

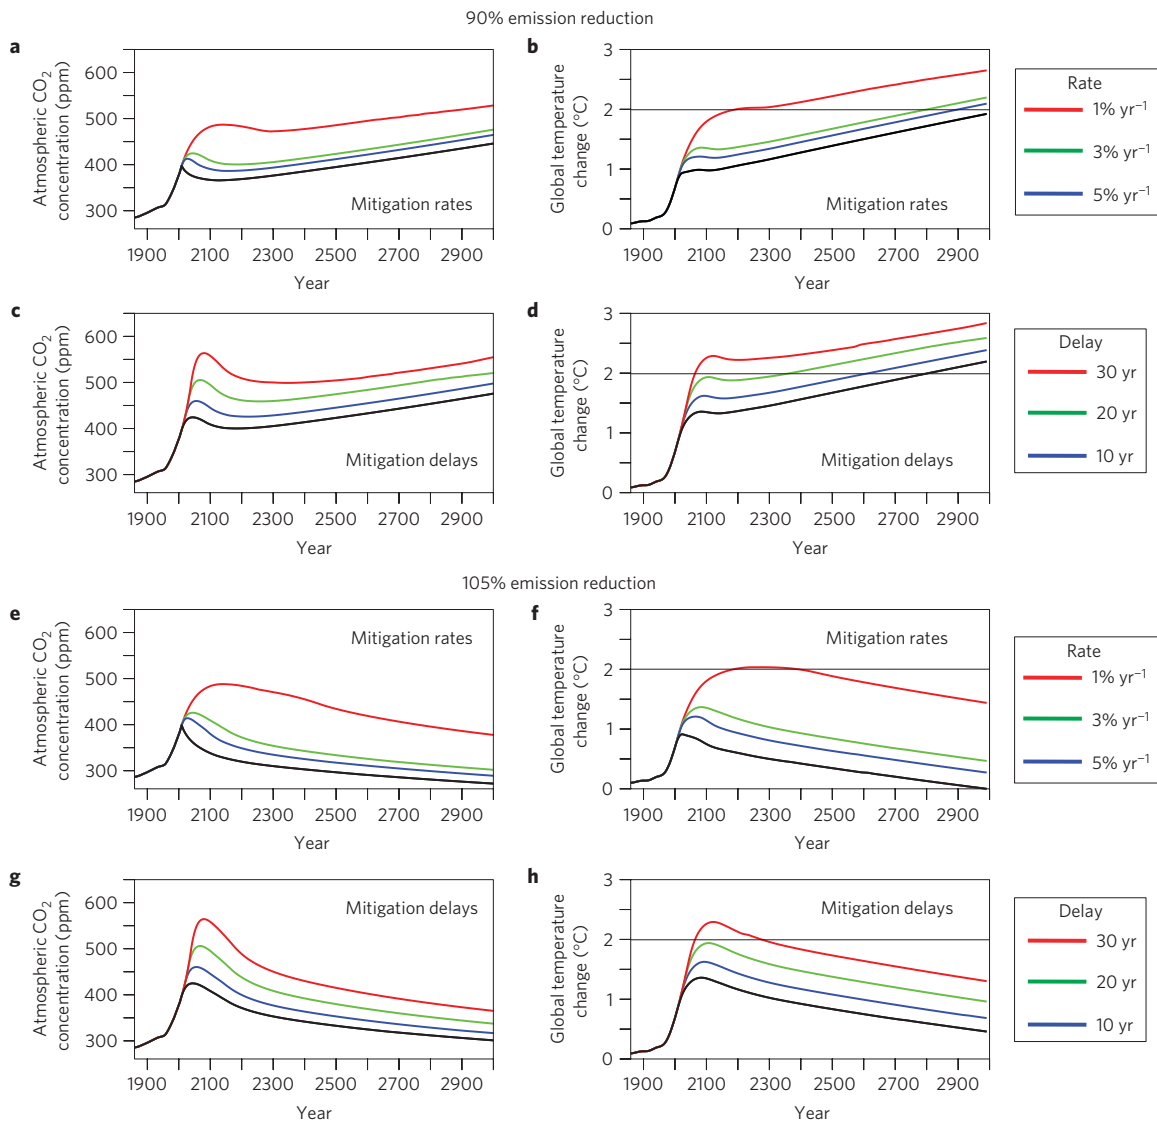


Figure 2 | Long-term effect of non-zero positive and negative long-term emissions. Illustrative scenarios where emissions eventually decrease by 90% relative to the present-day level (top four panels) or by 105% relative to the present-day level (negative emissions, through carbon capture and storage) (bottom four panels). As in Fig. 1, three mitigation rates are investigated as well as three delays in emission mitigation. The instantaneous mitigation cases are also shown on **a,b** and **e,f** (black lines) as the 'baseline' case (infinite mitigation rate). Similarly, the 3% mitigation rate with no delay cases (**a,b** and **e,f**, green lines) are shown on **c,d** and **g,h** respectively (black line) as the 'baseline' cases (no delay in mitigation). Also shown on each temperature panel is the 2°C isoline (thin black line).

mitigation rate of 5% yr⁻¹, global temperature by the year 3000 is projected to be about 0.2°C higher than in the case where emissions are cut instantaneously (Fig. 1d). For the medium and low mitigation rates towards a 100% emission cut, the projected additional warming, relative to the instantaneous reduction, reaches 0.4°C and 1.2°C respectively.

We next evaluate the long-term climate impact of a delay in the initiation of global emissions mitigation. We investigate scenarios where one, two or three decades pass before mitigation efforts begin at the same rates as above (see Methods). Figure 1e,f shows the effect of delaying mitigation for the medium mitigation rate of 3% yr⁻¹ and the long-term reduction level of 100% (relative to 2009). Even a relatively short delay of ten years has a noticeable impact on the simulated atmospheric CO₂: the additional 175 Pg C emission leads to an additional CO₂ increase of about 20 ppm to the year 3000 relative to the scenario with the same mitigation rate but without any delay. Similarly, waiting 20 or 30 years would add 390 and 650 Pg C respectively to the atmosphere. In terms of temperature

increase, waiting for one decade adds about 0.2°C, and waiting for three decades adds 0.8°C.

The above cases assume that emissions eventually drop to zero. For comparison, Fig. 2 (top 4 panels) shows a series of cases where emissions are reduced by only 90% relative to the present-day level (that is, down to 10% of the current value). As in the 100% reduction case, we explore the effect of both different mitigation rates and delays in mitigation. As the long-term emissions are positive, the near-term options for keeping global temperature rise below 2°C over the coming millennium are severely reduced. It is clear from this figure that long-term significant non-zero emission leads to a sustained increase of temperature and CO₂ concentration in the long-term (Fig. 2a,b). Land and ocean uptake will remove only a near-constant fraction of the emissions, with the remainder, the airborne fraction, accumulating in the atmosphere. CO₂ concentration and global temperature therefore increase monotonically. Even if emissions were instantaneously cut by 90% now, the long-term warming would reach 2°C (above

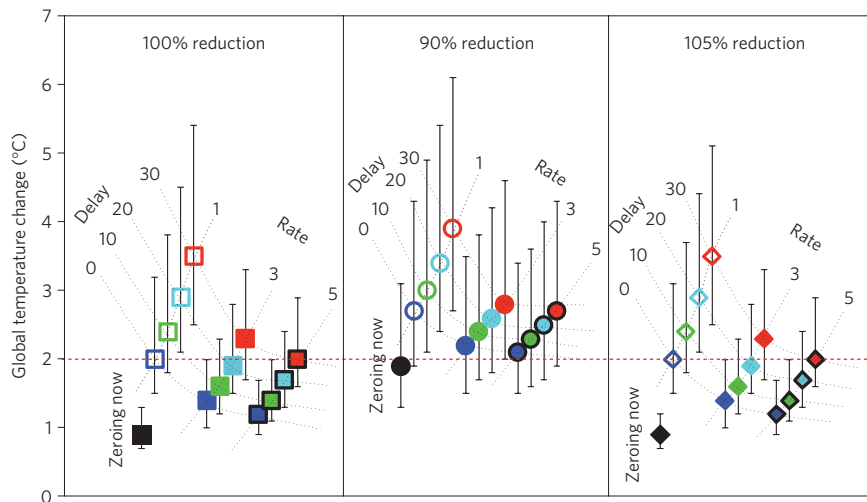


Figure 3 | Global peak temperature accounting for uncertainty in climate sensitivity. Global peak temperature change reached within the next 1,000 years for the three levels of long-term mitigation: 100% reduction (left panel, square symbols), 90% reduction (middle panel, circle symbols) and 105% reduction (right panel, diamond symbols). Symbols are for the median climate sensitivity (3 °C), while error bars shows the 'likely' range (more than 66% probability) for low (2 °C) and high (4.5 °C) climate sensitivities. Within each panel, black symbols represent immediate and instantaneous emissions reduction; dark blue, green, light blue and red symbols represent global emissions mitigation starting today or delayed by 10, 20 and 30 years respectively; mitigation rates are indicated as 1% (open symbols), 3% (filled symbols) and 5% (symbols bordered in black).

pre-industrial) by the end of the millennium (Fig. 2b, black line). All other cases with finite mitigation rates (Fig. 2a,b, coloured lines) and delays in mitigation (Fig. 2c,d, all lines) reach 2 °C before the year 3000. There is, however, a level of non-zero emission compatible with a climate stabilization, which is mainly defined by the strength of the long-term ocean carbon sink and the inertia of the ocean heat uptake. For our model, this level is below 0.1 Pg C yr⁻¹ by the year 3000, about 1% of present-day emissions.

Also shown in Fig. 2 are several cases where long-term emissions become negative (−5% of the current value), implying active removal of carbon from the atmosphere. Sustained negative emissions lead to a continuous decrease of atmospheric CO₂ (Fig. 2e), and hence of global temperature (Fig. 2f). If applied immediately (Fig. 2e,f, black line), CO₂ and climate return near to pre-industrial levels by the end of the millennium. It is clear from the figure that negative emissions provide the only option for a long-term reduction of global temperature compared with present. In the near-term, as in the previous cases, low mitigation rates or long delays will inevitably induce additional warming, by more than 2 °C in some cases. The climate overshoot persists for up to a couple of centuries, as in the 30-year delay case (Fig. 2h, red line).

Limiting future warming to less than 2 °C is hence severely constrained by the level of allowed residual emissions, the rate of emission mitigation and the time delay before these mitigation efforts begin to take place. As shown in Fig. 2, any non-zero positive residual emission will eventually lead to a sustained warming. Similarly, a finite near-term mitigation rate and a delay in action will add additional warming (Figs 1 and 2). Negative emissions will eventually reduce the warming, but weak near-term mitigation could still lead to an overshoot of the 2 °C target in this case.

However, it is critical to recognize how uncertainties in climate sensitivity and climate–carbon cycle feedbacks affect these projections. We performed further BERN2.5D simulations with a low-end (2 °C) and a high-end (4.5 °C) choice from the 'likely' (more than 66% probability) range of climate sensitivity values²⁶ (see Methods). Figure 3 shows the 'likely' range of simulated peak warming during this millennium for the 'likely' range of climate sensitivity. Whereas low climate sensitivity would allow for a longer delay or lower mitigation rate, high climate sensitivity would reduce the options available for meeting the 2 °C target.

Under a high climate sensitivity, there is no route that could lead to a warming of less than 2 °C for non-zero positive emissions. Even immediate emissions reductions as large as 5% yr⁻¹ would lead to a warming of more than 3 °C by the year 3000 for the high end of climate sensitivity. Similarly, if the real world climate sensitivity is at the high end of our range, negative emissions would lead only to a warming significantly lower than 2 °C for the high mitigation rates (5% yr⁻¹) and no delay in mitigation. Only if climate sensitivity is low can a delay of two decades followed by mitigation rates within the range of the current IAM literature meet the 2 °C target. However, this would still require at least 90% phase-out thereafter.

On millennial timescales, slow components of the Earth system could further enhance the climate response to the anthropogenic radiative forcing²⁷, but the quantitative effect of such amplification on the cases presented here is still highly uncertain (see Methods).

Previous studies have highlighted the need to limit cumulative anthropogenic CO₂ emissions below a given value to prevent a specific level of warming. About 1,000 to 1,500 Pg C have been estimated to be compatible with a 2 °C warming for a median climate sensitivity^{1,2}. Although our results, about 1,300 Pg C by 2100 (Supplementary Fig. S2), are broadly consistent with these previous findings, here we emphasize that climate sensitivity is the key parameter that determines the available mitigation options compatible with a climate target. For a given climate sensitivity, options for near-term mitigation rates and the onset of mitigation action are severely constrained by the level of long-term emission reduction below current values. Negative emissions might allow a low climate target to be reached in the long term, but the peak level of warming will be largely driven by the near-term parameters.

Methods

Bern2.5D model. The Bern2.5D model is a coupled climate–carbon cycle model of intermediate complexity^{5,6}. The atmosphere is a one-layer energy-balance model coupled to a zonally averaged dynamic ocean that consists of three ocean basins, connected in the south by a circumpolar Southern Ocean. The model includes a representation of the marine carbon cycle²⁸ and a simple four-box terrestrial carbon cycle²⁹. The global ocean heat uptake is near the average of existing comprehensive Atmosphere–Ocean General Circulation Models²³. For the land biosphere, a fertilization by elevated atmospheric CO₂ is taken into account by a logarithmic dependence of net primary production on CO₂. The BERN2.5D model does not consider interaction with sediments. This process becomes significant for ocean carbon uptake on timescales longer (multi-millennia) than the ones investigated here³⁰.

Simulation set-up. The BERN 2.5D model was first brought to pre-industrial (1765 AD) equilibrium with a 20,000 years spin-up under a prescribed atmospheric CO₂ concentration (278 ppmv). Then the model was run through the pre-instrumental atmospheric CO₂ measurements period (1765–1958) with imposed CO₂ concentration taken from ice-core data. For the 1959–2009 period, the model was forced with anthropogenic emissions (fossil fuel and land use) using global emissions from the Global Carbon Project (<http://lglmacweb.env.uea.ac.uk/lequere/co2>).

For the 2010–3000 period, the model was forced with anthropogenic CO₂ emissions following 39 different cases (Supplementary Fig. S1): three idealized cases with instantaneous emission reductions, where CO₂ emissions were set constant with a value equal to either 90, 100 or 105% reduction relative to the 2009 value; nine cases with illustrative mitigation rates, where starting from the 2009 value, CO₂ emissions were decreased by 1%, 3% or 5% yr⁻¹ until they reached either the 90%, 100% or 105% reduction level; and 27 cases with illustrative delays in emissions where, starting from the 2009 value, CO₂ emissions were first assumed to increase for 10, 20 or 30 years with a positive growth rate of 2% yr⁻¹ (similar to the growth rate of the past decade), then CO₂ emissions were decreased by 1, 3 or 5% until they reached either the 90%, 100% or 105% reduction level.

The 90% emission reduction leads to a long-term emission level of 0.9 Pg C yr⁻¹. The 105% emission reduction represents a net global negative emission of -0.45 Pg C yr⁻¹, comparable to what is proposed in the RCP2.6 scenario¹⁸. To achieve negative emissions we switched from the exponential decrease to a constant decrease when emissions reached 10% of 2009 emissions.

Uncertainty from climate sensitivity and climate-carbon cycle feedback. The BERN-2.5D model does not represent the physical feedbacks responsible for the uncertainty in climate sensitivity. However, the model has a feedback parameter that sets the model's climate response to a given CO₂ radiative forcing. In the standard simulations listed above, this parameter was set to give a median climate sensitivity of 3 °C. For the further simulations addressing the uncertainty on climate sensitivity, this parameter was set to values giving a low climate sensitivity of 2 °C or a high climate sensitivity of 4.5 °C (ref. 26). All 39 simulations presented above were performed with these low and high climate sensitivity values. We note that this range of climate sensitivity does not include Earth system processes such as vegetation or ice-sheet dynamics²⁷. Accounting for an 'Earth-system sensitivity' rather than a standard 'climate sensitivity' could increase the warming by 30–50% (ref. 27). This would lead virtually all illustrative cases presented here to exceed the 2 °C target. It is however unclear how much of this additional sensitivity could operate within the coming millennium, under relatively low forcing scenarios, as in our study.

The version of the BERN2.5D model used here has a positive climate-carbon cycle feedback, but is at the lower end of the published range of the coupled climate-carbon cycle models intercomparison project (C⁴MIP; ref. 31). Larger positive climate-carbon cycle feedback would lead to larger warming for a given emission scenario. The mitigation cases explored here lead to a relatively low forcing, and do not probe the range of possible magnitudes of the climate-carbon cycle feedback. Within the C⁴MIP study, 11 models investigated this feedback in the context of a high emission scenario (SRES A2) only, with atmospheric CO₂ reaching 800–1,000 ppm by 2100, and showed the gain to increase with the warming. Also, it has been shown that the climate-carbon cycle feedback depends on the rate of forcing³². We therefore felt that it might not be feasible to define a generic likely range of climate-carbon cycle feedbacks. It seems nevertheless improbable for this feedback to be very large here; hence the contribution of the climate-carbon cycle feedback uncertainty to the total uncertainty in our projections is expected to be small³³.

Received 20 May 2011; accepted 31 October 2011; published online 20 November 2011

References

- Allen, M. *et al.* Warming caused by cumulative carbon emissions towards the trillionth tonne. *Nature* **458**, 1163–1166 (2009).
- Zickfeld, K., Eby, M., Matthews, H. D. & Weaver, A. J. Setting cumulative emissions targets to reduce the risk of dangerous climate change. *Proc. Natl Acad. Sci. USA* **106**, 16129–16134 (2010).
- Matthews, H. D., Gillett, N. P., Stott, P. A. & Zickfeld, K. The proportionality of global warming to cumulative carbon emissions. *Nature* **459**, 829–832 (2009).
- Raupach, M. R. *et al.* The relationship between peak warming and cumulative CO₂ emissions, and its use to quantify vulnerabilities in the carbon-climate-human system. *Tellus B* **63**, 145–164 (2011).
- Stocker, T. F., Wright, D. G. & Mysak, L. A. A zonally averaged, coupled ocean-atmosphere model for paleoclimate studies. *J. Clim.* **5**, 773–797 (1992).
- Joos, F., Plattner, G.-K., Stocker, T. F., Marchal, O. & Schmittner, A. Global warming and marine carbon cycle feedbacks on future atmospheric CO₂. *Science* **284**, 464–467 (1999).
- UNEP *The Emissions Gap Report: Are the Copenhagen Accord Pledges Sufficient to Limit Global Warming to 2 °C or 1.5 °C?* (UNEP, 2010); available at <http://www.unep.org/publications/ebooks/emissionsgapreport/>.
- den Elzen, M. G. J., van Vuuren, D. P. & van Vliet, J. Postponing emission reductions from 2020 to 2030 increases climate risks and long-term costs. *Climatic Change* **99**, 313–320 (2010).
- UNFCCC, *Communications Received from Parties in Relation to the Listing in the Chapeau of the Copenhagen Accord* (UNFCCC, 2010); available at http://unfccc.int/meetings/cop_15/copenhagen_accord/items/5276.php.
- Friedlingstein, P. *et al.* Update on CO₂ emissions. *Nature Geosci.* **3**, 811–812 (2010).
- Hegerl, G. C. *et al.* in *IPCC Climate Change 2007: The Physical Science Basis* (eds Solomon, S. *et al.*) 663–745 (Cambridge Univ. Press, 2007).
- Hourcade, J.-C. & Crassous, R. Low-carbon societies: A challenging transition for an attractive future. *Clim. Policy* **8**, 607–612 (2008).
- Van Vuuren, D. P. *et al.* Stabilizing greenhouse gas concentrations at low levels: An assessment of reduction strategies and costs. *Climatic Change* **81**, 119–159 (2007).
- Van Vuuren, D. P., Bellevrat, E., Kitous, A. & Isaac, M. Bio-energy use and low stabilization scenarios. *Energy J.* **31**, 193–222 (2010).
- Van Vuuren, D. P. *et al.* RCP2.6: Exploring the possibility to keep global mean temperature change below 2 °C. *Climatic Change* **109**, 95–116 (2011).
- Edenhofer, O. *et al.* The economics of low stabilization: Model comparison of mitigation strategies and costs. *Energy J.* **31**, 11–48 (2010).
- Clarke, L. *et al.* International climate policy architectures: Overview of the EMF 22 international scenarios. *Energy Econ.* **31**, S64–S81 (2009).
- Edmonds, J., Clarke, L., Lurz, J. & Wise, M. Stabilizing CO₂ concentrations with incomplete international cooperation. *Clim. Policy* **8**, 355–376 (2008).
- Richels, R., Blanford, G. J. & Rutherford, T. F. International climate policy: A 'second best' solution for a 'second best' world? *Climatic Change* **97**, 289–296 (2009).
- Davis, S. J., Caldeira, K. & Matthews, H. D. Future CO₂ emissions and climate change from existing energy infrastructure. *Science* **329**, 1330–1333 (2010).
- Bosetti, V., Carraro, C. & Tavoni, M. Climate change mitigation strategies in fast-growing countries: The benefits of early action. *Energy Econ.* **31**, S144–S151 (2009).
- Weaver, A. J., Zickfeld, K., Montenegro, A. & Eby, M. Long term climate implications of 2050 emission reduction targets. *Geophys. Res. Lett.* **34**, L19703 (2007).
- Meehl, G. A. *et al.* in *IPCC Climate Change 2007: The Physical Science Basis* (eds Solomon, S. *et al.*) 747–845 (Cambridge Univ. Press, 2007).
- Solomon, S., Plattner, G.-K., Knutti, R. & Friedlingstein, P. Irreversible climate change due to carbon dioxide emissions. *Proc. Natl Acad. Sci. USA* **106**, 1074–1079 (2009).
- Solomon, *et al.* Persistence of climate changes due to a range of greenhouse gases. *Proc. Natl Acad. Sci. USA* **107**, 18354–18359 (2010).
- Knutti, R. & Hegerl, G. C. The equilibrium sensitivity of the Earth's temperature to radiation changes. *Nature Geosci.* **1**, 735–743 (2008).
- Lunt, D. J. *et al.* Earth system sensitivity inferred from Pliocene modeling and data. *Nature Geosci.* **3**, 60–64 (2010).
- Marchal, O., Stocker, T. F. & Joos, F. A latitude-depth circulation-biogeochemical ocean model for paleoclimate studies. Model development and sensitivity. *Tellus B* **50**, 290–316 (1998).
- Siegenthaler, U. & Oeschger, H. Biospheric CO₂ emissions during the past 200 years reconstructed by convolution of ice core data. *Tellus B* **39**, 140–154 (1987).
- Archer, D. Fate of fossil fuel CO₂ in geologic time. *J. Geophys. Res.* **110**, C09S05 (2005).
- Friedlingstein, P. *et al.* Climate-carbon cycle feedback analysis: Results from the C⁴MIP model intercomparison. *J. Clim.* **19**, 3337–3353 (2006).
- Gregory, J. M., Jones, C. D., Cadule, P. & Friedlingstein, P. Quantifying carbon cycle feedbacks. *J. Clim.* **22**, 5232–5250 (2009).
- Knutti, R., Stocker, T. F., Joos, F. & Plattner, G.-K. Probabilistic climate change projections using neural networks. *Clim. Dyn.* **21**, 257–272 (2003).

Acknowledgements

The authors wish to thank D. Matthews, J. Daniel and T. Sanford for helpful discussions and R. Flecker for creating Fig. 3.

Author contributions

P.F. and S.S. designed the work and the experiments, P.F. performed the model simulations, G.-K.P. provided the model code, G.-K.P. and R.K. gave guidance on the use of the model, P.F. led the writing of the paper with contributions from all other authors.

Additional information

The authors declare no competing financial interests. Supplementary information accompanies this paper on www.nature.com/natureclimatechange. Reprints and permissions information is available online at <http://www.nature.com/reprints>. Correspondence and requests for materials should be addressed to P.F.