

Risk of multiple interacting tipping points should encourage rapid CO₂ emission reduction

Yongyang Cai^{1,2†}, Timothy M. Lenton^{3*†} and Thomas S. Lontzek^{4*†}

Evidence suggests that several elements of the climate system could be tipped into a different state by global warming, causing irreversible economic damages. To address their policy implications, we incorporated five interacting climate tipping points into a stochastic-dynamic integrated assessment model, calibrating their likelihoods and interactions on results from an existing expert elicitation. Here we show that combining realistic assumptions about policymakers' preferences under uncertainty, with the prospect of multiple future interacting climate tipping points, increases the present social cost of carbon in the model nearly eightfold from US\$15 per tCO₂ to US\$116 per tCO₂. Furthermore, passing some tipping points increases the likelihood of other tipping points occurring to such an extent that it abruptly increases the social cost of carbon. The corresponding optimal policy involves an immediate, massive effort to control CO₂ emissions, which are stopped by mid-century, leading to climate stabilization at <1.5 °C above pre-industrial levels.

The social cost of carbon (SCC) represents the cost of all future climate damages stemming from a marginal emission of CO₂, discounted to the year of emission. The 2010 US Federal assessment¹ used three simple integrated assessment models (IAMs) to arrive at a SCC of US\$21 per tCO₂ for a tonne emitted in 2010, which was subsequently revised upwards² to US\$33 per tCO₂. Several other studies^{3–6} have argued for a higher SCC on various grounds. A key potential contributor to increasing the SCC is the possibility that ongoing climate change will cause elements of the climate system to pass 'tipping points' leading to irreversible damages^{7,8}.

Existing scientific studies suggest there are multiple climate tipping points that could be triggered this century or next if climate change continues unabated^{7,8}, and there are causal interactions between tipping events such that tipping one element affects the likelihoods of tipping others⁸ (Fig. 1). The likelihood of specific tipping events varies, but is generally expected to increase with global temperature^{7,8}. However, internal variability within the climate system, and relatively rapid anthropogenic forcing, mean that even if deterministic tipping points could be precisely identified, the actual systems could be tipped earlier or later⁹. Thus, any assessment of their policy implications needs to represent the stochastic uncertainty surrounding when tipping points could occur¹⁰. Furthermore, the impacts of passing different tipping points are expected to vary^{7,11}, and to unfold at different rates depending on the internal timescale of the part of the climate system being tipped^{7,11}.

Relative to this scientific understanding, most cost–benefit analyses of climate change allow for only simple and scientifically unrealistic representations of climate tipping points¹¹. Most previous IAM studies of climate catastrophes have treated them in a deterministic fashion, sometimes giving them a probability distribution^{5,12–15}. Some recent IAM studies have considered one stochastic climate tipping point impacting economic output¹⁰, non-market welfare¹⁶, climate sensitivity¹⁷, or carbon cycle feedbacks¹⁷. This can lead to up to 200% increases in the SCC in extreme cases¹⁰, with the results clearly sensitive to the timescale over which

tipping point impacts unfold, as well as the final magnitude of those impacts¹⁰. However, there has been little consideration of multiple tipping points and interactions between them, or of how an appropriate representation of risk aversion affects the optimal response to the prospect of future tipping points.

A recent IAM study¹⁸ has examined three loosely defined tipping points that instantaneously alter climate sensitivity, carbon cycle feedbacks, or economic output, and interact through their effects on atmospheric CO₂, global temperature, or economic output. Here we consider five carefully defined tipping points^{7,8} and the direct causal interactions between them identified by scientific experts⁸ (Fig. 1). These interactions occur primarily through aspects of the climate system that are not resolved in simple IAMs. The impacts of our tipping points unfold at a rate appropriate for the system being tipped, in contrast with instantaneous changes^{17,18} in climate sensitivity and carbon cycle feedbacks, which are scientifically questionable¹⁰. Our tipping points principally affect economic output, although we also consider their feedback effects on the carbon cycle. Instead of arbitrarily specifying the likelihood of the tipping points¹⁸, we calibrate their likelihoods (and the causal interactions between them) on the basis of the results of an existing expert elicitation⁸. Furthermore, in contrast to recent work¹⁸, we alter the specification of the social planner's preferences regarding risk aversion and intergenerational equity, in a manner appropriate for the stochastic uncertainty surrounding future tipping points.

Modelling tipping points

We use the dynamic stochastic integration of climate and economy (DSICE) framework¹⁹ to incorporate five stochastic tipping points and causal interactions between them into the 2013 version of the well-known DICE model²⁰ (see Methods and Supplementary Figs 1 and 2). This means solving a sixteen-dimensional stochastic model—the first time in the field of economics of climate change that an analysis on such a scale has been accomplished (our previous work¹⁰ solved a seven-dimensional system, whereas other simplified stochastic versions¹⁷ of DICE consider only four dimensions). In our stochastic version of the DICE model, we use annual time steps, and

¹Hoover Institution, Stanford University, Stanford, California 94305, USA. ²Becker Friedman Institute, University of Chicago, Chicago, Illinois 60636, USA. ³Earth System Science, College of Life and Environmental Sciences, University of Exeter, Exeter EX4 4QE, UK. ⁴Department of Business Administration, University of Zurich, Zürich CH 8008, Switzerland. †These authors contributed equally to this work. *e-mail: t.m.lenton@exeter.ac.uk; lontzek@gmail.com

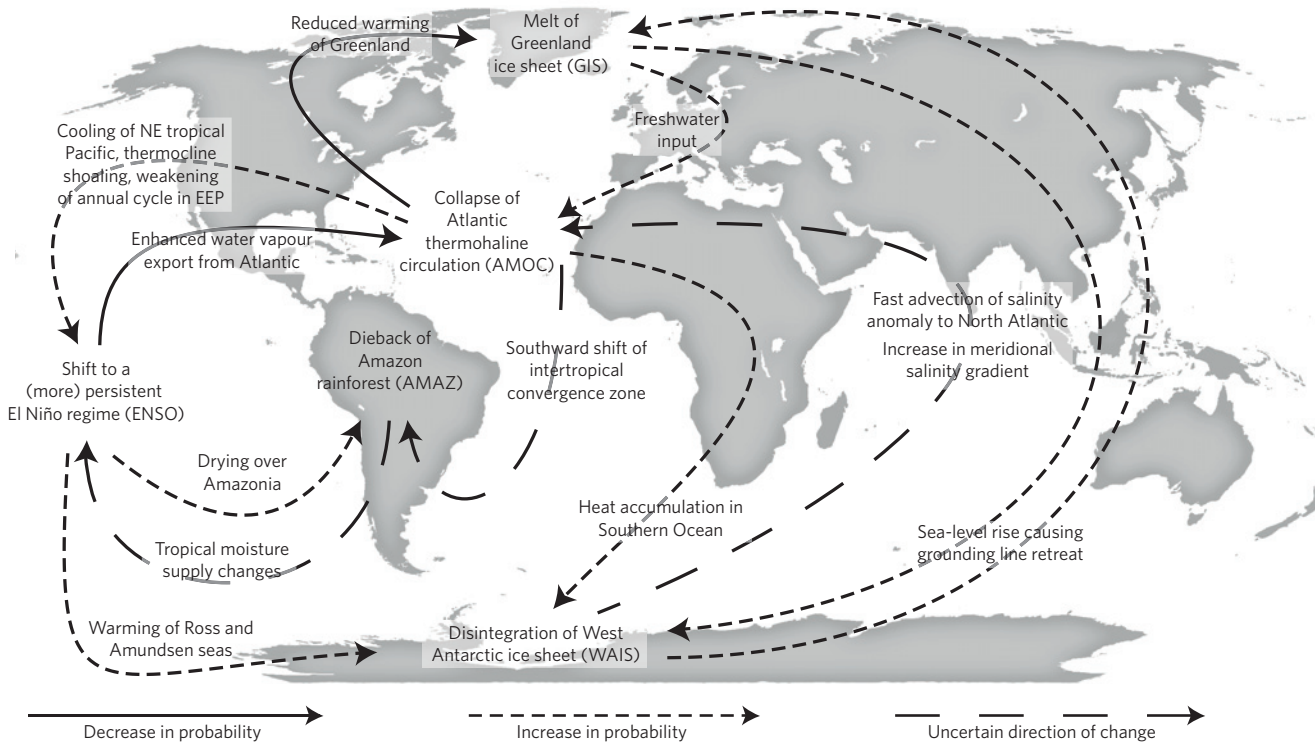


Figure 1 | Map of the five climate tipping events considered here and the causal interactions between them previously identified in an expert elicitation⁸.

calibrate parameters in the carbon cycle and temperature modules against the emulated median response of complex climate models for the four RCP (representative concentration pathway) scenarios²¹ (see Supplementary Methods). In a deterministic setting within our model (without considering climate tipping points) our calibration gives a SCC in 2010 of US\$15 per tCO₂ (all results are in 2010 US dollars). For reference, the DICE-2013R model²⁰ that uses five-year time steps and is calibrated against one RCP scenario also has a 2010 SCC of US\$15 per tCO₂.

In IAMs such as DICE, greater emission control at present mitigates damages from climate change in the future but limits consumption and/or capital investment today. A ‘social planner’ is assumed to weigh these costs and benefits of emission control to maximize the expected present value of global social welfare. When faced with stochastic uncertainty about future tipping events, the social planner’s response will depend on their preferences regarding risk and smoothing consumption. DICE adopts a specification of risk aversion that is inversely tied to the decision-maker’s preferences to smooth consumption over time (that is, the inter-temporal elasticity of substitution). Thus, a high inter-temporal elasticity of substitution is taken to imply a low risk aversion. In the baseline DICE model, risk aversion $RA = 1.45$, and inter-temporal elasticity of substitution $IES = 1/1.45$. However, empirical economic data do not support this inverse proportionality (implying time-separable utility) and suggest instead decoupling these preferences²². Hence, we incorporated ‘Epstein–Zin’ (EZ) preferences²² using default parameter settings²³ of $RA = 3.066$ and $IES = 1.5$, which are consistent with empirical findings²³ (implying non-time-separable utility). Estimates of $IES > 1$ have been obtained from, for example, stockholder data²⁴, $IES = 1.5$ is used in a long-run risk model^{19,25}, and the upper bound is considered²³ to be $IES \sim 2$. Using $IES = 1.5$, equity returns data²³ suggest $RA = 3.066$, which is in the range $RA = 3–4$ from a separate study of equity premiums of rare disasters²⁶, with the upper bound considered²⁵ to be $RA \sim 10$.

The five interacting, stochastic, potential climate tipping points^{7,8} (Fig. 1 and Table 1) represent reorganization of the Atlantic

meridional overturning circulation (AMOC), disintegration of the Greenland ice sheet (GIS), collapse of the West Antarctic ice sheet (WAIS), dieback of the Amazon rainforest (AMAZ), and shift to a more persistent El Niño regime (ENSO). We used published expert elicitation results⁸ to derive the likelihoods (see Methods) of each of the five tipping events (Table 1), and the causal interactions between them (Fig. 1 and Supplementary Table 1). By causal interaction, we mean that the hazard rate of each tipping point depends on the state of the others.

For each tipping event, we specified a transition timescale¹⁰ (Table 1, see Methods)—that is, how long it would take for the full impacts to unfold, based on current scientific understanding of the timescales of the systems being tipped^{7,11} (for example, ice sheets melt more slowly than the ocean circulation can reorganize). Recognizing the scientific uncertainty surrounding transition times, we explore a factor of 5 uncertainty range in either direction. We must also specify a final damage for each tipping event (Table 1, see Methods), taken to be an irreversible percentage reduction in world GDP (gross domestic product). This is the most problematic and debatable part of the parameterization, because of a gross shortage of scientific and economic estimates of tipping point damages¹¹. We can make some scientific inferences about relative damages (for example, based on the eventual contributions of different ice sheets to sea-level rise). Past studies with DICE have loosely associated a 25–30% reduction in GDP comparable with the Great Depression with a collapse of the AMOC^{27,28}, but when combined with other tipping points this could lead to excessively high overall damages. Our assigned damages for individual tipping points range from 5 to 15% reduction in GDP with a combined reduction in GDP if all five tipping events occur and complete their transitions of 38%. However, owing to relatively low probabilities and long transition timescales, the expected tipping point damages in our default scenario amount to only 0.53% of GDP in 2100 and 1.89% of GDP in 2200. In our sensitivity analysis, we consider a factor of 2–3 total uncertainty range in final damages for each tipping point. Finally, we include some conservative

Table 1 | Hazard rate, transition time, final damages and carbon cycle effect for each tipping element, with uncertainty ranges (in parentheses) considered in the sensitivity analysis.

Tipping element	Hazard rate (% yr ⁻¹ K ⁻¹)	Transition time (yr)	Final damages (% of GDP)	Carbon cycle effect
AMOC	0.063	50 (10–250)	15 (10–20)	No effect
GIS	0.188	1,500 (300–7,500)	10 (5–15)	100 GtC over transition
WAIS	0.104	500 (100–2,500)	5 (2.5–7.5)	100 GtC over transition
AMAZ	0.163	50 (10–250)	5 (2.5–7.5)	50 GtC over transition
ENSO	0.053	50 (10–250)	10 (5–15)	0.2 GtC yr ⁻¹ permanent

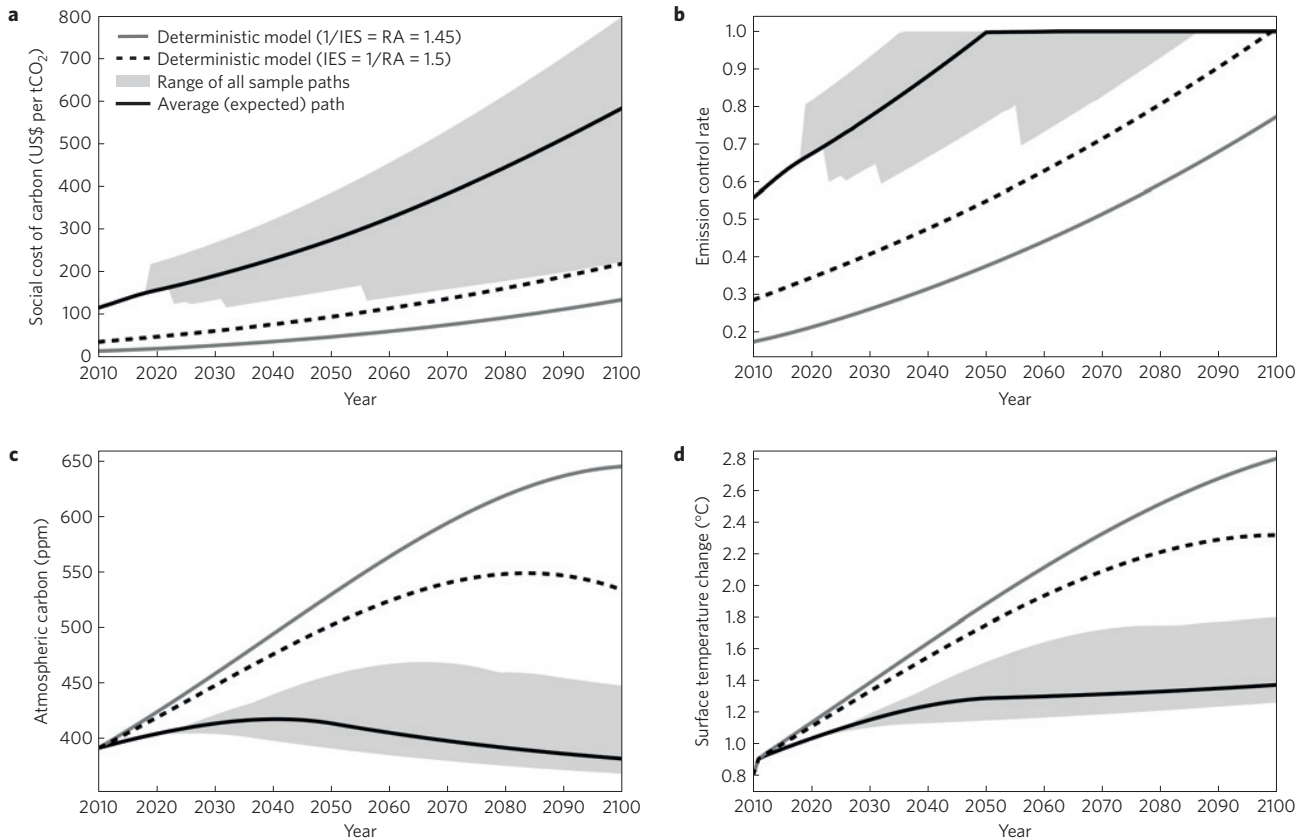


Figure 2 | Effect of multiple interacting tipping points and altered preferences on optimal policy. **a–d**, Results for the SCC (**a**), emissions control policy (**b**), atmospheric carbon (**c**) and surface temperature change (above pre-industrial levels) (**d**) in the baseline deterministic model, the deterministic model with $IES = 1.5$ and the expected path of the stochastic model with multiple interacting tipping points and EZ preferences ($IES = 1.5$, $RA = 3.066$). The grey shaded area shows the range of sample paths from 10,000 simulations of the stochastic model (see Supplementary Fig. 3 for the analogous case without interaction).

effects of tipping particular systems on the carbon cycle (Table 1, see Methods).

Optimal policy

The result of including multiple interacting tipping points under the default EZ preferences (Fig. 2) is a nearly eightfold increase in the initial SCC from US\$15 per tCO₂ in the baseline model (grey line) to US\$116 per tCO₂ (black line). Across 10,000 sample paths of the model, there are cases where one or more tipping points still occur, leading to uncertainty ranges for the key variables (grey-shaded areas). The emissions control rate jumps from ~18 to ~56% in 2010 and rises to 100% by 2050, effectively shutting down fossil fuel CO₂ emissions—whereas in the baseline model, emissions continue into the next the century. The average atmospheric carbon peaks in the 2030s at 415 ppm and then declines (owing to ongoing ocean carbon uptake)—whereas in the baseline model, atmospheric CO₂ continues to rise to ~650 ppm by 2100. Temperature rise slows down and is almost stable around 1.4 °C above pre-industrial

by 2100—whereas in the baseline model, warming continues and approaches 3 °C by 2100. Following the expected path (black line) there is only an 11% probability of one or more tipping events by 2100, reduced from 46% in the baseline model, or 87% under a prescribed RCP8.5 emissions scenario (Table 2).

A factor of 2.4 increase from the baseline SCC to US\$36 per tCO₂ is just due to the change to $IES = 1.5$ (dashed black line, Fig. 2), with a further factor of 3.2 increase due to the potential for multiple tipping points. With just $IES = 1.5$ (and no stochastic tipping points), the initial emissions control rate increases from ~18 to ~29% with 100% emissions control in 2100. Atmospheric carbon peaks around 550 ppm, with surface temperature stabilizing around 2.3 °C above pre-industrial.

Tipping point interactions

In the full model, there are both positive and negative causal interactions between tipping points (Fig. 1 and Supplementary Table 1), which are conservatively calibrated (see Methods). Hence,

Table 2 | Expected tipping point probabilities (%) by years 2100 and 2200, based on 10,000 model runs of the DSICE model¹⁹ with five stochastic tipping points, and those that would be obtained from the temperature paths in the deterministic baseline model without tipping points, or under prescribed RCP 8.5 emissions.

Number of tipping events	Stochastic tipping points (interacting)		Stochastic tipping points (no interaction)		Baseline model temperature path*		RCP8.5 temperature path†	
	2100	2200	2100	2200	2100	2200	2100	2200
1	10.8	24.38	12.04	26.88	34.28	23.03	29.69	0
2	0.65	4.14	0.72	4.08	10.03	31.31	30.73	0
3	0.04	0.42	0.05	0.41	1.81	24.7	19.08	0.33
4	0	0.02	0	0.02	0.18	10.1	6.76	16.87
5	0	0.01	0	0	0	2.29	0.85	82.80
Cumulative probability	11.49	28.97	12.81	31.39	46.30	91.43	87.11	100

*2.8°C warming in 2100, 2.76°C in 2200. †4.7°C warming in 2100, 7.5°C in 2200.

their inclusion has only a modest net effect on the expected SCC, increasing it from US\$109 per tCO₂ to US\$116 per tCO₂ (see also Supplementary Fig. 3). However, a specific sample path where multiple tipping events occur before 2200 (Fig. 3, solid line) reveals that some tipping point interactions can have a strong effect on the time evolution of the SCC. Considering a no-interactions sample path (Fig. 3, dashed line) shows that, in general, passing a tipping point reduces the incentive to mitigate and therefore lowers the SCC, because it can no longer be avoided. However, with interactions, tipping of the GIS significantly increases the likelihood of AMOC tipping (which is assumed to be the most damaging event); hence, this causes a large increase in the SCC to try to avoid AMOC tipping. (This is consistent with previous suggestions^{29,30} that tipping points can create multiple optima—here for the SCC and corresponding emissions³⁰.) Subsequent tipping of AMOC greatly reduces the SCC. Tipping of ENSO causes a small increase in the SCC because it increases the likelihood of tipping the Amazon. Subsequent tipping of the Amazon halves the SCC because there is now an unavoidable extra source of carbon to the atmosphere and only WAIS left to tip. There are other sample paths where the first tipping event does not increase the likelihood of others, so the SCC drops—for example, when the Amazon rainforest tips first (Supplementary Fig. 4).

The SCC therefore depends on whether tipping events occur and in which order. This can also be seen by looking at the sample paths for the earliest and sole tipping before 2100 of each element (Supplementary Fig. 5). If the GIS tips first, this leads to the highest SCC path and the most stringent emission control, reaching 100% before 2040, because of the increased risk of AMOC collapse. If the AMOC tips first, this gives the lowest SCC path because it has the greatest damages, which can no longer be avoided—yet emission control remains above 60% and the SCC remains above US\$110 per tCO₂. If the Amazon tips first, this also lowers SCC and emission control, but it leads to the highest atmospheric carbon and temperature trajectory because of an accompanying carbon source. If ENSO tips first, this slightly increases emission control because the likelihood of the AMAZ tipping is increased. If the WAIS tips first, there is little effect on emission control because it only slightly increases the likelihood of tipping the AMOC and GIS. CO₂ emissions trajectories (Supplementary Fig. 6) therefore depend on the contemporaneous state of tipping elements.

Sensitivity analysis

The high SCC is robust to sensitivity analyses (see Methods). Combined variations in assumed transition times and final damages of the tipping points give a full range in initial SCC of US\$50–166 per tCO₂ (Supplementary Table 2). With pessimistic settings for the expert assessment of interactions between tipping elements (Supplementary Table 3), the SCC increases from US\$116 per tCO₂ to US\$121 per tCO₂. Including an endogenous transition

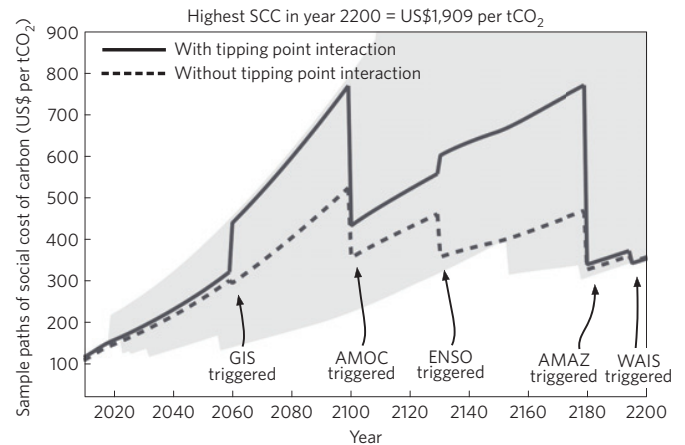


Figure 3 | Effect of causal interactions between tipping events on the social cost of carbon. Example sample paths of the SCC in US dollars per tCO₂ with multiple tipping points interacting and not interacting.

time for the GIS gives only a slight reduction in SCC to US\$114 per tCO₂ because its damages tend to be discounted away anyway. Allowing all tipping elements to have an endogenous transition time reduces SCC to US\$94 per tCO₂.

Retaining an inter-temporal elasticity of substitution IES = 1.5 but increasing risk aversion to RA = 10 increases the SCC from US\$116 per tCO₂ to US\$146 per tCO₂. With the original RA = 3.066 and an upper limit of IES = 2, the SCC increases to US\$151 per tCO₂. Using the default DICE settings of IES = 1/1.45 and RA = 1.45 gives an SCC of US\$28 per tCO₂, a factor 1.9 increase from the default US\$15 per tCO₂ due to the five interacting tipping points. Thus, EZ preferences magnify the effect of including potential future tipping points, causing a factor 3.2 (rather than 1.9) increase in the SCC. To disentangle the effect of IES and RA, we also investigate a case with IES = 1.5 and RA = 1/1.5, which gives an SCC of US\$104 per tCO₂. That is, when we incorporate the climate tipping risks, using time-separable preferences as in DICE, an increase from IES = 1/1.45 (and RA = 1.45) to IES = 1.5 (and RA = 1/1.5) leads to a factor 3.7 increase in the SCC, and the additional change to our default time non-separable EZ preferences (IES = 1.5, RA = 3.066) leads to an extra SCC of US\$12 per tCO₂.

Discussion and conclusion

Putting our results in scientific context, there is already evidence that major ice sheets are losing mass at an accelerating rate^{31,32}. GIS mass loss is estimated to be contributing ~0.7 mm yr⁻¹ to sea-level rise³³, with a corresponding increase in freshwater flux to the North Atlantic³⁴ since 1990 of ~0.01 Sv. Although modest

at present, this and other contributors to increasing freshwater input to the North Atlantic³⁵ are thought⁸ to increase the likelihood of AMOC tipping, and our results suggest that this should be increasing the incentive to control CO₂ emissions. WAIS mass loss is contributing ~ 0.35 mm yr⁻¹ to sea-level rise³², and there is evidence that parts of the West Antarctic ice sheet are already in irreversible retreat^{36–38}. If the WAIS has already passed a tipping point, then mitigation cannot avoid it, but our results suggest that this should not significantly reduce the incentive to mitigate to try to avoid other tipping events.

Our results and policy recommendations differ considerably from another recent study considering multiple tipping points¹⁸, which recommends at most a doubling of the SCC that allows CO₂ emissions to continue to grow past mid-century, with temperature ultimately peaking at just under 3 °C. In contrast, our results recommend a nearly eightfold increase in the SCC to drive a cessation of CO₂ emissions by mid-century, which limits warming to <1.5 °C. This very different outcome is a result of our different specification of tipping points together with our change in decision-maker preferences to something more appropriate for such stochastic climate risks.

There are several caveats with the DICE modelling approach used here (and the simplified version of DICE used elsewhere¹⁸). In the climate component of the model, the ocean carbon sink is too strong³⁹, causing it to overestimate the effect of emissions reductions on atmospheric CO₂ and temperature, especially beyond 2100. We consider only one value for equilibrium climate sensitivity (2.9 °C following DICE-2013), whereas the Intergovernmental Panel on Climate Change likely range⁴⁰ spans 1.5–4.5 °C. Nevertheless, the DICE prediction that a shutdown of CO₂ emissions by mid-century will lead to ~ 1.5 °C warming is compatible with more detailed probabilistic projections^{41,42} varying climate sensitivity (noting that DICE shuts down emission faster but then does not allow for net carbon dioxide removal in the second half of this century^{41,42}).

The economic component of DICE allows for an unrealistic instantaneous adjustment of emissions (to, for example, a control rate >0.5), whereas in reality emissions control rates are low and there are lags in ramping them up, for example due to the lifetime of coal-fired power stations. However, recent energy-economic model studies^{41,42} show that it is technologically feasible to increase the emissions control rate to 100%, and thus achieve net zero CO₂ emissions, by mid-century. The assumed costs of mitigation options in DICE are also relatively low⁴³, whereas energy-economic models⁴¹ indicate that limiting warming to 1.5 °C would be considerably more expensive than limiting it to 2 °C, especially between now and 2030. Despite these uncertainties, in a real options analysis framework⁴⁴, paying up front now to minimize the future risk of climate tipping points can still be the logical and cost-effective option for societies. Furthermore, acknowledging that society also faces other potential tipping points (for example, disease pandemics) should increase the willingness to pay to avert any one of them⁴⁵, even though we should not necessarily avert all of them⁴⁵. The decision to try to avert climate tipping points depends crucially on a relatively high risk aversion⁴⁵, consistent with our findings.

In summary, our results illustrate that the prospect of multiple interacting climate tipping points with irreversible economic damages ought to be provoking very strong mitigation action, on the part of 'social planners'—including governments signed up to the United Nations Framework Convention on Climate Change. Under realistic preferences under uncertainty, the optimal policy involves a shutdown of carbon emissions by mid-century.

Methods

Methods and any associated references are available in the [online version of the paper](#).

Received 2 October 2015; accepted 16 February 2016;
published online 21 March 2016

References

1. Interagency Working Group on Social Cost of Carbon *Social Cost of Carbon for Regulatory Impact Analysis—Under Executive Order 12866* (US Government, 2010).
2. Interagency Working Group on Social Cost of Carbon *Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis* (US Government, 2013).
3. Dietz, S. High impact, low probability? An empirical analysis of risk in the economics of climate change. *Climatic Change* **108**, 519–541 (2011).
4. Kopp, R. E., Golub, A., Keohane, N. O. & Onda, C. The influence of the specification of climate change damages on the social cost of carbon. *Economics* **6**, 1–40 (2012).
5. Ackerman, F. & Stanton, E. A. Climate risks and carbon prices: revising the social cost of carbon. *Economics* **6**, 1–25 (2012).
6. van den Bergh, J. C. J. M. & Botzen, W. J. W. A lower bound to the social cost of CO₂ emissions. *Nature Clim. Change* **4**, 253–258 (2014).
7. Lenton, T. M. *et al.* Tipping elements in the Earth's climate system. *Proc. Natl Acad. Sci. USA* **105**, 1786–1793 (2008).
8. Kriegler, E., Hall, J. W., Held, H., Dawson, R. & Schellnhuber, H. J. Imprecise probability assessment of tipping points in the climate system. *Proc. Natl Acad. Sci. USA* **106**, 5041–5046 (2009).
9. Lenton, T. M. Early warning of climate tipping points. *Nature Clim. Change* **1**, 201–209 (2011).
10. Lontzek, T. S., Cai, Y., Judd, K. L. & Lenton, T. M. Stochastic integrated assessment of climate tipping points indicates the need for strict climate policy. *Nature Clim. Change* **5**, 441–444 (2015).
11. Lenton, T. M. & Ciscar, J.-C. Integrating tipping points into climate impact assessments. *Climatic Change* **117**, 585–597 (2013).
12. Mastrandrea, M. D. & Schneider, S. H. Probabilistic integrated assessment of 'dangerous' climate change. *Science* **304**, 571–575 (2004).
13. Kosugi, T. Integrated assessment for setting greenhouse gas emission targets under the condition of great uncertainty about the probability and impact of abrupt climate change. *J. Environ. Inform.* **14**, 89–99 (2009).
14. Ackerman, F., Stanton, E. A. & Bueno, R. Fat tails, exponents, extreme uncertainty: simulating catastrophe in DICE. *Ecol. Econ.* **69**, 1657–1665 (2010).
15. Weitzman, M. L. GHG targets as insurance against catastrophic climate damages. *J. Public Econ. Theor.* **14**, 221–244 (2012).
16. Cai, Y., Judd, K. L., Lenton, T. M., Lontzek, T. S. & Narita, D. Environmental tipping points significantly affect the cost-benefit assessment of climate policies. *Proc. Natl Acad. Sci. USA* **112**, 4606–4611 (2015).
17. Lemoine, D. & Traeger, C. Watch your step: optimal policy in a tipping climate. *Am. Econ. J.* **6**, 137–166 (2014).
18. Lemoine, D. & Traeger, C. P. Economics of tipping the climate dominoes. *Nature Clim. Change* <http://dx.doi.org/10.1038/nclimate2902> (2016).
19. Cai, Y., Judd, K. L. & Lontzek, T. S. The social cost of carbon with economic and climate risks. Preprint at <http://arxiv.org/abs/1504.06909> (2015).
20. Nordhaus, W. Estimates of the social cost of carbon: concepts and results from the DICE-2013R model and alternative approaches. *J. Assoc. Environ. Res. Econ.* **1**, 273–312 (2014).
21. Meinshausen, M. *et al.* The RCP greenhouse gas concentrations and their extensions from 1765 to 2300. *Climatic Change* **109**, 213–241 (2011).
22. Epstein, L. G. & Zin, S. E. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework. *Econometrica* **57**, 937–969 (1989).
23. Pindyck, R. S. & Wang, N. The economic and policy consequences of catastrophes. *Am. Econ. J.* **5**, 306–339 (2013).
24. Vissing-Jørgensen, A. & Attanasio, O. P. Stock-market participation, intertemporal substitution, and risk-aversion. *Am. Econ. Rev.* **93**, 383–391 (2003).
25. Bansal, R. & Yaron, A. Risks for the long run: a potential resolution of asset pricing puzzles. *J. Finance* **59**, 1481–1509 (2004).
26. Barro, R. J. Rare disasters, asset prices, and welfare costs. *Am. Econ. Rev.* **99**, 243–264 (2009).
27. Nordhaus, W. D. Expert opinion on climatic change. *Am. Sci.* **82**, 45–51 (1994).
28. Nordhaus, W. D. & Boyer, J. *Warming the World. Models of Global Warming* (MIT Press, 2000).
29. Baumol, W. J. On taxation and the control of externalities. *Am. Econ. Rev.* **62**, 307–322 (1972).
30. Kopp, R. E. & Mignone, B. K. The U.S. government's social cost of carbon estimates after their first two years: pathways for improvement. *Economics* **6**, 1–41 (2012).
31. Khan, S. A. *et al.* Sustained mass loss of the northeast Greenland ice sheet triggered by regional warming. *Nature Clim. Change* **4**, 292–299 (2014).

32. Harig, C. & Simons, F. J. Accelerated West Antarctic ice mass loss continues to outpace East Antarctic gains. *Earth Planet. Sci. Lett.* **415**, 134–141 (2015).
33. Csatho, B. M. *et al.* Laser altimetry reveals complex pattern of Greenland Ice Sheet dynamics. *Proc. Natl Acad. Sci. USA* **111**, 18478–18483 (2014).
34. Bamber, J., van den Broeke, M., Ettema, J., Lenaerts, J. & Rignot, E. Recent large increases in freshwater fluxes from Greenland into the North Atlantic. *Geophys. Res. Lett.* **39**, L19501 (2012).
35. Peterson, B. J. *et al.* Increasing river discharge to the Arctic Ocean. *Science* **298**, 2171–2173 (2002).
36. Joughin, I., Smith, B. E. & Medley, B. Marine ice sheet collapse potentially under way for the Thwaites Glacier Basin, West Antarctica. *Science* **344**, 735–738 (2014).
37. Rignot, E., Mouginot, J., Morlighem, M., Seroussi, H. & Scheuchl, B. Widespread, rapid grounding line retreat of Pine Island, Thwaites, Smith, and Kohler glaciers, West Antarctica, from 1992 to 2011. *Geophys. Res. Lett.* **41**, 3502–3509 (2014).
38. Wouters, B. *et al.* Dynamic thinning of glaciers on the Southern Antarctic Peninsula. *Science* **348**, 899–903 (2015).
39. Glotter, M. J., Pierrehumbert, R. T., Elliott, J. W., Matteson, N. J. & Moyer, E. J. A simple carbon cycle representation for economic and policy analyses. *Climatic Change* **126**, 319–335 (2014).
40. Bindoff, N. L. *et al.* in *Climate Change 2013: The Physical Science Basis* (eds Stocker, T. F. *et al.*) Ch. 10 (IPCC, Cambridge Univ. Press, 2013).
41. Rogelj, J. *et al.* Energy system transformations for limiting end-of-century warming to below 1.5 °C. *Nature Clim. Change* **5**, 519–527 (2015).
42. Rogelj, J. *et al.* Zero emission targets as long-term global goals for climate protection. *Environ. Res. Lett.* **10**, 105007 (2015).
43. Nordhaus, W. *The Climate Casino: Risk, Uncertainty, and Economics for a Warming World* (Yale Univ. Press, 2013).
44. Anda, J., Golub, A. & Strukova, E. Economics of climate change under uncertainty: benefits of flexibility. *Energy Policy* **37**, 1345–1355 (2009).
45. Martin, I. W. R. & Pindyck, R. S. Averting catastrophes: the strange economics of scylla and charybdis. *Am. Econ. Rev.* **105**, 2947–2985 (2015).

Acknowledgements

We thank K. L. Judd and participants of the 2015 Annual Conference of the European Association of Environmental and Resource Economics for comments. Y.C. was supported by NSF (SES-0951576 and SES-146364). T.S.L. was supported by the Züricher Universitätsverein, the University of Zurich, and the Ecosciencia Foundation. T.M.L. was supported by a Royal Society Wolfson Research Merit Award and the European Commission HELIX project (ENV.2013.6.1-3). Supercomputer support was provided by Blue Waters (NSF awards OCI-0725070 and ACI-1238993, and the state of Illinois).

Author contributions

Y.C., T.M.L. and T.S.L. designed research, performed research and wrote the paper.

Additional information

Supplementary information is available in the [online version of the paper](#). Reprints and permissions information is available online at www.nature.com/reprints. Correspondence and requests for materials should be addressed to T.M.L. or T.S.L.

Competing financial interests

The authors declare no competing financial interests.

Methods

Summary. We use the DSICE model^{10,19} (Supplementary Fig. 1) to compute the socially optimal reduction of global greenhouse gas emissions under the possibility of five interacting climate tipping points. The baseline deterministic model without tipping points is based on the 2013 version of DICE²⁰, but uses parameters in the carbon cycle and temperature system calibrated against all four RCP scenarios (see Supplementary Methods), and solves on an annual time step. DICE comprises one state variable for the capital stock, representing the world economy, a three-box carbon cycle module, and a two-box climate. To this we add a ten-dimensional system of interacting tipping elements.

For each of five tipping elements, we have a discrete binary state indicating whether its corresponding tipping process has been already triggered or not, and a continuous state variable indicating the contemporaneous length of the transition process. The occurrence of each climate tipping point is modelled by a Markov process and its timing is not known at the times of decisions. The endogenous hazard rate ($\text{yr}^{-1} \text{K}^{-1}$) for each tipping event is assumed zero up to 1°C warming above pre-industrial levels (reached in about 2015 in the model) and increases linearly with global warming above 1°C at a rate derived from published expert elicitation results⁸. The conditional probabilities representing changes to the other hazard rates should a particular system tip are conservatively specified given wide ranges in the expert assessment⁸. The transition timescale¹⁰ of each tipping element is based on current scientific understanding of the timescales at which specific climate subsystems can transition into an alternative state, with a factor of 5 uncertainty range in either direction considered in the sensitivity analysis. Tipping points are assumed to directly impact economic output and their relative final damages are based on scientific understanding. The absolute final damages of individual tipping events are highly uncertain and are varied in the sensitivity analysis over a factor of 2–3 range, giving a range in total reduction in GDP if all five tipping events occur of 23–50%. In addition to the impacts of tipping points on economic output we also include conservative effects of tipping particular systems on the carbon cycle, implemented as exogenous emissions to the atmosphere. The stochastic model is solved using a supercomputer^{19,46}, to generate 10,000 stochastic sample paths, with the expected path calculated as the average of all paths.

In the following, we detail the specific modifications to the DICE-2013R model and refer to Nordhaus⁴³ for calibration and formulations of the remaining parts of the model.

Calibration of tipping elements and interactions between them. As in previous work¹⁰, we define three phases to the tipping process for each tipping element (Supplementary Fig. 2). In the first, pre-trigger phase, the additional damage from a tipping point is 0. In the second, transition phase, there is a positive, but not stationary additional damage level. In the third and final, post-tipping phase the tipping element is in a new, absorbing state, with a constant (irreversible) damage level.

For each tipping element, i , after a tipping point is passed, the additional damage factor $J_{i,t}$ will increase continuously from a minimal level (that is, $J_{i,t} = 0$) to some maximum and persistent climate impact level ($\bar{J}_i > 0$), implying that $J_{i,t+1} = \min\{J_{i,t} + \Delta_{i,t}, \bar{J}_i\}$, where $\Delta_{i,t}$ is the incremental impact level from stage t to $t + 1$ of tipping element i . In our default case, $\Delta_{i,t}$ denotes linear increments, but these increments become nonlinear in the sensitivity case with endogenous transition time. We use $I_{i,t}$ as the indicator function to denote for each tipping element i the pre-trigger state of the world as $I_{i,t} = 0$ and the post-trigger state of the world as $I_{i,t} = 1$, where $I_{i,t}$ is a jump process with a Markovian hazard rate. The latter is endogenous with respect to the contemporaneous level of global average atmospheric temperature, T_t^{AT} . Furthermore, to model causal relationships between the tipping elements, the Markovian hazard rate for tipping element i also depends on whether a tipping process of climate tipping element j has been triggered. We do not explicitly consider other indicators for tipping, for example, the gradient of temperature⁴⁷. The transition function for $I_{i,t}$ from stage t to stage $t + 1$ is $I_{i,t+1} = g_i^t(\mathbf{I}_t, T_t^{\text{AT}}, \omega_{i,t}^t)$, where \mathbf{I}_t is the vector of the indicator functions for the five climate tipping elements ($I_{1,t}, \dots, I_{5,t}$) and $\omega_{i,t}^t$ is a random process. With $J_{i,t+1} = \min\{J_{i,t} + \Delta_{i,t}, \bar{J}_i\}$ the impact factor on the economy becomes

$$\Omega_i(T_t^{\text{AT}}, \mathbf{I}_t, \mathbf{I}_t) = \frac{\prod_i (1 - I_{i,t} J_{i,t})}{1 + \pi_2 (T_t^{\text{AT}})^2} \quad (1)$$

where T_t^{AT} is the average global atmospheric temperature and π_2 is a coefficient in the damage function. (The impact of global warming on the economy is reflected by a convex damage function of atmospheric temperature, which is a standard feature of the DICE model—a deterministic model specification would simply be to fix all $I_{i,t}$ at 0.) We specify the probability transition matrix of the tipping process i at time t as

$$\begin{bmatrix} 1 - p_{i,t} & p_{i,t} \\ 0 & 1 \end{bmatrix} \quad (2)$$

where its (n, m) element is the transition probability from state n to m for $I_{i,t}$, and $p_{i,t} = 1 - \exp(-B_i(\mathbf{I}) \max\{0, T_t^{\text{AT}} - 1\})$, where $B_i(\mathbf{I})$ is the hazard rate function

for tipping element i , depending on whether other tipping elements have tipped. A general formula for the hazard rate function is given by

$$B_i(\mathbf{I}) = b_i \left(1 + \sum_j (I_j f_{ij}) \right) \quad (3)$$

We calibrated the values for b_i using the expert opinions reported in Kriegler *et al.*⁸ and our previously described methodology¹⁰. Specifically, we calibrated b_i to match the average expert's cumulative trigger probabilities for each tipping element by the year 2200 for the medium temperature corridor in Kriegler *et al.*⁸, which implies 2.5°C warming in 2100 and 3°C warming in 2200. These probabilities are 22% for AMOC, 52% for GIS, 34% for WAIS, 48% for AMAZ and 19% for ENSO. The corresponding values for b_i are $b_{\text{AMOC}} = 0.00063064$, $b_{\text{GIS}} = 0.00188445$, $b_{\text{WAIS}} = 0.00103854$, $b_{\text{AMAZ}} = 0.00163443$ and $b_{\text{ENSO}} = 0.000526678$ (Table 1).

To model the interaction component of tipping point likelihood, we introduce f_{ij} as an additional probability factor, which describes by how much the hazard factor for tipping element j is affected if tipping element i has tipped (when it is negative, it implies a decrease in probability). The parameter matrix f_{ij} is calibrated for $i, j \in \{\text{AMOC, GIS, WAIS, AMAZ, ENSO}\}$. Again we use the results in Kriegler *et al.*⁸ as the source for our calibration of the interaction effects between tipping elements. In particular, we consider the core experts' assessment of the interaction effects for the 'medium' temperature corridor. Our aim is to implement the interactions as direct, conditional alterations to the hazard rate of individual tipping events. Supplementary Table 1 summarizes our calibrated factors, f_{ij} . For some of the interaction effects, experts assessed ambiguous effects. For example, in the case of WAIS affecting AMOC the interaction factor ranges between <0 and >0 among the experts and among the average optimistic and pessimistic opinions of the core experts. In such an ambiguous case, although it might be worthwhile incorporating this uncertainty in the direction of interaction, we leave that as a possible avenue for further research and focus here, as in the non-ambiguous cases, solely on the average core experts' assessment.

The order of the tipping sequence is important for the overall impact of any individual tipping element, owing to asymmetric causal relationships between some of the tipping events (Fig. 1 and Supplementary Table 1). For example, when GIS tipping is triggered first, the likelihood of AMOC is increased, but if instead a tipping point in the AMOC is triggered first, the likelihood of GIS tipping is reduced.

Specification of transition times, final damages and carbon cycle effects. In addition to calibrating the hazard rate (described above), we have to specify the transition time, final damage levels and the effect on the carbon cycle for each tipping element (Table 1). We base this on reviews of the literature, updated from previous work¹¹. Recognizing the scientific and economic uncertainties in these choices, the transition times are given a common factor of 5 range of uncertainty in either direction from default values, and the final damages are given a factor of 2–3 total uncertainty range. The values chosen are briefly justified as follows.

AMOC. Past abrupt climate changes linked to reorganizations of the AMOC have occurred in a decade or less, but future AMOC collapse in model simulations can take a couple of centuries. Hence, we opt for a 50-year default transition time and 10–250 year range. The AMOC collapse is often viewed as the archetype of a climate catastrophe; hence, we assign it the highest final damage (accepting that others will question this). Past studies with DICE have suggested that a collapse of the AMOC might result in a 25–30% reduction in GDP comparable with the Great Depression^{27,28}. However, when combined with other tipping events this could lead to excessively high damages, so we opt for a 15% GDP reduction with a range of 10–20%. We considered the potential for the AMOC collapse to reduce both ocean heat⁴⁸ and carbon^{49,50} uptake. However, quantitative estimates of these effects based on existing studies^{48–50} suggest that they are small; hence, they are ignored here.

GIS. Irreversible meltdown of the Greenland ice sheet typically takes millennia in model simulations^{51,52}, but models are unable to explain the speed of recent ice loss⁷. To cover the uncertainty, we opt for a default timescale of 1,500 years, with a minimum timescale⁷ of 300 years and an upper limit of 7,500 years. The final damages from the GIS melt will largely be due to sea-level rise⁷ of around 7 m, which is roughly twice what can come from WAIS disintegration⁵³. Hence, we give the GIS twice the default final damages of the WAIS, noting that the spatial pattern of sea-level rise will be greatest farthest away from each ice sheet (owing to gravitational effects). As well as flooding low-lying cities and agricultural land, flooding of large areas of low-lying permafrost (especially in Siberia) could ultimately release large amounts of carbon¹¹. We conservatively assume an exogenous emission of 100 GtC over the duration of the transition, which is only ~6% of the total permafrost carbon reservoir⁵⁴.

WAIS. The West Antarctic ice sheet is grounded largely below sea level and has the potential for more rapid disintegration than the Greenland ice sheet⁷, ultimately leading to up to 3.3 m sea-level rise⁵⁵. Past sea-level rise in the penultimate Eemian

inter-glacial period is estimated to have occurred⁵⁵ at rates >1 m per century and must have come from Antarctica and/or Greenland. We assign a minimum timescale of 100 years for WAIS disintegration, with a default setting of 500 years, and an upper limit of 2,500 years. Noting that the effect of GIS meltdown on Arctic sea level is greatly suppressed by gravitational adjustment⁵⁶, whereas that of WAIS disintegration is not⁵³, we assign WAIS the same potential to release 100 GtC from low-lying permafrost over the duration of the transition.

AMAZ. Dieback of the Amazon rainforest in future model simulations⁵⁷ takes around 50 years, which we use as our default. However, if drought and corresponding fires respond very nonlinearly to climate change⁵⁸ dieback could conceivably occur on a minimum timescale of 10 years, whereas if the forest is more resilient it could take centuries, consistent with a maximum timescale of 250 years. The Amazon rainforest is estimated to store 150–200 GtC in living biomass and soils⁵⁹ and we conservatively assume that dieback will release 50 GtC over the duration of the transition.

ENSO. In the past the frequency and amplitude of ENSO variability has changed on decadal to centennial timescales⁷, and in the future the amplitude of ENSO variability is expected to increase with more frequent extreme El Niño and extreme La Niña events⁶⁰. Past El Niño and La Niña events have had large impacts, especially on the agricultural sector, and their more global footprint than Amazon dieback leads us to assign higher damages to ENSO. The observational record shows that individual strong El Niño events can cause anomalous emissions of carbon by fire⁶¹ of ~2 GtC. Hence, we assume that an increase in El Niño amplitude could readily cause an average increase in land carbon emissions (exogenous) by 0.2 GtC yr⁻¹ that is essentially permanent on the timescale of our integrations.

The combined effect on final damages if all tipping points occur is 38%, with a 23–50% range in our sensitivity analysis. However, the timescale for all damages to be felt in our default case is over 1,000 years, and our tipping probabilities are relatively low. Only two tipping elements (GIS, AMAZ) have an expected tipping time around 2200 (when it is as likely as not that their tipping process will be triggered), with the remaining three elements being less likely to tip. Furthermore, slow transition times mean that damages tend to be discounted away. As we have shown previously¹⁰, a tipping point with 2.5% damage to GDP and a 5-year transition time will have much larger impact on the SCC today than a tipping point with 25% damage to GDP and a 500-year transition time. Other integrated assessment model studies that treat tipping points have tended to assume instantaneous transitions and double-digit percentage damages. Thus, we argue that overall our model is conservatively calibrated with relatively low expected damages, which amount to 0.53% of GDP in 2100 and 1.89% of GDP in 2200 in our default model parameterization.

The couplings to the carbon cycle lead to the following new specification of the exogenous land carbon source (in GtC) in DSICE:

$$E_{Land,t} = 0.9e^{-0.04t} + I_{GIS}I_{-}\{J_{GIS} < \bar{J}_{GIS}\} \frac{100}{1500} + I_{WAIS}I_{-}\{J_{WAIS} < \bar{J}_{WAIS}\} \frac{100}{500} + I_{AMAZ}I_{-}\{J_{AMAZ} < \bar{J}_{AMAZ}\} \frac{50}{50} + 0.2 (J_{ENSO}/\bar{J}_{ENSO}) \quad (4)$$

where the first term on the right-hand side is from the DICE model and all remaining terms are our modifications. Here, $I_{-}\{\}$ serves as an indicator function.

The dynamic programming problem. In the following we present the dynamic programming problem of the social planner:

$$V_t(\mathbf{S}) = \max_{C_t, \mu_t} u(C_t, L_t) + \beta \left[E \left\{ \left(V_{t+1}(\mathbf{S}^+) \right)^{\frac{1-\gamma}{1-\beta\gamma}} \right\} \right]^{\frac{1-\beta}{1-\beta\gamma}} \quad (5)$$

$$\text{s.t.} \quad K^+ = (1-\delta)K + Y_t(K, T^{AT}, \mathbf{I}, \mathbf{J}) - C_t - \Psi_t \quad (6)$$

$$\mathbf{M}^+ = \Phi^M \mathbf{M} + (\varepsilon_t(K, \mu), 0, 0)^T \quad (7)$$

$$\mathbf{T}^+ = \Phi^T \mathbf{T} + (\xi_t F_t(M^{AT}), 0)^T \quad (8)$$

$$I_i^+ = g_i(\mathbf{I}, T^{AT}, \omega_i) \quad (9)$$

$$J_i^+ = \min \{ J_i + \Delta_i, \bar{J}_i \} I_i \quad (10)$$

where $V_t(\mathbf{S})$ denotes the time t value function, which is endogenous in the sixteen-dimensional state vector denoted by \mathbf{S} . Furthermore, C_t, μ_t are the control variables for consumption and mitigation. Each period's utility u depends on consumption and exogenous labour supply L_t . With β we denote the utility discount rate. The expectation operator is over the next period's value function with γ and ψ denoting the risk aversion parameter and the elasticity of inter-temporal substitution, respectively. The utility also depends on the elasticity of inter-temporal substitution, but not on the risk aversion parameter, so that the risk aversion parameter will play a role only for stochastic cases as the second term of the objective function of the maximization problem (5) will be simplified to be the discounted next-period value function for deterministic cases. In our default parameter case, we follow the calibration by Pindyck and Wang²³ and specify: $\gamma = 3.066$ and $\psi = 1.5$. Furthermore, K, \mathbf{M} and \mathbf{T} denote the capital stock, the three carbon stocks and the two temperatures (M^{AT} and T^{AT} represent carbon concentration and temperature in the atmosphere), respectively and a '+' superscript denotes a variable's next period value. Y_t denotes world gross product net of damages and ε_t denotes non-mitigated emissions into the atmosphere. Finally, Ψ_t is the expenditure on mitigation, and F_t is a term related to radiative forcing. The model is solved for the next 300 years with a terminal value function approximating the welfare of future years from 301 to an infinite horizon (see Supplementary Methods). Our SCC is computed using

$$SCC_t = -1,000 \left(\frac{\partial V_t}{\partial M_t^{AT}} \right) / \left(\frac{\partial V_t}{\partial K_t} \right)$$

as in DSICE¹⁹, denoting the marginal rate of substitution between atmospheric carbon concentration and capital.

After solving the dynamic programming problem using parallel backward value function iteration⁴⁶ (see Supplementary Methods), we use these approximated value functions V_t to simulate 10,000 paths in the following way: at the initial time, its state vector \mathbf{S}_0 is known as the observed market values; then we can get the optimal consumption and emission control rate at time 0 by solving the dynamic programming problem with the previously computed V_1 . Using sample realization of shocks, we can obtain the next state vector \mathbf{S}_1 ; using the same method to iterate forward, we get one simulated path of states and optimal policies that depend on realization of shocks. Repeating this process, we get 10,000 sample paths for our analysis.

Numerical implementation of the model. We have found that for the relatively short time horizon, when recalibrating the carbon cycle and temperature modules to match all four RCP scenarios closely we can omit the deep ocean stock of carbon without any loss of accuracy in the carbon-to-temperature relationship. Thus, the numerical implementation of the model is fifteen-dimensional. The computational task required to solve this fifteen-dimensional problem goes far beyond what has previously been achieved in truly stochastic climate-economy models, where three- to four-dimensional problems are considered the current frontier. We solve the model with parallel dynamic programming methods⁴⁶ on 312,500,000 approximation nodes for the ten-dimensional continuous state space and degree-4 complete Chebyshev polynomials for each of the five discrete state vectors. It takes about 3 h to solve the model for a single set of parameter values on 10,560 cores at the Blue Waters supercomputer. The estimated error bound of the optimal solution is 0.1–1% for policy functions and 0.01–0.1% for the value functions.

Sensitivity analyses. We conducted several sensitivity analyses. First, we varied the transition times and/or damages of all five tipping elements across their assigned uncertainty ranges. Second, we took a more pessimistic assessment of the interaction between the tipping elements (Supplementary Table 3), which uses the upper bounds of the core experts' assessment.

Third, some more complex sensitivity studies were also conducted exploring the effect of endogenous transition times for tipping elements. In our model, the transition time for tipping element i is inversely tied to $\Delta_{i,t}$, the annual damage increase during the transition phase. Thus, the transition time for element i is proportional to $(1/\Delta_{i,t})$ and also its final damage level \bar{J}_i . In the case of an endogenous transition time, we let the annual damage increase be $\Delta_{i,t} = \bar{J}_i \exp(a_i T_t^{AT} - b_i)$, where a_i and b_i are parameters calibrated such that $\bar{J}_i/\Delta_{i,t}$ results in the upper-bound value of the range of transition times considered for $T_t^{AT} = 0$, and in the lower-bound value of that range for $T_t^{AT} = 6$. Thus, the endogenous transition time is equal to $\int_0^\infty \exp(a_i T_t^{AT} - b_i) I_{i,t} I_{-}(J_{i,t} < \bar{J}_i) dt$.

As a general rule, transition timescales should be governed by the internal dynamical timescale(s) of the system in question, so it may not be appropriate to include a temperature dependence of the transition timescale for all tipping elements. However, endogenous transition times have some backing for the major ice sheets, where models^{51,52} show that the rate of ice sheet meltdown depends on the amount by which a temperature threshold is exceeded.

References

46. Cai, Y., Judd, K. L., Thain, G. & Wright, S. J. Solving dynamic programming problems on a computational grid. *Comput. Econ.* **45**, 261–284 (2015).
47. Nævdal, E. & Oppenheimer, M. The economics of the thermohaline circulation—a problem with multiple thresholds of unknown location. *Resour. Energy Econ.* **29**, 262–283 (2007).
48. Kostov, Y., Armour, K. C. & Marshall, J. Impact of the Atlantic meridional overturning circulation on ocean heat storage and transient climate change. *Geophys. Res. Lett.* **41**, 2108–2116 (2014).
49. Perez, F. F. *et al.* Atlantic Ocean CO₂ uptake reduced by weakening of the meridional overturning circulation. *Nature Geosci.* **6**, 146–152 (2013).
50. Zickfeld, K., Eby, M. & Weaver, A. J. Carbon-cycle feedbacks of changes in the Atlantic meridional overturning circulation under future atmospheric CO₂. *Glob. Biogeochem. Cycles* **22**, GB3024 (2008).
51. Huybrechts, P. & De Wolde, J. The dynamic response of the Greenland and Antarctic ice sheets to multiple-century climatic warming. *J. Clim.* **12**, 2169–2188 (1999).
52. Robinson, A., Calov, R. & Ganopolski, A. Multistability and critical thresholds of the Greenland ice sheet. *Nature Clim. Change* **2**, 429–432 (2012).
53. Bamber, J. L., Riva, R. E. M., Vermeersen, B. L. A. & LeBrocq, A. M. Reassessment of the potential sea-level rise from a collapse of the West Antarctic ice sheet. *Science* **324**, 901–903 (2009).
54. Tarnocai, C. *et al.* Soil organic carbon pools in the northern circumpolar permafrost region. *Glob. Biogeochem. Cycles* **23**, GB2023 (2009).
55. Rohling, E. J. *et al.* High rates of sea-level rise during the last interglacial period. *Nature Geosci.* **1**, 38–42 (2008).
56. Mitrovica, J. X., Tamislea, M. E., Davis, J. L. & Milne, G. A. Recent mass balance of polar ice sheets inferred from patterns of sea-level change. *Nature* **409**, 1026–1029 (2001).
57. Huntingford, C. *et al.* Towards quantifying uncertainty in predictions of Amazon 'dieback'. *Phil. Trans. R. Soc. B* **363**, 1857–1864 (2008).
58. Brando, P. M. *et al.* Abrupt increases in Amazonian tree mortality due to drought–fire interactions. *Proc. Natl Acad. Sci. USA* **111**, 6347–6352 (2014).
59. Feldpausch, T. R. *et al.* Tree height integrated into pantropical forest biomass estimates. *Biogeosciences* **9**, 3381–3403 (2012).
60. Cai, W. *et al.* ENSO and greenhouse warming. *Nature Clim. Change* **5**, 849–859 (2015).
61. van der Werf, G. R. *et al.* Continental-scale partitioning of fire emissions during the 1997 to 2001 El Niño/La Niña period. *Science* **303**, 73–76 (2004).