

# Economics of tipping the climate dominoes

Derek Lemoine<sup>1</sup> and Christian P. Traeger<sup>2,3\*</sup>

**Greenhouse gas emissions can trigger irreversible regime shifts in the climate system, known as tipping points. Multiple tipping points affect each other's probability of occurrence, potentially causing a 'domino effect'. We analyse climate policy in the presence of a potential domino effect. We incorporate three different tipping points occurring at unknown thresholds into an integrated climate–economy model. The optimal emission policy considers all possible thresholds and the resulting interactions between tipping points, economic activity, and policy responses into the indefinite future. We quantify the cost of delaying optimal emission controls in the presence of uncertain tipping points and also the benefit of detecting when individual tipping points have been triggered. We show that the presence of these tipping points nearly doubles today's optimal carbon tax and reduces peak warming along the optimal path by approximately 1 °C. The presence of these tipping points increases the cost of delaying optimal policy until mid-century by nearly 150%.**

The threat of climate tipping points plays a major role in calls for aggressive emission reductions to limit warming to 2 °C (refs 1–4). The scientific literature is particularly concerned with the possibility of a 'domino effect' from multiple interacting tipping points<sup>5–10</sup>. For instance, reducing the effectiveness of carbon sinks amplifies future warming, which in turn makes further tipping points more likely. Nearly all of the preceding quantitative economic studies analyse optimal policy in the presence of a single type of tipping point that directly reduces economic output<sup>11–14</sup>. This type of tipping point affects the potential for further tipping points only indirectly: the resulting reduction in emissions will generally reduce the likelihood of triggering further tipping points. So far, only a single paper analyses optimal climate policy in the presence of tipping points that alter the physical climate system<sup>15</sup>, specifically a temperature feedback tipping point and the carbon sink tipping point described above. These two tipping points could interact directly; however, that paper considers only a single type of tipping point at a time. The present study synthesizes the tipping point models from the previous literature to provide the first analysis of optimal climate policy when tipping points can directly interact.

Our study integrates all three types of previously modelled tipping points into a single integrated assessment model that combines smooth and reversible changes with irreversible regime shifts. Each tipping point is stochastically triggered at an unknown threshold. We solve for the optimal policy under Bayesian learning. Optimality means that resources within and across periods are distributed to maximize the expected stream of global welfare from economic consumption over time under different risk states. The optimal policy must anticipate all possible thresholds, interactions and future policy responses. The anticipation of learning acknowledges that future policymakers will have new information about the location of temperature thresholds and can also react to any tipping points that may have already occurred. Learning over the threshold location also avoids the assumption implicit in ref. 14 that tipping will eventually occur with certainty if temperatures stay permanently above the level where tipping points are possible. Finally, going beyond the conventional focus on optimal policy, we also calculate the welfare cost of delaying optimal climate policy. We demonstrate the value of monitoring for tipping points that have already been triggered, so that policy can adjust

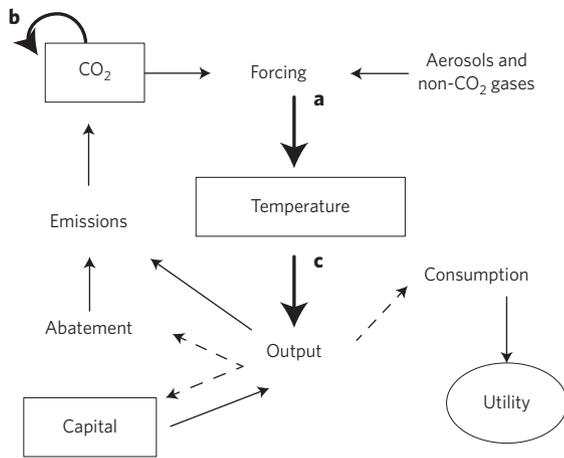
and reduce the probability of a single tipping point turning into a domino effect.

## Tipping points and economic model

We integrate three tipping points into a stochastic dynamic programming version<sup>16</sup> of the DICE integrated assessment model<sup>17,18</sup> (see Methods). DICE combines an economic growth model with a climate module (Fig. 1). The policymaker decides how to allocate output between consumption, investment in capital, and emission reductions. Unabated emissions accumulate in the atmosphere, where they change the radiative balance, induce further warming feedbacks, increase global average surface temperature, and thereby cause economic damages. In addition to integrating tipping points, uncertainty and learning, we also modify DICE's damage relationship to avoid double-counting: because we explicitly model tipping points, we eliminate an *ad hoc* adjustment that approximately doubles DICE's damages to account for the (originally unmodelled) possibility of tipping points<sup>18</sup>. Today's optimal policy has to foresee all tipping possibilities and anticipate how they affect future welfare, which in turn depends on how future policy responds to tipping at a given time and state (Fig. 2). A simplified decision tree with just 50 time periods illustrates the complexity of the problem: finding today's optimal mitigation policy requires the simultaneous solution of over 20 million coupled optimization problems. In the Methods, we explain how we solve the infinite horizon problem by solving a set of coupled non-autonomous multi-state dynamic programming problems.

The bold arrows in the schematic of Fig. 1 illustrate how our tipping points alter the climate system. The first tipping point (labelled a) makes temperature more sensitive to carbon dioxide (CO<sub>2</sub>) emissions. It reflects the possibility that warming mobilizes large methane stores locked in permafrost and in shallow ocean clathrates<sup>19–22</sup>. It also reflects the possibility that land ice sheets begin to retreat on decadal timescales: the resulting loss of reflective ice could double the long-term warming predicted by models that hold land ice sheets fixed<sup>1</sup>. DICE characterizes the temperature response to CO<sub>2</sub> through a climate sensitivity parameter that represents the equilibrium warming from doubling CO<sub>2</sub>. The value of 3 °C used in DICE is inferred from climate models that hold land ice sheets and most methane stocks constant. We represent a climate

<sup>1</sup>Department of Economics, University of Arizona, McClelland Hall 401, Tucson, Arizona 85721-0108, USA. <sup>2</sup>Department of Agricultural and Resource Economics, University of California Berkeley, 207 Giannini Hall #3310, UC Berkeley, California 94720-3310, USA. <sup>3</sup>Department of Economics, University of Oslo, Postbox 1095 Blindern, 0317 Oslo, Norway. \*e-mail: [traeger@berkeley.edu](mailto:traeger@berkeley.edu)

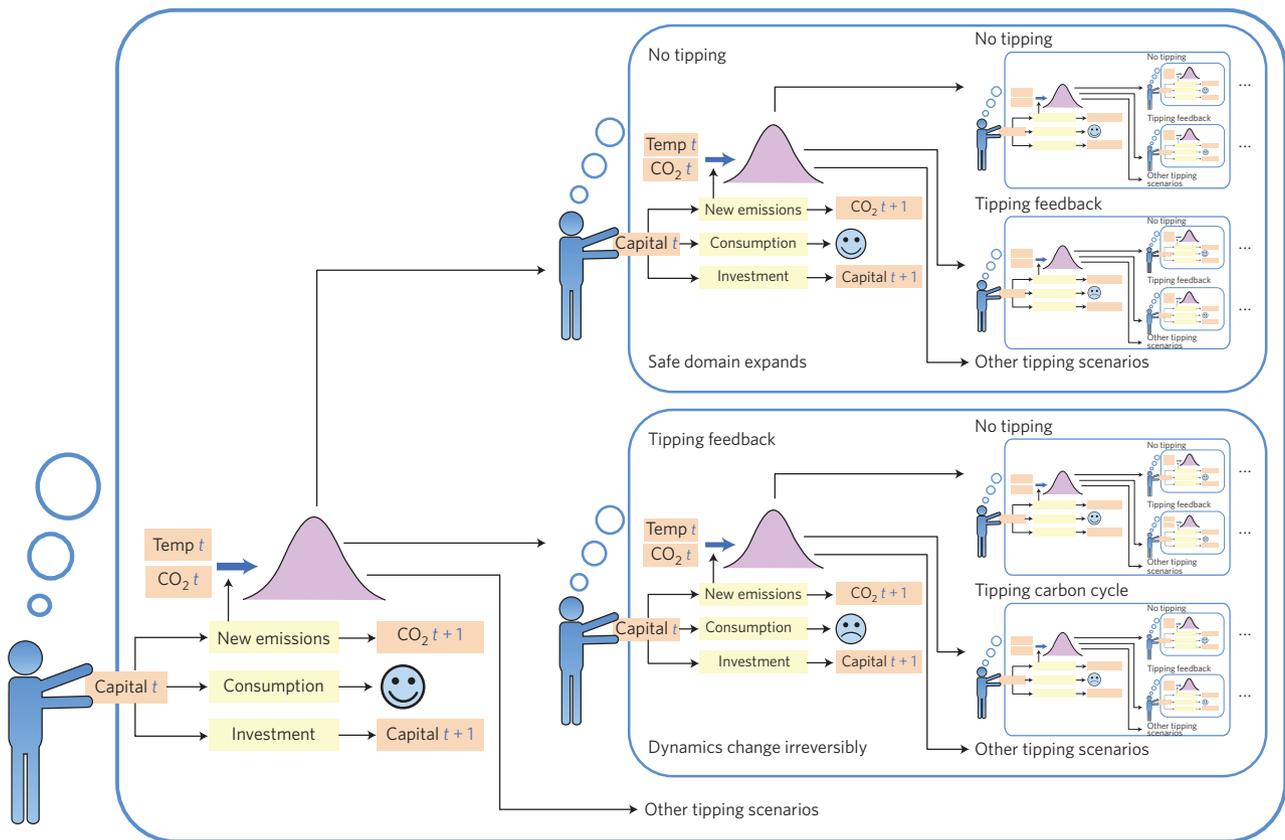


**Figure 1 | A simplified schematic of our DICE-based integrated assessment model.** Unabated emissions from economic production accumulate in the atmosphere. Together with other greenhouse gases, the emission stock changes the energy balance of the planet, induces various feedback processes, and increases global surface temperature. The bold arrows indicate the relationship changed by the warming feedback (a), carbon sink (b), and damage (c) tipping points described in the main text. Dashed lines indicate the decision variables (controls).

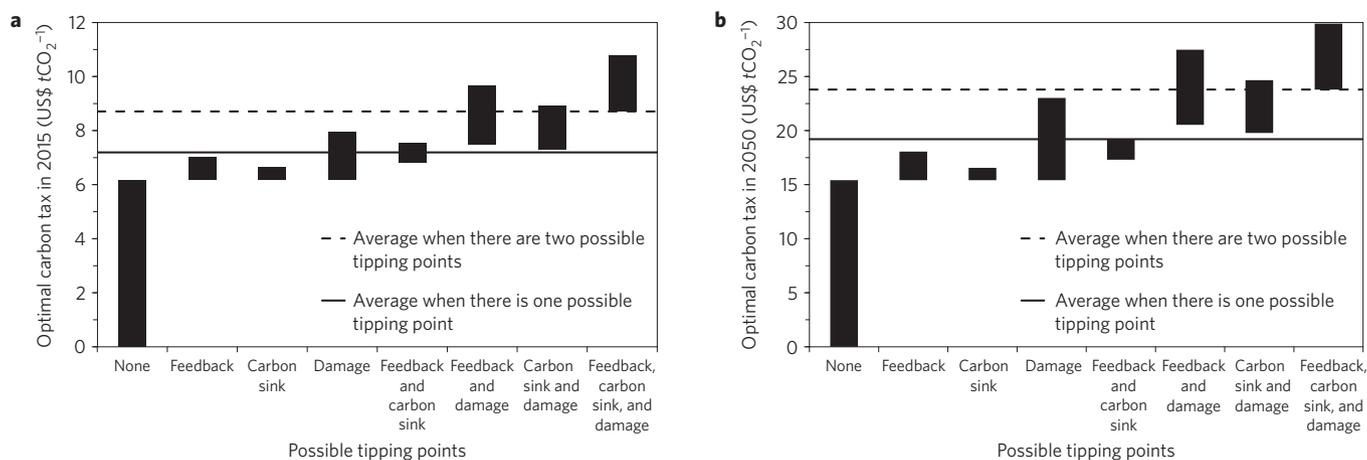
feedback tipping point as increasing climate sensitivity to 5 °C. The second tipping point (labelled b) increases the residence time of atmospheric CO<sub>2</sub>. Warming-induced changes in oceans<sup>23</sup>, soil carbon dynamics<sup>24</sup>, and standing biomass<sup>25</sup> could affect the uptake

of CO<sub>2</sub> from the atmosphere. We represent such a weakening of carbon sinks as reducing the rate of atmospheric CO<sub>2</sub> removal by 50%. These first two tipping points modify the (simplified) climate dynamics of welfare-optimizing integrated assessment models to allow them to jump into a new climate regime.

The third tipping point (labelled c) directly affects the economic damage function. This damage function encapsulates all impacts from warming, including damages from sea-level rise, from habitat loss, and from a weakening of the Atlantic conveyor belt (Gulf Stream). Any of these channels is subject to potential abrupt changes that would cause substantial economic damages. For example, if the West Antarctic or Greenland ice sheets collapsed, sea levels could rise quickly and drastically<sup>26–29</sup>. These higher sea levels would interact with the existing pathways by which warming causes damages. By stressing adaptive capacity, higher sea levels make damages increase faster with warming. We model such a tipping point as changing the DICE damage function from the conventional assumption of a quadratic in temperature to a cubic in temperature: if this tipping point occurs, then doubling anthropogenic warming increases damages eightfold rather than fourfold. Previous models of an economic damage tipping point have represented it as triggering a discontinuous jump in damages<sup>12–14</sup>. Our implementation combines this jump with a shift into a regime of more convex damages, which reflects that tipping points are likely to have more severe consequences when they impact a system that is already stressed by higher temperatures. Such convex damages have been increasingly studied in conventional integrated assessment models that lack tipping points<sup>30,31</sup>. The Supplementary Information contains a detailed discussion of the tipping points and the underlying hazard rate



**Figure 2 | Schematic illustrating the recursive decision problem underlying the optimal policy choice.** The policymaker anticipates that a tipping point might happen; future policymakers learn about the location of temperature thresholds by observing whether past warming triggered a tipping point; future policymakers adjust to the new climate dynamics in the case of tipping; these adjustments and their consumption and welfare effects depend on the climate and capital passed on to the future policymakers; and each of these future policymakers in turn anticipates the actions of policymakers further in the future.



**Figure 3 | The optimal carbon tax in the years 2015 and 2050, assuming no tipping point has yet occurred. a,b,** Optimal tax in 2015 (a) and 2050 (b). The columns plot scenarios without any possible tipping points, scenarios with a single possible tipping point, scenarios with two possible tipping points, and scenarios with three possible tipping points. The top of each bar depicts the optimal carbon tax. The bottom of each bar depicts the average optimal tax when removing one tipping point from the set indicated in the column: it is the no-tipping value when facing only a single tipping point; it is the average of the two single-tipping values when facing two tipping points; and it is the average of the two-tipping values when facing three tipping points. The length of each bar therefore gives the (average) contribution of adding the last tipping point. The horizontal lines show the average carbon tax among the scenarios with only one (solid) or two (dashed) tipping points.

model, which is calibrated to trigger, in expectation, a first tipping point at 2.5 °C of warming relative to 1900.

### Optimal policy

Figure 3 reports the optimal tax on a ton of CO<sub>2</sub> in 2015 and 2050 in settings without any potential tipping points, with only a single potential tipping point, with two potential tipping points, and with all three potential tipping points. The optimal tax in 2015 nearly doubles from US\$6 per ton of CO<sub>2</sub> in the absence of tipping concerns to US\$11 when taking the full range of tipping possibilities and their interactions into account. Note that the optimal carbon tax without tipping points is almost half of that in DICE (ref. 18) because, as discussed above, we remove DICE's *ad hoc* damage adjustment that compensates for not modelling tipping points. It turns out that DICE's *ad hoc* adjustment yields a similar year 2015 carbon tax to the one in our model with interacting tipping points. This coincidence no longer holds in later years and at different temperature levels.

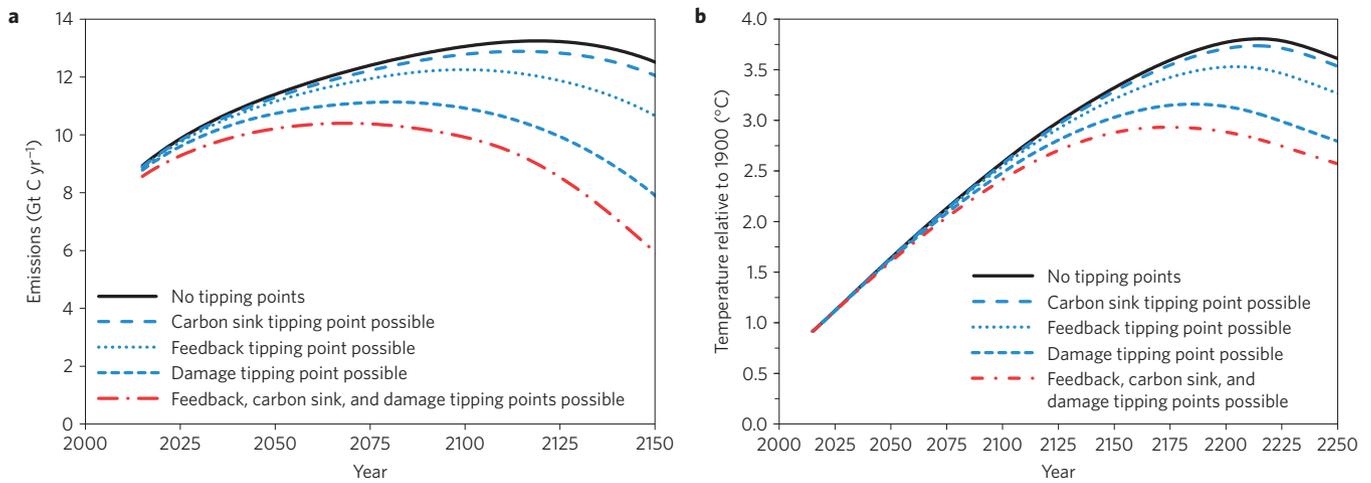
The economic damage tipping point has the strongest individual effect, increasing the optimal emission tax by 30%. The feedback tipping point increases the optimal emission tax by 14%, and the carbon sink tipping point increases it by 8%. Similar findings hold for 2050: the possibility of three interacting tipping points approximately doubles the optimal carbon tax from US\$15 to US\$30 per ton of CO<sub>2</sub>, and the damage tipping point remains the most important one.

Figure 4 presents the optimal emission and temperature trajectories, conditional on not having crossed a threshold by the corresponding date. Again, recognizing the potential for a carbon sink tipping point only slightly reduces optimal emissions and temperature. Recognizing the potential for a feedback tipping point reduces optimal peak annual emissions by 1 Gt C and peak temperature by 0.3 °C. Recognizing the potential for a damage tipping point reduces optimal peak annual emissions by more than 2 Gt C and peak temperature by over 0.6 °C. Anticipating all three potential tipping points reduces optimal peak annual emissions by almost 3 Gt C and reduces the optimal peak temperature from close to 4 °C to just below 3 °C. Anticipating these potential tipping points makes optimal emissions and temperatures peak more than half a century earlier.

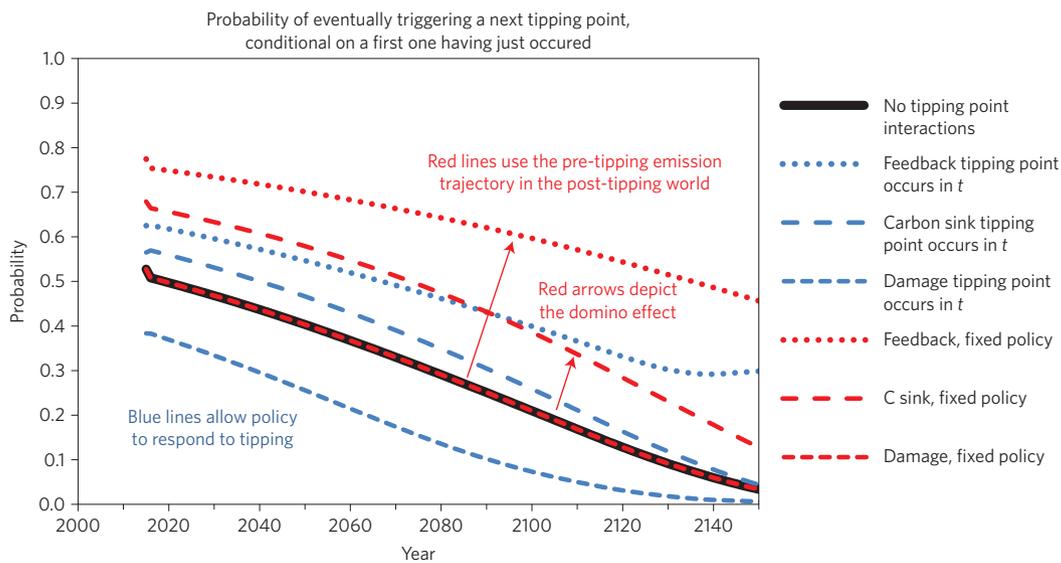
### Tipping point interactions

Our results indicate strong interactions between the different tipping points. A simple analytic theory model cannot determine whether the joint policy effect of multiple tipping points is larger or smaller than the sum of the adjustments derived from individual tipping point models. On the one hand, a given policy intervention simultaneously reduces the probability of several tipping points in the joint model. On the other hand, the greater probability of tipping makes the decision maker prefer to invest (instead of reducing emissions) because the expected benefit from future consumption possibilities is higher. Both the optimal temperature reduction (preceding the peak temperature) and the optimal carbon tax at a given time are larger than the sum of the adjustments obtained from the individual tipping point models. The optimal carbon tax adjustment in the joint model is 50% higher in 2015 (and almost 27% higher in 2050) than suggested by simply adding the tax adjustments of the three individual tipping point models. The primary interaction is between the temperature feedback and the damage tipping points: combining these two tipping points into a single model raises the optimal 2015 tax adjustment by 35% above the sum of the individual tax adjustments. The interaction between the carbon sink and the damage tipping points raises the optimal tax adjustment by 22% over the sum of the adjustments deriving from the individual tipping point models. The interaction between the temperature feedback and carbon sink tipping points has the smallest effect, raising the carbon tax by 6% in 2015 above the sum of the individual adjustments. The ranking of the relative importance of the tipping point interactions persists throughout the century, with absolute adjustments increasing and relative adjustments slightly decreasing.

Figure 5 shows the domino effect underlying these policy interactions. Each line depicts the probability of eventually crossing another threshold if the stated tipping point occurs in period  $t$ . The thick black line depicts this probability in the absence of tipping point interactions or any policy response. The red lines allow tipping points to interact, but they maintain the emission trajectory from the optimal pre-tipping policy. The change from the black line to a red line shows the domino effect in a narrow or physical sense (no policy or economic response). The damage tipping point does not change the likelihood of crossing another temperature threshold



**Figure 4 | Optimal emissions and temperature with a single potential tipping point and with all three potential tipping points. a, b**, Optimal emissions (a) and temperature (b) in settings with a single potential tipping point (blue lines) and with all three potential tipping points (red lines). Both panels assume that the policymaker is aware of potential tipping points, but nature’s draws are such that no tipping point happens. The damage tipping point (by itself) causes the strongest emission reduction. In the present century, the reductions in emissions and temperature in the setting with three potential tipping points are greater than the sum of the three individual reductions in the settings with single tipping points. When there are three potential tipping points, emissions peak half a century earlier and the optimal peak temperature is approximately 1 °C lower than in the case without potential tipping points.

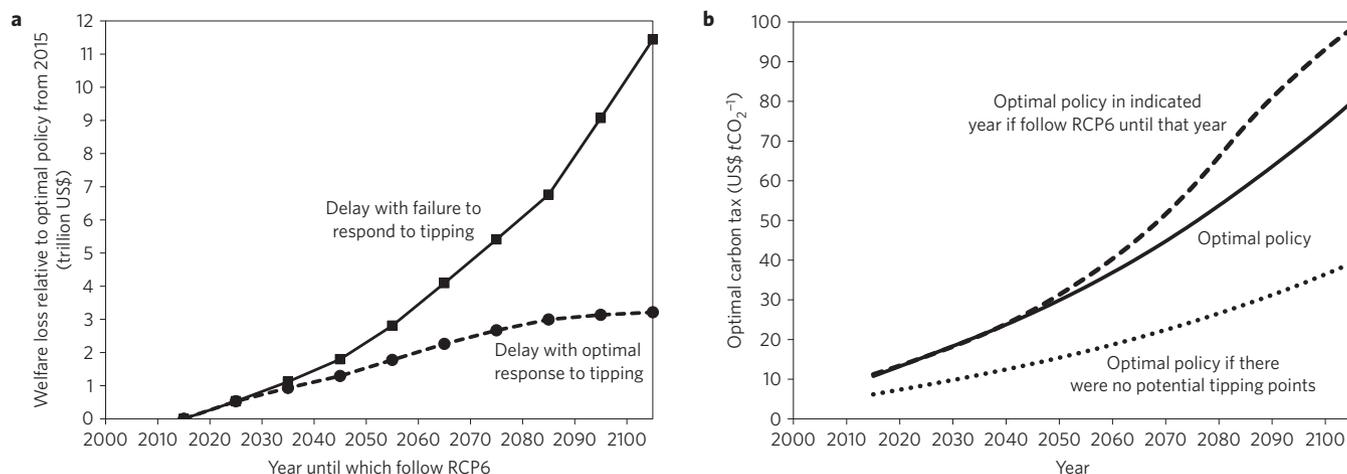


**Figure 5 | The probability of eventually triggering a next tipping point, conditional on having just triggered a first tipping point (out of three potential ones) in a given year.** The black line depicts the case in which tipping points do not interact through system dynamics or policy. The red lines combine the pre-tipping emission trajectory with post-tipping dynamics. The change from the black to the red lines gives each tipping point’s (physical) domino effect. The blue lines allow policy and the economy to respond to having tipped. We find that triggering one of the two climate regime shifts substantially increases the probability of triggering a further tipping point. In contrast, triggering the economic damage tipping point does not increase the probability of triggering another tipping point. The optimal policy response to tipping reduces the domino effect.

because it does not change the climate’s dynamics. Triggering the feedback tipping point induces the strongest domino effect: it approximately doubles the probability of further tipping because any given CO<sub>2</sub> concentration now generates more warming. The carbon sink tipping point also has a domino effect, but it is approximately half as strong as the domino effect of the feedback tipping point.

The blue lines allow the policymaker and the economy to respond to the tipping point that has occurred. Each of the tipping points makes the optimal emission policy more stringent (see Supplementary Figs 2 and 3). The greatest reduction in emissions arises after crossing the damage tipping point, and the smallest reduction in emissions arises after crossing the carbon sink

tipping point. In addition, economic damages respond to the new temperature trajectory and, in the case of the damage tipping point, the new damage regime. The policy response and the economic response both work to reduce the probability of triggering another tipping point by between 10 and 20 percentage points in the present century. Under optimal response to a tipping of the carbon sink and the economic damage regime, the probabilities of further tipping approach zero by the end of the depicted time horizon, whereas in the temperature feedback regime, temperatures would still be increasing and the likelihood of triggering a further tipping point would remain high. The carbon sink tipping point induces a positive net domino effect, but it is small (only 5 percentage points) if the



**Figure 6 | The welfare and policy consequences of delaying optimal climate policy.** **a**, ‘Delay with optimal response to tipping’ depicts the welfare cost from switching to optimal policy only after period  $t$  or after a first tipping point occurs (if that happens before period  $t$ ). ‘Delay with failure to respond to tipping’ depicts the welfare cost from switching to optimal policy only in period  $t$ , regardless of whether tipping points occur before then. **b**, Each year’s optimal carbon tax if the policymaker begins to optimize only in that year, assuming that no tipping point has yet been triggered.

tipping point is triggered this century. The temperature feedback tipping point induces the strongest net domino effect, raising the probability of tipping by 15–20 percentage points. In contrast, the economic damage tipping point, which is the most commonly modelled one, actually reduces the probability of triggering a second tipping point by around 15 percentage points. Nonetheless, we saw in Fig. 3 that the optimal pre-tipping carbon tax is most sensitive to the combination of the feedback tipping point and the economic damage point. The reason is that the economic damage tipping point is the most painful and the strong domino effect of the feedback tipping point works to make the damage tipping point more likely.

### The cost of delaying policy

Our results so far demonstrate the implications of potential tipping points for optimal mitigation policy. But actual emission policy has deviated substantially from model recommendations. Figure 6 analyses the cost of postponing optimal policy in the presence of tipping points. It shows the expected welfare loss from delaying policy (left), as well as the optimal carbon tax that the policymaker would have to set in year  $t$  if beginning to optimize only at that point (right). Before year  $t$ , the emission pathway is defined by the Representative Concentration Pathway scenario that stabilizes total radiative forcing at  $6 \text{ W m}^{-2}$  by 2100 (RCP6; refs 32,33) and by a 22% investment rate. RCP6 initially requires only moderate emission reductions, but eventually undertakes significant emission reductions towards the end of the century (see Supplementary Information). The initial carbon tax is not affected much by a delay of less than 50 years. However, delay is expensive. Suppose policy follows a ‘delay with optimal response to tipping’ scenario: the policymaker follows the exogenous RCP6 emission pathway until 2050 if no tipping occurs, but switches to the optimal policy immediately in case of tipping before. Such a delay costs US\$1.5 trillion in current dollars, approximately the GDP of Australia. This cost is more than 60% greater than the cost of delaying policy until mid-century in the absence of tipping points. If the policymaker follows this scenario not just until mid-century, but until the first (potential) occurrence of a tipping point, the expected cost more than doubles to US\$3.3 trillion, almost the GDP of Germany. Supplementary Fig. 1 shows that approximately a century from today the RCP6 emission trajectory catches up with the optimal emission trajectory (conditional on no tipping point having occurred). Thus, our numbers reflect the cost of delaying (not abandoning) policy even when the decision maker never switches

from RCP6 to the optimal policy (beyond the year 2100, we use the IPCC’s RCP6 extension).

Now suppose that the policymaker follows a ‘delay with failure to respond to tipping’ scenario. First, assume that the policymaker fails to react to the first tipping point that occurs but does switch to the optimal policy in the case of triggering a second tipping point. Then the welfare loss more than doubles again, to US\$7.3 trillion. We can also view this scenario as a proxy for observing a threshold crossing only with major delay. Second, let us assume that the policymaker never switches to the optimal policy, even if all three tipping points occur. Then the welfare loss from following the suboptimal RCP6 emission pathway nearly doubles once more, increasing to US\$13.5 trillion. The social cost of emitting a ton of carbon today in this suboptimal policy scenario increases by 50% to US\$15 per ton of  $\text{CO}_2$  compared to the case of optimal policy (or even to the case of optimal response to a first tipping point). These results demonstrate the value of a reactive policy. They also emphasize the importance of evaluating policy in a framework that anticipates not only tipping possibilities but also how policy responds to future information.

### Discussion

The precise quantitative results from any integrated assessment model are sensitive to a number of assumptions. Economic integrated assessments of optimal policy under uncertainty cannot incorporate the full scientific details of the tipping process. We employed the most widely used integrated assessment model (DICE) for our analysis, retaining all of its assumptions except for our explicit treatment of regime shifts in the climate and their damage implications. Our quantitative results have to be interpreted within this framework. For instance, alternate discount rate calibrations can greatly increase the optimal emission tax in models without tipping points<sup>30,34,35</sup>, and would also increase the effect of tipping points. In addition, our results are sensitive to our assumptions about the support and shape of our Bayesian prior. Ref. 15 contains a discussion of these sensitivities in the simpler setting where only one of the regime shifts can occur. Our model provides initial insight into how interactions among tipping points affect optimal policy. We find that tipping points are of major quantitative importance for climate change policy. The presence of the three potential tipping points in our Bayesian learning model nearly doubles the optimal carbon tax and reduces optimal peak warming by approximately  $1^\circ\text{C}$ . Moreover, delaying optimal policy in the face of irreversible tipping points is costly, even in a standard economic model that

exhibits a possibly conservative damage function and discounts the future utility loss from climate change.

Our results suggest that adding the optimal tax (and temperature) adjustments across separate studies provides a reasonable first-order approximation to the effect of jointly interacting tipping points, although the sum does slightly underestimate the effect on optimal policy. The temperature feedback tipping point poses the greatest danger of triggering further tipping points, even after accounting for the optimal mitigation response to triggering the tipping point. The economic damage tipping point is the most economically painful, immediately reducing output and consumption. The interaction between the temperature feedback and the economic damage tipping point therefore delivers the strongest interaction, raising the carbon tax by 35% over the sum of these two individual tipping points' tax adjustments.

We also demonstrate the value of the ability to detect tipping points even after they have been irreversibly triggered. Tipping points make delaying policy especially expensive, but if detecting a triggered tipping point would spur nations to collectively respond so as to reduce the chance of triggering further tipping points, then the detection capability cuts the cost of delaying policy by 75%. The scientific literature has thus far focused on early warning signals for tipping points<sup>36,37</sup>, but our results call for a complementary focus on fingerprinting already-triggered tipping points.

## Methods

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

Received 31 May 2015; accepted 25 November 2015;  
published online 18 January 2016

## References

- Hansen, J. *et al.* Target atmospheric CO<sub>2</sub>: where should humanity aim? *Open Atmos. Sci. J.* **2**, 217–231 (2008).
- Ramanathan, V. & Feng, Y. On avoiding dangerous anthropogenic interference with the climate system: formidable challenges ahead. *Proc. Natl Acad. Sci. USA* **105**, 14245–14250 (2008).
- Rockström, J. *et al.* A safe operating space for humanity. *Nature* **461**, 472–475 (2009).
- National Research Council *Abrupt Impacts of Climate Change: Anticipating Surprises* (National Academies Press, 2013).
- Lenton, T. M. *et al.* Tipping elements in the Earth's climate system. *Proc. Natl Acad. Sci. USA* **105**, 1786–1793 (2008).
- 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).
- Levermann, A. *et al.* Potential climatic transitions with profound impact on Europe. *Climatic Change* **110**, 845–878 (2012).
- Lenton, T. M. & Ciscar, J.-C. Integrating tipping points into climate impact assessments. *Climatic Change* **117**, 585–597 (2013).
- Kopits, E., Marten, A. & Wolverton, A. Incorporating 'catastrophic' climate change into policy analysis. *Clim. Policy* **14**, 637–664 (2014).
- Rocha, J. C., Peterson, G. D. & Biggs, R. Regime shifts in the Anthropocene: drivers, risks, and resilience. *PLoS ONE* **10**, e0134639 (2015).
- Gjerde, J., Grepperud, S. & Kverndokk, S. Optimal climate policy under the possibility of a catastrophe. *Resour. Energy Econ.* **21**, 289–317 (1999).
- Keller, K., Bolker, B. M. & Bradford, D. F. Uncertain climate thresholds and optimal economic growth. *J. Environ. Econ. Manag.* **48**, 723–741 (2004).
- van der Ploeg, F. & de Zeeuw, A. *Climate Policy and Catastrophic Change: Be Prepared and Avert Risk* OxCarre Working Paper 118 (Oxford Centre for the Analysis of Resource Rich Economies, University of Oxford, 2013).
- 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).
- Lemoine, D. & Traeger, C. Watch your step: optimal policy in a tipping climate. *Am. Econ. J. Econ. Policy* **6**, 137–166 (2014).
- Traeger, C. A 4-stated DICE: Quantitatively addressing uncertainty effects in climate change. *Environ. Resour. Econ.* **59**, 1–37 (2014).
- Nordhaus, W. D. An optimal transition path for controlling greenhouse gases. *Science* **258**, 1315–1319 (1992).
- Nordhaus, W. D. *A Question of Balance: Weighing the Options on Global Warming Policies* (Yale Univ. Press, 2008).
- Zimov, S. A., Schuur, E. A. G. & Chapin, F. S. Permafrost and the global carbon budget. *Science* **312**, 1612–1613 (2006).
- Hall, D. C. & Behl, R. J. Integrating economic analysis and the science of climate instability. *Ecol. Econ.* **57**, 442–465 (2006).
- Archer, D. Methane hydrate stability and anthropogenic climate change. *Biogeosciences* **4**, 521–544 (2007).
- Schaefer, K., Zhang, T., Bruhwiler, L. & Barrett, A. P. Amount and timing of permafrost carbon release in response to climate warming. *Tellus B* **63**, 165–180 (2011).
- Le Qur, C. *et al.* Saturation of the Southern Ocean CO<sub>2</sub> sink due to recent climate change. *Science* **316**, 1735–1738 (2007).
- Eglin, T. *et al.* Historical and future perspectives of global soil carbon response to climate and land-use changes. *Tellus B* **62**, 700–718 (2010).
- Huntingford, C. *et al.* Towards quantifying uncertainty in predictions of Amazon 'dieback'. *Phil. Trans. R. Soc. B* **363**, 1857–1864 (2008).
- Oppenheimer, M. Global warming and the stability of the West Antarctic Ice Sheet. *Nature* **393**, 325–332 (1998).
- Oppenheimer, M. & Alley, R. B. The West Antarctic Ice Sheet and long term climate policy. *Climatic Change* **64**, 1–10 (2004).
- Vaughan, D. G. West Antarctic Ice Sheet collapse—the fall and rise of a paradigm. *Climatic Change* **91**, 65–79 (2008).
- Notz, D. The future of ice sheets and sea ice: between reversible retreat and unstoppable loss. *Proc. Natl Acad. Sci. USA* **106**, 20590–20595 (2009).
- Stern, N. *The Economics of Climate Change: The Stern Review* (Cambridge Univ. Press, 2007).
- Ackerman, F., Stanton, E. A. & Bueno, R. Fat tails, exponents, extreme uncertainty: simulating catastrophe in DICE. *Ecol. Econ.* **69**, 1657–1665 (2010).
- Moss, R. H. *et al.* The next generation of scenarios for climate change research and assessment. *Nature* **463**, 747–756 (2010).
- Masui, T. *et al.* An emission pathway for stabilization at 6 W m<sup>2</sup> radiative forcing. *Climatic Change* **109**, 59–76 (2011).
- Schneider, M. C., Winkler, R. & Traeger, C. P. Trading off generations: equity, discounting, and climate change. *Eur. Econ. Rev.* **56**, 1621–1644 (2012).
- Crost, B. & Traeger, C. P. Optimal CO<sub>2</sub> mitigation under damage risk valuation. *Nature Clim. Change* **4**, 631–636 (2014).
- Scheffer, M. *et al.* Early-warning signals for critical transitions. *Nature* **461**, 53–59 (2009).
- Scheffer, M. *et al.* Anticipating critical transitions. *Science* **338**, 344–348 (2012).

## Acknowledgements

C.P.T. gratefully acknowledges support by the National Science Foundation through the Network for Sustainable Climate Risk Management GEO-1240507.

## Author contributions

Both authors contributed in equal parts to the conceptual and numerical design of the study, as well as to the interpretation and presentation of the results. D.L. coded the novel tipping point interaction part of the model.

## 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](http://www.nature.com/reprints). Correspondence and requests for materials should be addressed to C.P.T.

## Competing financial interests

The authors declare no competing financial interests.

## Methods

Here we summarize the novel components of the model and our solution method. The Supplementary Information contains additional background and results, provides the complete model equations and parameterization, and details the solution method. Ref. 15 provides the Matlab code extended for this analysis.

We reformulate the DICE-2007 model of ref. 18 into a recursive dynamic programming problem, following refs 15,16,38. To mitigate the curse of dimensionality in dynamic programming, we reduce the state space of the climate system's representation. Ref. 16 shows that this simplification does not impair the model's ability to reproduce the DICE model's climate response. We follow ref. 15 in using an infinite planning horizon, unlike other approaches to modelling tipping points<sup>12,14</sup>.

We make one substantive change to the DICE-2007 parameterization. Ref. 18 adjusts a coefficient in the damage function to compensate for not explicitly modelling climate catastrophes and tipping points. This damage adjustment contributes about half of DICE's damages at 1 °C of warming. We eliminate this adjustment because we now explicitly model tipping points.

We introduce tipping points as regime shifts following ref. 15. That paper provides more details about the modelling of the probability of tipping and of learning. In each period, the climate system might irreversibly trigger any of the possible tipping regimes. In a given state of the world, the hazard of crossing a threshold is identically and independently distributed for each tipping point, and this hazard is a function of the temperature increase. The thresholds are identically and independently distributed because there is no particular knowledge that one is more likely than another. However, the probability of different tipping sequences is not symmetric because the hazards interact through the endogenous system dynamics and policy responses. The Bayesian policymaker assesses the probability of tipping by updating their previous beliefs based on whether they have just observed a tipping point or not. Ref. 39 shows that it is critical to model learning when modelling thresholds. Further, modelling learning means that our policymaker might avoid tipping if they eventually stop temperatures from increasing, whereas settings without learning can imply that tipping will eventually occur with probability one (ref. 14). Ref. 40 gives an overview of different uncertainties in climate change and how they have been treated in the literature (see their Table 2.2).

We calibrate the threshold distributions so that the policymaker with year 2005 information expects a first temperature threshold at 2.5 °C of warming relative to 1900—that is, 2.5 °C is the expectation of the random variable 'first threshold crossed' in the setting with three possible tipping thresholds. This requirement implies an upper bound  $\bar{T}$  of 7.8 °C, which is among the highest values explored in the sensitivity analysis for the single-tipping runs in ref. 15. An expected trigger of 2.5 °C is consistent with the political 2 °C limits for avoiding dangerous anthropogenic interference. Further, in the 'burning embers' diagram<sup>41</sup>, the yellow

(medium-risk) shading for the 'risk of large-scale discontinuities' begins around 1.6 °C of warming relative to 1900, and the red (high-risk) shading begins around 3.1 °C of warming relative to 1900.

We solve the tipping scenarios recursively, starting with a world where all tipping thresholds have been crossed. We solve that world's Bellman equation by value function iteration. We approximate the value function by combining collocation methods with a Chebyshev tensor basis containing  $10^4$  basis functions<sup>42–44</sup>. Each iteration uses the Knitro solver to optimize abatement and consumption decisions at each Chebyshev node. Once we have successfully approximated that value function, we use it to define the post-tipping continuation value in each scenario where all but one tipping point has been crossed. We then approximate the value functions for these scenarios by following the same procedure. Once we have successfully approximated these value functions, they provide additional continuation values in scenarios where all but two tipping points have been crossed. We then approximate these scenarios' value functions. We proceed in this fashion until we approximate the value function for a scenario in which no tipping point has yet occurred. At this point, we have approximated the value function for any possible tipping regime. Solving a model with three interacting tipping points requires the solution of eight nested dynamic programming problems representing the different combinations in which tipping points might possibly occur.

We obtain each year's optimal carbon tax and emission policy by maximizing the Bellman equation using the solved continuation values. The optimal carbon tax is composed of the marginal change of the value function in CO<sub>2</sub> in the current regime, the (hazard rate-weighted) marginal change of the value function in CO<sub>2</sub> in the possible tipping regimes, and the marginal impact of emissions on the hazard rate of tipping weighted by the resulting value change<sup>15</sup>.

## References

38. Kelly, D. L. & Kolstad, C. D. Bayesian learning, growth, and pollution. *J. Econ. Dynam. Control* **23**, 491–518 (1999).
39. Lorenz, A., Schmidt, M. G. W., Kriegler, E. & Held, H. Anticipating climate threshold damages. *Environ. Model. Assess.* **17**, 163–175 (2012).
40. Kunreuther, H. *et al.* in *Climate Change 2014: Mitigation of Climate Change* (eds Edenhofer, O. *et al.*) 151–205 (IPCC, Cambridge Univ. Press, 2014).
41. Smith, J. B. *et al.* Assessing dangerous climate change through an update of the Intergovernmental Panel on Climate Change (IPCC) "reasons for concern". *Proc. Natl Acad. Sci. USA* **106**, 4133–4137 (2009).
42. Judd, K. L. Projection methods for solving aggregate growth models. *J. Econ. Theory* **58**, 410–452 (1992).
43. Judd, K. L. *Numerical Methods in Economics* (MIT Press, 1998).
44. Miranda, M. J. & Fackler, P. L. *Applied Computational Economics and Finance* (MIT Press, 2002).