

Optimal CO₂ mitigation under damage risk valuation

Benjamin Crost¹ and Christian P. Traeger^{2*}

The current generation has to set mitigation policy under uncertainty about the economic consequences of climate change. This uncertainty governs both the level of damages for a given level of warming, and the steepness of the increase in damage per warming degree. Our model of climate and the economy is a stochastic version of a model employed in assessing the US Social Cost of Carbon (DICE). We compute the optimal carbon taxes and CO₂ abatement levels that maximize welfare from economic consumption over time under different risk states. In accordance with recent developments in finance, we separate preferences about time and risk to improve the model's calibration of welfare to observed market interest. We show that introducing the modern asset pricing framework doubles optimal abatement and carbon taxation. Uncertainty over the level of damages at a given temperature increase can result in a slight increase of optimal emissions as compared to using expected damages. In contrast, uncertainty governing the steepness of the damage increase in temperature results in a substantially higher level of optimal mitigation.

The DICE integrated assessment model^{1,2} couples a growing global economy to a simple climate model. The economy produces emissions that accumulate in the atmosphere, change the radiative forcing, and warm the planet's surface. This warming feeds back into economic production and consumption (goods and services humans care about). The climate–economy interactions are nonlinear and delayed. Integrated assessment models such as DICE inform us about the long-term economic loss resulting from current carbon emissions and help us evaluate climate policy. Many of the interactions within and between climate and the economy are uncertain and an increasing number of studies simulate the consequences of given policies under uncertainty^{2–10}. We follow Kelly and Kolstad¹¹, Keller *et al.*¹² and Leach¹³ in building a stochastic integrated assessment model that evaluates the optimal policy response to uncertainty. Optimality means that resources within and across periods are distributed to maximize the expected stream of global welfare. Our contribution is twofold. First, our optimal trade-off among consumption, emissions, and capital investment treats the damage parameters governing the relation between climate change and economic impact as stochastic. Second, we use a more sophisticated approach to evaluate the uncertain impact of climate change on economic and human welfare.

Uncertainty evaluation

Nordhaus² estimates the DICE damages as a function of the global average temperature increase T_t above the level prevailing in 1900. In these estimates, the fraction of global economic production lost to climate change is

$$D(T_t) = b_1 T_t^{b_2} \quad (1)$$

DICE's damage function and similar variations are widespread in the integrated assessments of climate change. The damage coefficient b_1 captures the level of damages at a 1 °C warming (damage level). The damage exponent b_2 captures the steepness of the damage increase as temperature rises (damage convexity).

A damage exponent $b_2 = 2$ (or 3) implies that damages for a 3 °C warming increase to $3^2 = 9$ (or $3^3 = 27$) times the damages of a 1 °C warming. Nordhaus² estimates that a 1 °C warming reduces global production by 0.28% ($b_1 = 0.0028$), and that a 3 °C warming reduces production by 2.6% ($b_2 = 2$). These estimates are based on collecting and extrapolating damage data for 0 °C, 2.5 °C, and 6 °C of global warming across different regions of the world. These estimates contain major uncertainties due to the lack of observation of warming above a one degree change, limited data with which to assess damages around the world, and various identification issues summarized by Hanemann¹⁴. We analyse uncertainty about both the level of damages, characterized by b_1 , and the convexity of damages, characterized by b_2 .

We optimize policy for a 'low'-uncertainty and a 'high'-uncertainty scenario. Our 'low'-uncertainty scenario is based on a sensitivity study by Nordhaus² governing the damage coefficient b_1 . It implies a 1.4% and a 3.7% production loss for minus/plus one standard deviation at a 3 °C warming (as opposed to a 2.6% loss without uncertainty). We assess both the 'low'-uncertainty and the 'high'-uncertainty scenarios for both damage coefficient uncertainty and damage exponent uncertainty. We make the two different types of uncertainty comparable by calibrating the damages for minus/plus one standard deviation at a 3 °C warming to the same, cited values. Our 'high'-uncertainty scenario is based on a model survey by Tol¹⁵ and his estimate of damage exponent uncertainty b_2 . It induces a global production loss of 0.3% and 4.8% at a 3 °C warming for minus/plus one standard deviation. For more details, we refer to Methods.

Integrated assessment models either evaluate climate change based on normative principles, or they connect the welfare function to people's actual preferences by calibrating a representative agent's preference parameters to observed interest rates. We follow this second, widespread, observation-based approach. It ubiquitously applies the standard economic model (discounted expected utility model), and we base our first set of results on the DICE calibration by Nordhaus² of this standard model. We label the corresponding

¹Department of Economics, University of Colorado Denver, Campus Box 181, Denver, Colorado 80217-3364, USA, ²Department of Agricultural & Resource Economics, 207 Giannini Hall #3310, University of California Berkeley, Berkeley, California 94720-3310, USA. *e-mail: traeger@berkeley.edu

results ‘scenario N’, for Nordhaus. Two important dimensions of preferences and welfare are risk aversion and the desire to smooth consumption over time. Risk aversion captures that economic agents prefer a given consumption level with certainty over a coin-toss lottery that increases consumption if heads comes up, but decreases consumption if tails comes up (yielding the same consumption level as under certainty in expectation). Intertemporal consumption smoothing captures that economic agents prefer a consumption trajectory where they consume evenly over time to a trajectory where they consume more in one period and less in another (while consuming the same on average along both consumption trajectories). The standard model assumes that risk aversion coincides with the desire to smooth consumption over time, and calibrates a joint parameter for both preference characteristics.

Our comprehensive evaluation of uncertain climate damages relaxes the standard model’s assumption that risk aversion coincides with the desire to smooth consumption over time. First, observed preferences do not support this assumption^{16–21}. Second, risk and time are different dimensions, and a distinct aggregation of welfare across these different dimensions is fully rational^{22–24}. A leading approach in the recent finance literature attributes two famous asset pricing puzzles to the standard model’s failure to distinguish the two dimensions of welfare. The so-called risk-free rate and equity premium puzzles state that calibrating the standard model to asset pricing data results either in risk premia that are too low, or in a (risk-free) consumption discount rate that is too high, or in both. Evaluating climate change with such a model implies that we take too much risk, or we pay too little attention to the long-run future, or both—a theoretic argument with important implication for climate policy^{25–27}.

A breakthrough in explaining these and several other asset pricing phenomena is the long-run risk model of Bansal and Yaron¹⁷. Building on earlier work^{22,23,28,29}, the authors distinguish risk aversion from the desire to smooth consumption over time (and account for small, persistent consumption shocks). This more comprehensive asset pricing framework finds that agents are significantly more averse to risk than to intertemporal consumption fluctuations^{16–21}. Economic agents prefer to sacrifice consumption in one period and get it back with certainty in a different period over a gamble where they either end up in a world with higher consumption or a world with consumption sacrifices. We base our second set of runs on this more comprehensive evaluation framework and label the results ‘scenario D’, abbreviating the fact that we disentangle risk aversion from the propensity to smooth consumption over time. The welfare function underlying this second set of runs calibrates much more closely to observed market data than does the Nordhaus calibration of DICE, which uses a joint risk aversion and consumption smoothing parameter to reproduce an average market return.

The optimal policy in any given period not only has to account for all possible future evolutions of the climate and the economy, but it also has to anticipate the optimal policy decisions that future policy makers will take, conditional on the actual damages that are realized in the meantime. In Crost and Traeger³⁰ we illustrate that even the sign of how uncertainty affects optimal policy can be wrong when separately optimizing and averaging Monte Carlo simulations. We therefore solve simultaneously for the optimal policies in all periods, conditional on all possible system states. Whereas the optimal policy in the present is uniquely determined by welfare maximization, the optimal policies in the future depend on the realization of damages up to that period. Instead of presenting multidimensional control rules that reflect the optimal future policy for every state of the world, we present trajectories assuming that nature happens to draw expected values in every period. Therefore, the differences between our graphs depicting certainty and uncertainty directly reflect the consequences of the decision maker’s awareness of

uncertainty, and they do not rely on differences in actual damage realizations. It turns out that these ‘expected draw’ trajectories for the policy variables are virtually indistinguishable from the median and the mean paths of 10,000 truly stochastic runs simulated in Supplementary Section A.4.

Results

Figure 1 compares the optimal carbon tax, abatement rate, emission level, and temperature trajectory for the different scenarios. The line types distinguish between the type of uncertainty, and the line colours distinguish the underlying modelling framework: the trajectories that are optimal in the standard model (N) are in blue, and the trajectories that are optimal under the comprehensive preference model (D) are in green. Table 1 states the numbers for optimal present, mid-century, and end-of-century policy, as well as the optimal peak levels and years of carbon concentration and temperature.

Our first result compares the optimal policies and climate trajectories resulting from a standard evaluation (N) to trajectories resulting from the comprehensive evaluation framework (D). Using the comprehensive evaluation framework, the optimal carbon tax more than doubles, to values averaging around US\$130 per tonne of carbon in the present and increasing to around US\$500 by the end of the century. Similarly, the optimal present abatement rate almost doubles in scenario D as compared to scenario N, from around 17% to around 30% of business-as-usual emissions. The optimal peak temperature drops by 1 °C and the optimal carbon concentration drops by 120–160 ppm, to significantly below a doubling of pre-industrial levels. Absolute emissions in the present should be approximately 1 gigatonne of carbon (GtC) lower under the comprehensive evaluation. Moreover, scenario D implies that emissions by the end of the century should fall by a further 1.5–2 GtC. In contrast, optimality in the standard model N results in an increase of emissions by a further 1–1.5 GtC by the end of the century. Figure 1 shows that full abatement (zero industrial CO₂ emissions) is optimal about half a century earlier under the comprehensive evaluation approach, and Table 1 states the same result for peak carbon concentration and peak temperature.

The intuition for our first finding rests on the reduction of the consumption smoothing parameter (from $\eta = 2$ in the original DICE model to $\eta = 2/3$ in the disentangling preference framework). This reduction is a consequence of the observed risk-free rate being significantly lower than the 5.5% average interest rate to which the original DICE model is calibrated. We briefly refer to Nordhaus³¹ to avoid potential misunderstandings. Here, the author points out that the choice of the consumption smoothing parameter η does not matter as long as the rate of pure time preference is chosen so that the overall consumption discount rate is 5.5%. Pure time preference is a measure of intrinsic impatience. However, estimating asset returns with the disentangled model implies a reduction of the consumption smoothing parameter η without increasing the impatience parameter. The approach better captures both the lower risk-free discount rate and the higher, risky rate. In fact, many of the cited estimates of the disentangled model imply an even lower rate of pure time preference than the 1.5% used in DICE-2007.

We note that our observation-based calibration implies a similar if not slightly higher optimal carbon tax as compared to the normatively motivated evaluation in the Stern Review³². Stern suggests a higher social cost of carbon, exceeding US\$200 per tonne C, along a business-as-usual trajectory. However, as Stern explains, this social cost of carbon is not the optimal carbon tax; it only specifies the damage done by the last tonne of carbon emitted in a world without a carbon policy. Stern cites a cost of carbon of US\$110 for the case where CO₂ concentrations are stabilized at 550 ppm, and a cost of US\$92 if CO₂ concentrations were stabilized at 450 ppm. The Stern Review does not assess the optimal carbon tax, which

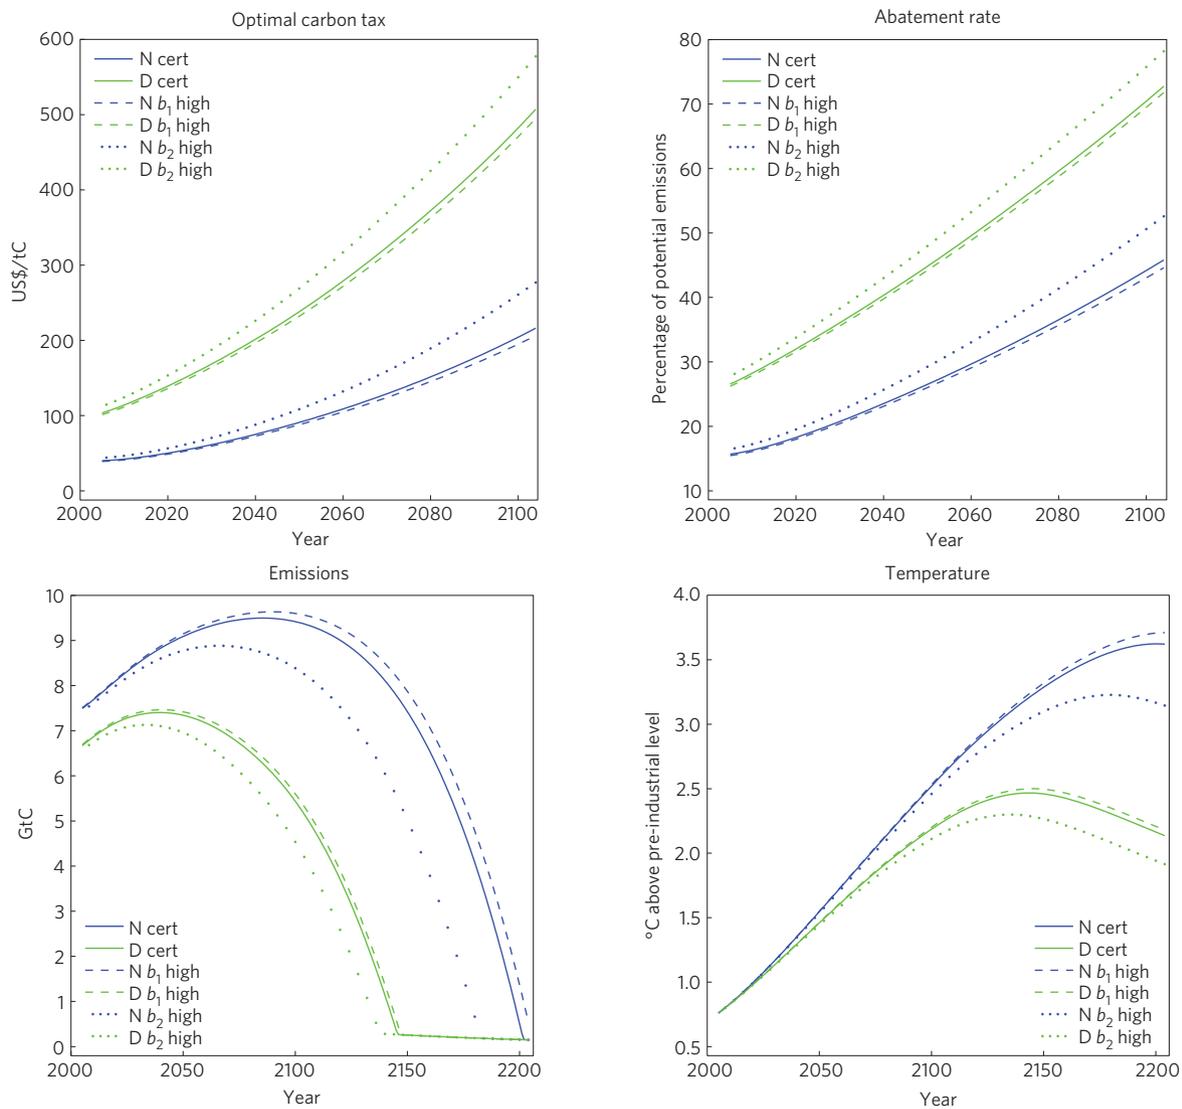


Figure 1 | The optimal carbon tax in US\$ per tonne of carbon and the abatement rate as a percentage of business-as-usual emissions (top, 100 years), as well as the CO₂ emissions from fossil fuel use and the temperature trajectories (bottom, 200 years), for different uncertainty specifications and evaluation frameworks. The blue lines represent evaluations using the standard model (scenario N), which equates risk aversion with the desire to smooth consumption over time ($\eta = RRA = 2$ as in DICE-2007). The green lines use a more comprehensive assessment (scenario D) that disentangles risk aversion from consumption smoothing over time ($\eta = 2/3$ and $RRA = 9.5$). It follows the recent finance literature calibrating to both the low risk-free interest (and consumption discount) rate and the relatively high-risk premia observed in the market. ‘cert’ denotes assessment under certainty (solid lines), ‘b₁ high’ introduces uncertainty over the damage level (dashed lines), whereas ‘b₂ high’ introduces uncertainty over the damage convexity (dotted lines). Both the optimal carbon tax and the optimal abatement rate approximately double under the state-of-the-art evaluation approach (D), and fossil fuels are phased out approximately half a century earlier than in scenario (N). Uncertainty over the damage level slightly reduces optimal mitigation, whereas uncertainty over the damage convexity results in a substantially higher optimal mitigation level.

balances the costs from carbon emissions against the costs from abating carbon. Our comprehensive preference evaluation finds that the optimal present-day carbon tax should be above US\$110 and should reduce emissions below 550 ppm.

The second result compares the effects of uncertainty about the damage level with the implications of uncertainty about the damage convexity—that is, the steepness of the damage increase in global warming. The dashed lines show that uncertainty about the level of damages at a 1 °C warming, measured by the parameter b_1 , slightly reduces the optimal abatement level and carbon tax. In contrast, uncertainty about the damage convexity, measured by the parameter b_2 , increases the optimal abatement level and carbon tax (dotted lines). This finding is independent of whether we use the standard model or the comprehensive evaluation approach. Moreover, the

abatement increase under uncertainty over the damage convexity is significantly stronger than the abatement reduction under damage level uncertainty. Table 1 also reports the optimal policy for the case of joint uncertainty over b_1 and b_2 ; the result is a clear strengthening of optimal mitigation policy. Figure 1 shows only the lines for high uncertainty, whereas Table 1 and Supplementary Fig. 3 also state our findings for the case of low as well as joint uncertainty.

Under the standard model’s evaluation, absolute emissions differ across the different uncertainty specifications by 0.1Gt in the present and 1 Gt at the end of the century; the peak temperature differs by 0.5 °C, the peak concentration by approximately 50 ppm, and the optimal carbon tax by US\$7 in the present and US\$60 by the end of the century. Note that low damage level uncertainty (N, b_1) is the type of damage uncertainty that Nordhaus² uses in his

Table 1 | Optimal policy variables and the peak carbon and temperature levels for different uncertainties and evaluation frameworks.

N	Tax (US\$/tC)			Abatement rate (%)			Emissions (GtC)			Peak carbon		Peak temperature	
	2015	2050	2100	2015	2050	2100	2015	2050	2100	ppm	Year	°C >1900	Year
Certainty	46	91	204	17.2	26.5	44.2	7.9	9.1	9.4	682	2174	3.6	2201
b_1 low	45	89	200	17.0	26.2	43.6	7.9	9.1	9.5	688	2176	3.7	2203
b_1 high	44	88	195	17.0	26.0	43.0	7.9	9.1	9.6	695	2178	3.7	2205
b_2 low	48	99	233	17.6	27.8	47.5	7.9	8.9	8.9	650	2163	3.4	2189
b_2 high	51	108	261	18.3	29.2	50.6	7.8	8.8	8.4	625	2155	3.2	2180
b_1 and b_2 low	47	98	227	17.6	27.6	46.8	7.9	9.0	9.0	655	2165	3.4	2191
b_1 and b_2 high	49	104	247	18.0	28.6	49.1	7.8	8.8	8.6	636	2158	3.3	2184
D													
Certainty	126	238	482	30.1	44.8	70.4	7.0	7.3	5.4	527	2120	2.5	2144
b_1 low	124	234	475	29.8	44.4	69.9	7.1	7.4	5.5	529	2121	2.5	2145
b_1 high	123	232	471	29.7	44.2	69.5	7.1	7.4	5.6	531	2122	2.5	2145
b_2 low	131	253	516	30.8	46.4	73.1	7.0	7.2	5.0	516	2115	2.4	2139
b_2 high	138	269	550	31.7	48.0	75.7	6.9	7.0	4.5	507	2111	2.3	2135
b_1 and b_2 low	130	250	507	30.7	46.0	72.4	7.0	7.2	5.1	519	2117	2.4	2141
b_1 and b_2 high	135	262	531	31.3	47.2	74.3	6.9	7.0	4.8	511	2114	2.3	2138

Scenario N uses the standard economic model ($\eta = RRA = 2$ as in DICE-2007), whereas scenario D distinguishes risk aversion from consumption smoothing over time ($\eta = 2/3$ and $RRA = 9.5$). Specification b_1 introduces uncertainty over the damage level, whereas specification b_2 introduces uncertainty over the damage convexity. Low and high refer to low and high uncertainty.

Monte Carlo experiment averaging deterministic DICE runs. In the comprehensive evaluation approach, the absolute policy differences across the different uncertainty specifications are larger, but, in relative terms they are similar to the uncertainty effects observed in the standard model. The present optimal tax decreases around 2% under damage level uncertainty and increases 5% or 10% for low or high uncertainty about the damage convexity. The case of joint, high uncertainty raises the optimal carbon tax 8% over its deterministic level in both evaluation frameworks. Given that peak temperature is already 1 °C lower in the comprehensive evaluation framework, the change in the carbon tax has a smaller impact on peak temperature: it drops a further 0.2 °C (0.1 °C) under high (low) uncertainty.

The intuition for our second finding is that uncertainty about the damage level affects production linearly, whereas uncertainty about the steepness of the damage increase has a nonlinear effect that emphasizes high damage realizations. The slight negativity of the mitigation impact of damage level uncertainty results from the precise way that Nordhaus² translates his estimated damage function equation (1) into the model equations. In every period, the world production net of climate damages, Y_t^{net} , derives from gross production Y_t^* as

$$Y_t^{\text{net}} = \frac{Y_t^*}{1 + D(T_t)} = \frac{Y_t^*}{1 + b_1 T_t^{b_2}} \quad (2)$$

The uncertainty over b_1 corresponds to a linear variation in the denominator of equation (2). The function characterizing net output is therefore convex in b_1 , and expected output under uncertainty over b_1 is higher than the output using the expected coefficient (Jensen's inequality). Thus, on average, the world looks a little better under uncertainty about the damage level at a given temperature increase. The damage convexity parameter b_2 enters in the denominator of equation (2) as the exponent of temperature. A straightforward calculation shows that the resulting transformation of b_2 into net output is concave in the relevant temperature region, and expected output is lower than the output for the expected coefficient. On average, the world looks worse under uncertainty about the steepness of the damage function.

The preceding paragraph gives a first-order intuition for the results. We briefly note that the actual analytic mechanism driving optimal mitigation is more complicated. First, the economic cost

of carbon depends on the damage caused by the last emitted tonne of carbon, and its effects on net production through temperature increase. Supplementary Section A.2 shows that our convexity reasoning also holds for the more precise marginal reasoning. Second, temperatures are a nonlinear, delayed response to emissions. Third, optimal policy maximizes expected welfare over an infinite time horizon, not economic production in a given period. Welfare depends on consumption, and future production in equation (2) depends on the endogenous temperature and production levels. The following experiment, however, underpins that the nonlinearity of equation (2) indeed drives the quantitative results.

In calibrating DICE, Nordhaus² estimates damages as the percentage of gross production that is lost because of climate change. A damage formulation of the form $Y_t^{\text{net}} = (1 - D(T_t))Y_t^* = (1 - b_1 T_t^{b_2})Y_t^*$ instead of equation (2) would more appropriately translate this estimation into the model's equations. Supplementary Section A.2 presents the optimal policies in such a model, where damages truly represent the fraction of world production lost because of climate change. As we expect from our discussion in the preceding paragraph, we find that the negative impact of damage level uncertainty on optimal mitigation vanishes under this linear reformulation of the damage equation. In defence of Nordhaus's reformulation, we note that placing damages in the denominator, relying on the approximation that for small damages $1/(1 + \epsilon) \approx 1 - \epsilon$, his equation (2) gains the convenient property that output converges to zero as temperatures go to infinity.

Concluding remarks

We found that stochastic damages affect evaluation mostly through the nonlinear impact on production. Supplementary Section A.3 shows that, indeed, the consumption-based risk aversion parameter plays only a minor direct role in evaluating uncertainty of climate damages. The reason lies in the assumption of strong economic growth underlying DICE and most other integrated assessment models, in combination with the common damage formulation that affects production levels more moderately as compared to economic growth. Complementing the present analysis, Jensen and Traeger³³ analyse the effect of growth uncertainty on optimal climate policy. Our present finding regarding the relative importance of the

consumption smoothing parameter and the relative unimportance of the risk aversion parameter has an immediate consequence for the integrated assessments of climate change in general. Almost all large scale integrated assessment models are deterministic. Our recursive dynamic programming solution is computationally involved, even for the relatively simple DICE model. Computation time is exponential in the number of state variables, making the approach infeasible for large integrated assessment models, at least at present^{6,34–36}. Our simulations show that observation-based deterministic models should calibrate the entangled preference parameter η to the (significantly lower) consumption smoothing preference instead of calibrating it to the (much higher) risk aversion parameter or a mixture of the two (such as $\eta = 2$ in DICE-2007).

We have shown the impact of damage uncertainty and a state-of-the-art evaluation framework on optimal mitigation policy. DICE's precise impact forecasts have changed significantly with every revision since 1994. We calibrated our baseline to the DICE-2007 model, but emphasize that the effects shown here result from generic features shared by all versions of DICE and many if not most integrated assessment models. Our contribution focuses on introducing damage risk and a comprehensive evaluation framework, and our policy maker reacts optimally to the stochastic evolution of the climate–economy. A more sophisticated decision maker will also anticipate structural learning that reduces (anticipated) randomness over time as more observations and better models become available. Finally, our approach follows the majority of integrated climate change assessments by using a welfare function that relies on a calibration to observed market outcomes. Climate change will affect future generations. Therefore, several scholars argue that the welfare function to evaluate climate change should be based on normative reasoning rather than observation^{32,37,38}. We leave a normative application of the disentangled approach to future research.

Methods

The policy maker in our model is uncertain about the damage realization every year over an infinite time horizon. An optimal policy choice that accounts for future reactions to the stochastic evolution requires a recursive dynamic programming formulation, following Kelly and Kolstad¹¹. We map the infinite time horizon onto the unit interval and reformulate the DICE model in effective labour units so that capital, consumption and production converge within a bounded numeric support, following Traeger³⁹. We step down DICE's time step to one year, interpolating DICE's exogenous processes in continuous time. The exogenous processes in DICE include the non-industrial CO₂ emissions from land-use change and forestry as well as the joint radiative forcing from other greenhouse gases and aerosols. The annual time step not only gives a better resolution of optimal policy, but most importantly enables us to adopt the annual estimates for consumption smoothing and risk preferences from the macro and finance literature. The decadal time step of the original DICE model would imply fewer, but much larger jumps in both risk and time. The smoothing and aversion parameters cannot simply be adjusted to such large time steps because the model is not in a steady state.

Given the curse of dimensionality in dynamic programming, we reduce the state variables of the DICE model through analytic approximations, increasing numeric precision and reducing computation time. These simplifications in the representation of DICE's climate module are discussed in detail in the accompanying paper Traeger³⁹. The simplifications replace the explicit model of the ocean's temperature and carbon states by time-dependent cooling and decay functions. Traeger³⁹ shows that the resulting climate model performs just as well as the original DICE model in replicating the temperature response to the Intergovernmental Panel on Climate Change's SRES and RCP scenarios. His assessment compares our model's and the DICE model's climate module to the emulation of MAGICC 6.0 (ref. 40) of the temperature response of the Atmosphere–Ocean General Circulation Model used in the Intergovernmental Panel on Climate Change Assessment Reports.

Our (Hamilton–Jacobi–)Bellman equation disentangles intertemporal consumption smoothing from risk aversion, following Epstein and Zin²³ and Weil²⁸, using the numerically more convenient time additive formulation from Traeger⁴¹. We solve the non-autonomous Bellman equation by function iteration (using time as an additional state), subject to the equations of motion of our DICE version, summarized in the Supplementary Methods and discussed in

detail in Traeger³⁹. We employ Chebyshev polynomials to approximate the value function and use Howard's method to speed up the algorithm. We increase node numbers and reduce tolerance until we find no more changes in the optimal policies.

Nordhaus's² Monte Carlo study motivates our 'low'-uncertainty scenario. The author endows the damage coefficient b_1 with a normal distribution and discards negative draws. Instead, we use a lognormal distribution to ensure positivity of the damage range. Supplementary Fig. 2 shows that using a normal distribution instead, even when permitting negative draws of the damage coefficient, results in virtually the same optimal policy as our lognormal distribution. We calibrate our lognormal distribution on the damage coefficient and a normal distribution on the damage exponent to Nordhaus² by requiring similar damages for a one sigma deviation at a 3 °C warming level. We find a standard deviation of 0.13% for the distribution governing the damage coefficient, and a standard deviation of 0.35 for the damage exponent. We base the 'high'-uncertainty scenario on Tol's¹⁵ (rounded) estimate of damage exponent uncertainty with a standard deviation of 0.5, which he derives from a comparison across different integrated assessment models. We choose a corresponding 'high'-uncertainty standard deviation for coefficient uncertainty by requiring similar damages for plus/minus a one sigma standard deviation at a 3 °C warming level, resulting in a standard deviation of 0.25%. Numerically, we implement the continuous distributions using Gauss–Legendre quadrature nodes matching the first nine moments of the distribution. We test that further increases in the number of Gauss–Legendre nodes do not affect the results.

We adopt DICE's original preference parameter of $\eta = 2$ in the standard economic evaluation (N). It simultaneously captures the degree of risk aversion and the preference for intertemporal consumption smoothing. The comprehensive evaluation framework disentangles these two preference characteristics. The disentangled consumption smoothing preference estimate converges over the recent years to the elasticity $\eta = 2/3$ (refs 16–21). The same studies estimate the disentangled Arrow–Pratt measure of relative risk aversion RRA in the range from 8 to 10. Our disentangled set of simulations (D) picks $\eta = 2/3$ and RRA = 9.5, based on a best guess by Vissing-Joergensen and Attanasio¹⁶.

Received 5 April 2013; accepted 24 April 2014;
published online 15 June 2014

References

- Nordhaus, W. D. *Managing the Global Commons: The Economics of the Greenhouse Effect* (MIT Press, 1994).
- Nordhaus, W. D. *A Question of Balance: Economic Modeling of Global Warming* (Yale Univ. Press, 2008).
- Richels, R. G., Manne, A. S. & Wigley, T. M. Moving beyond concentrations: The challenge of limiting temperature change *AEI-Brookings Joint Center for Regulatory Studies* 04-11 (2004).
- Hope, C. The marginal impact of CO₂ from PAGE2002: An integrated assessment model incorporating the IPCC's five reasons for concern. *Integr. Assess. J.* **6**, 19–56 (2006).
- Dietz, S. High impact, low probability? An empirical analysis of risk in the economics of climate change. *Climatic Change* **108**, 519–541 (2009).
- Anthoff, D., Tol, R. S. J. & Yohe, G. W. Risk aversion, time preference, and the social cost of carbon. *Environ. Res. Lett.* **4**, 1–7 (2009).
- Anthoff, D. & Tol, R. S. J. The impact of climate change on the balanced growth equivalent: An application of FUND. *Environ. Res. Econom.* **43**, 351–367 (2009).
- Interagency Working Group on Social Cost of Carbon, U. S. G. *Technical Support Document: Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866* (Department of Energy, 2010).
- Pycroft, J., Vergano, L., Hope, C. W., Paci, D. & Ciscar, J. C. A tale of tails: Uncertainty and the social cost of carbon dioxide. *Econom. E-J.* **5**, 1–29 (2011).
- 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. *Econom. E-J.* **6**, 1–40 (2012).
- Kelly, D. L. & Kolstad, C. D. Bayesian learning, growth, and pollution. *J. Econom. Dynam. Control* **23**, 491–518 (1999).
- Keller, K., Bolker, B. M. & Bradford, D. F. Uncertain climate thresholds and optimal economic growth. *J. Environ. Econom. Manage.* **48**, 723–741 (2004).
- Leach, A. J. The climate change learning curve. *J. Econom. Dynam. Control* **31**, 1728–1752 (2007).
- Hanemann, M. in *Climate Change Science and Policy* (eds Schneider, S. H., Rosencranz, A., Mastrandrea, M. & Kuntz-Duriseti, K.) Ch. 17, 185–193 (Island Press, 2009).
- Tol, R. S. The economic effects of climate change. *J. Econom. Perspect.* **23**, 29–51 (2009).
- Vissing-Joergensen, A. & Attanasio, O. P. Stock-market participation, intertemporal substitution, and risk-aversion. *Am. Econom. Rev.* **93**, 383–391 (2003).

17. Bansal, R. & Yaron, A. Risks for the long run: A potential resolution of asset pricing puzzles. *J. Finan.* **59**, 1481–509 (2004).
18. Bansal, R., Kiku, D. & Yaron, A. Long run risks, the macroeconomy, and asset prices. *Am. Econom. Rev.: Papers Proc.* **100**, 542–546 (2010).
19. Chen, X., Favilukis, J. & Ludvigson, S. C. An estimation of economic models with recursive preferences. *Quant. Econom.* **4**, 39–83 (2013).
20. Bansal, R., Kiku, D. & Yaron, A. An empirical evaluation of the long-run risks model for asset prices. *Crit. Finance Rev.* **1**, 183–221 (2012).
21. Nakamura, E., Steinsson, J., Barro, R. & Ursua, J. Crises and recoveries in an empirical model of consumption disasters. *Am. Econom. J.: Macroeconom.* **5**, 35–74 (2013).
22. Kreps, D. M. & Porteus, E. L. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica* **46**, 185–200 (1978).
23. 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).
24. Traeger, C. P. Recent developments in the intertemporal modeling of uncertainty. *ARRE* **1**, 261–285 (2009).
25. Gollier, C. Discounting an uncertain future. *J. Public Econom.* **85**, 149–166 (2002).
26. Ha-Duong, M. & Treich, N. Risk aversion, intergenerational equity and climate change. *Environ. Res. Econom.* **28**, 195–207 (2004).
27. Traeger, C. P. Why uncertainty matters—Discounting under intertemporal risk aversion and ambiguity. *Econom. Theory* (in the press, 2014).
28. Weil, P. Nonexpected utility in macroeconomics. *Quart. J. Econom.* **105**, 29–42 (1990).
29. Epstein, L. G. & Zin, S. E. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis. *J. Polit. Economy* **99**, 263–286 (1991).
30. Crost, B. & Traeger, C. P. Optimal climate policy: Uncertainty versus Monte-Carlo. *Econom. Lett.* **120**, 552–558 (2013).
31. Nordhaus, W. D. A review of the Stern review on the economics of climate change. *J. Econom. Literat.* **45**, 686–702 (2007).
32. Stern, N. (ed.) *The Economics of Climate Change: The Stern Review* (Cambridge Univ. Press, 2007).
33. Jensen, S. & Traeger, C. Growth uncertainty in the integrated assessment of climate change. *Eur. Econom. Rev.* (in the press, 2014).
34. Nordhaus, W. D. & Yang, Z. A regional dynamic general-equilibrium model of alternative climate-change strategies. *Am. Econom. Rev.* **86**, 741–765 (1996).
35. Bosetti, V., Carraro, C., Galeotti, M., Massetti, E. & Tavoni, M. Witch: A world induced technical change hybrid model. *Energy J.* **27**, 13–38 (2006).
36. Bauer, N., Baumstark, L. & Leimbach, M. The ReMIND-R model: The role of renewables in the low-carbon transformation—First best vs second-best worlds. *Climatic Change* **114**, 145–168 (2012).
37. Broome, J. *Counting the Cost of Global Warming* (White Horse Press, 1992).
38. Asheim, G. Intergenerational equity. *Annu. Rev. Econom.* **2**, 197–222 (2010).
39. Traeger, C. P. A 4-stated DICE: Quantitatively addressing uncertainty effects in climate change. *Environ. Res. Econom.* (in the press, 2014).
40. Meinshausen, M., Raper, S. & Wigley, T. Emulating coupled atmosphere–ocean and carbon cycle models with a simpler model, MAGICC6—part 1: Model description and calibration. *Atmos. Chem. Phys.* **11**, 1417–1456 (2011).
41. Traeger, C. P. Intertemporal risk aversion. *CUDARE Working Paper No. 1102* (2010).
42. Croce, M. M., Kung, H., Nguyen, T. T. & Schmid, L. Fiscal policies and asset prices. *Rev. Financ. Stud.* **25**, 2635–2672 (2012).
43. McGrattan, E. R. & Prescott, E. C. Taxes, regulations, and the value of US and UK corporations. *Rev. Econom. Stud.* **72**, 767–796 (2005).

Acknowledgements

We are grateful to L. Karp, S. Jensen, A. Butz, M. Hanemann, K. Smith, D. Lemoine, D. Farber and J. Harte, and for financial support by the Giannini Foundation.

Author contributions

C.P.T. developed the research question and model and wrote the paper. Both authors together designed the algorithms and the graphical output; C.P.T. assisted and B.C. took the lead in programming and running 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. Correspondence and requests for materials should be addressed to C.P.T.

Competing financial interests

The authors declare no competing financial interests.