



Declining CO₂ price paths

Kent D. Daniel^{a,b,1}, Robert B. Litterman^{c,1}, and Gernot Wagner^{d,1,2,3}

^aColumbia Business School, New York, NY 10027; ^bNational Bureau of Economic Research, Cambridge, MA 02138; ^cKeos Capital, New York, NY 10018; and ^dHarvard University Center for the Environment, Cambridge, MA 02138

Edited by Jose A. Scheinkman, Columbia University, New York, NY, and approved September 9, 2019 (received for review October 9, 2018)

Pricing greenhouse-gas (GHG) emissions involves making trade-offs between consumption today and unknown damages in the (distant) future. While decision making under risk and uncertainty is the forte of financial economics, important insights from pricing financial assets do not typically inform standard climate–economy models. Here, we introduce EZ-Climate, a simple recursive dynamic asset pricing model that allows for a calibration of the carbon dioxide (CO₂) price path based on probabilistic assumptions around climate damages. Atmospheric CO₂ is the “asset” with a negative expected return. The economic model focuses on society’s willingness to substitute consumption across time and across uncertain states of nature, enabled by an Epstein–Zin (EZ) specification that delinks preferences over risk from intertemporal substitution. In contrast to most modeled CO₂ price paths, EZ-Climate suggests a high price today that is expected to decline over time as the “insurance” value of mitigation declines and technological change makes emissions cuts cheaper. Second, higher risk aversion increases both the CO₂ price and the risk premium relative to expected damages. Lastly, our model suggests large costs associated with delays in pricing CO₂ emissions. In our base case, delaying implementation by 1 y leads to annual consumption losses of over 2%, a cost that roughly increases with the square of time per additional year of delay. The model also makes clear how sensitive results are to key inputs.

climate risk | asset pricing | cost of carbon

For over 25 y, the dynamic integrated climate–economy (DICE) model (1–3) has been the standard tool for analyzing CO₂ emissions-reductions pathways, and for good reason. One attraction is its simplicity, turning a “market failure on the greatest scale the world has seen” (4) and “the mother of all externalities” (5) into a model involving fewer than 20 main equations, 3 representing the climate system (6). DICE has spawned many variants (7). It has also helped set the tone for what many consider “optimal” CO₂ price paths. The core trade-off between economic consumption and climate damages leads to relatively low CO₂ prices today rising over time.

DICE and models like it have well-known limitations, including how they represent climate risk and uncertainty (7–15). DICE, for example, is not an optimal-control model, as commonly understood by economists employing modern dynamic economic analysis, even though it lends itself to those extensions (9–12). The underlying structure all but prescribes a rising CO₂ price path over time.

One important limitation is the form of the utility function. Constant relative risk aversion (CRRA) preferences, standard in most climate–economy models (1, 7, 16), assume that economic agents have an equal aversion to variation in consumption across states of nature and over time. Evidence from financial markets suggests that this is not the case (17). The risk premium (RP) of equities over bonds points to a fundamental difference in how much society is willing to pay to substitute consumption risk across states of nature compared to over time (18, 19). Some have explained the discrepancy by allowing for extreme events (20–22), and others have looked to more flexible preferences (23–26) or both (27). Our own preference specification follows Epstein and Zin (EZ) (24, 25).

EZ Preferences

Here, we use EZ preferences and focus on climate uncertainties. We approach climate change as an asset pricing problem with atmospheric CO₂ as the “asset.” The value of an investment in reducing CO₂ emissions depends on the state of nature, represented by its fragility θ_t . That, in turn, helps determine the discount rate applied to the damages that would have occurred without the investment.

Our representative agent maximizes a recursive utility U_t based on consumption c_t and expectations E_t over future utility for times $t \in \{0, 1, 2, \dots, T - 1\}$:

$$U_t = \left[(1 - \beta)c_t^\rho + \beta(E_t [U_{t+1}^\alpha]^\rho)^\frac{1}{\rho} \right]^\frac{1}{\alpha} \quad [1]$$

Parameters α and ρ measure the agent’s willingness to substitute consumption across states of nature and across time, respectively. (See *Methods* for the final-period utility U_T and further derivations.) CRRA preferences are a special case, with $\alpha = \rho$. Unlike with CRRA, Eq. 1 implies that CO₂ prices no longer collapse to zero with increasing risk aversion (RA) and equity risk premia (Fig. 1A). The same goes for the portion of CO₂ prices explained by RA (Fig. 1B).

EZ preferences have since found their way into the climate–economic literature (9–12, 28–35). Some have embedded EZ into DICE (28, 35), and others employ supercomputers to solve (9–12). The complexity typically does not allow for analytic solutions (34). We here follow a simple binomial-tree model with a long history in financial modeling application (36). It is precisely this modeling choice—standard in financial economics but novel to climate–economic applications—that leads to our fundamentally differing CO₂ price paths. Mitigating climate risk provides

Significance

Risk and uncertainty are important in pricing climate damages. Despite a burgeoning literature, attempts to marry insights from asset pricing with climate economics have largely failed to supplement—let alone supplant—decades-old climate–economy models, largely due to their analytic and computational complexity. Here, we introduce a simple, modular framework that identifies core trade-offs, highlights the sensitivity of results to key inputs, and helps pinpoint areas for further work.

Author contributions: K.D.D., R.B.L., and G.W. designed research, performed research, contributed new analytic tools, and wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹K.D.D., R.B.L., and G.W. contributed equally to this work.

²Present addresses: Department of Environmental Studies, New York University, New York, NY 10003; and Robert F. Wagner Graduate School of Public Service, New York University, New York, NY 10012.

³To whom correspondence may be addressed. Email: gernot@wagner.com

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1817444116/-DCSupplemental.

First published October 1, 2019.

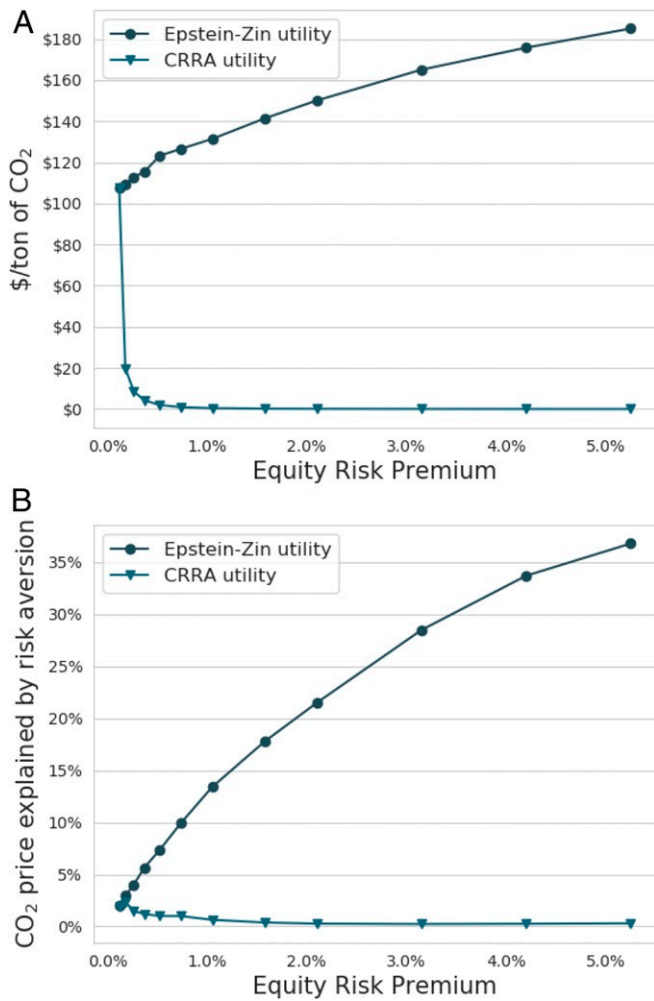


Fig. 1. Risk calibration. *A* shows how using EZ preferences, unlike CRRA, results in increasing 2015 CO₂ prices, in 2015 US\$, with increasing RA, translated into the implied equity RP using Weil’s conversion (19), while holding implied market interest rates stable at 3.11%. *B* shows how the percentage of the 2015 CO₂ price explained by RA, as opposed to expected damages (EDs), increases with equity RP for EZ utility, while decreasing for CRRA (Risk Decomposition).

a hedge, leading to high CO₂ prices early on. As uncertainties decline over time, so do CO₂ prices.

Model

The setting for EZ-Climate is a standard endowment economy (37). In each period, the agent is endowed with a certain amount of the consumption good \bar{c}_t . However, she is not able to consume the full \bar{c}_t for 2 reasons: climate change and climate policy. In periods $t \in \{1, 2, \dots, T\}$, a portion of \bar{c}_t may be lost due to climate-change damages, which are, in turn, a function of cumulative radiative forcing (CRF_t) up to time t and of fragility θ_t : $D_t(CRF_t, \theta_t)$. Up to period $T - 1$, the agent may elect to spend some of \bar{c}_t to reduce her impact on the future climate, κ_t , which, in turn, depends on mitigation x_t . The resulting consumption c_t , after D_t and κ_t are taken into account, is given by:

$$c_0 = \bar{c}_0 (1 - \kappa_0(x_0)), \tag{2}$$

$$c_t = \bar{c}_t (1 - \kappa_t(x_t))(1 - D_t(CRF_t, \theta_t)), \tag{3}$$

for $t \in \{1, 2, \dots, T - 1\}$, and

$$c_T = \bar{c}_T (1 - D_T(CRF_T, \theta_T)). \tag{4}$$

CRF_t is a function of mitigation, x_s , in each period from 0 to t , calibrated to a combination of Representative Concentration Pathway (RCP) scenarios (Climate Damages).

The agent maximizes utility given by Eq. 1 in each of T periods by selecting, at each time and in each state, a level of mitigation x_t , creating in essence a “ $2^{T+1} - 1$ ”-dimensional optimization problem. Fig. 2*A* shows our base case, which uses a 7-period tree with decision nodes from 2015 through 2300. At each node, more information about θ_t and the resulting climate damages is revealed, before uncertainty is resolved at the beginning of the penultimate period in 2300. In the “easy” spirit of EZ-Climate, the limited number of decision points makes the solution both tractable and quickly solvable.

EZ-Climate provides an accessible, modular framework (7) that is dependent on key economic inputs—chiefly, RA and the elasticity of intertemporal substitution (EIS)—and 2 main climate-related ones: mitigation costs and climate damages. Costs depend primarily on assumptions around backstop

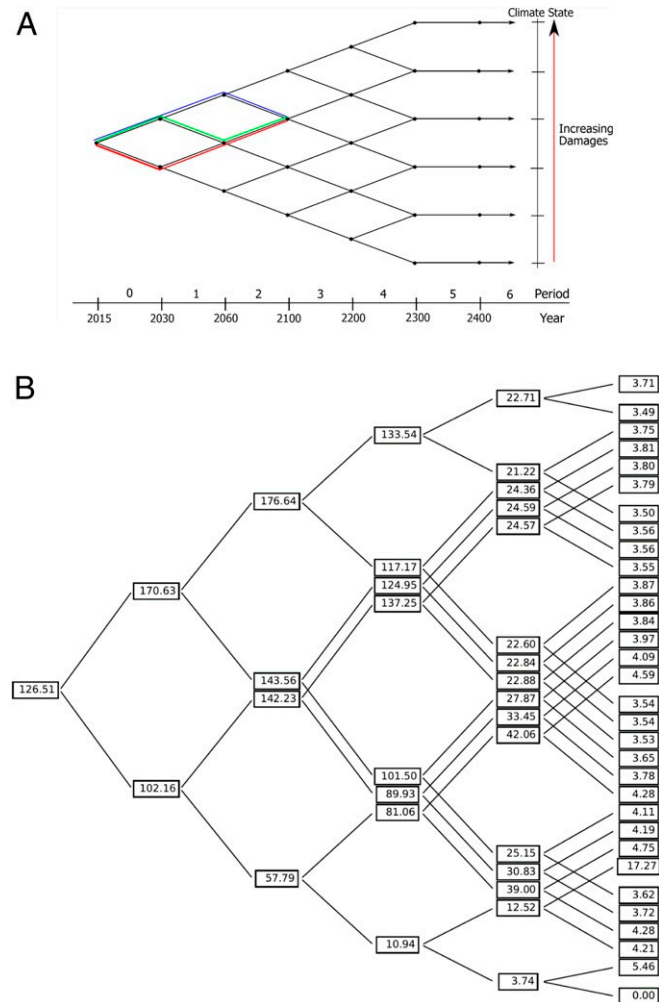


Fig. 2. Model tree structure. *A* shows a diagram of the binomial tree structure (with probability $p = \frac{1}{2}$) used in solving the model for each state of nature θ_t across time $t \in \{1, 2, \dots, T\}$, corresponding to years 2015, 2030, 2060, 2100, 2200, 2300, and final period 2400. Note the “recombining” tree structure, highlighted in the first 2 periods: Damage functions in any particular state are independent of the path taken, but x_t and the resulting U_t are path-dependent. *B* shows the CO₂ price in the base case ($peakT = 6$, $disaster\ tail = 18$, and $EIS = 0.9$).

technologies (38) and technological change. Damages depend on the full climate–economic chain from economic output to CO₂ emissions, from emissions to concentrations, from concentrations to *CRF*, and from *CRF* to climate damages lowering economic consumption. While all of these calibrations are important, and uncertainties abound, a key addition is allowing for potentially catastrophic risk in form of climatic tipping points (TPs) (39, 40) (Climate Damages).

Results and Discussion

Fig. 2*B* shows CO₂ prices for each node of the tree in our base case. The 2015 CO₂ price comes from a single node in the tree. In each subsequent period, the price is set in expectation over all possible states of nature θ_t in that given period. All grouped nodes at a given time have the same θ_t and, thus, the same damage for a given amount of *CRF*. The price itself is path-dependent. Fig. 2*B* also shows the costs associated with bad θ_t draws in latter periods. Bad news is costly. Bad news late, when it is more difficult to counteract with more active policy, is worse. It is precisely the inability to know upfront when good or bad news arrives that accounts for the insurance value of early mitigation and, thus, the role that the resolution of risk over time plays in the declining CO₂ price.

Declining CO₂ Price. Unlike most modeled CO₂ price paths, ours typically rise briefly before declining over time. One partial explanation is the move from CRRA to EZ preferences. CRRA preferences duplicate the decline only with a $RA = \frac{10}{9} \sim 1.1$, when EZ collapses to CRRA. For higher RA, consistent with those estimated from models calibrated to financial-market data, early CRRA prices collapse to near zero (Fig. 3*A*). But

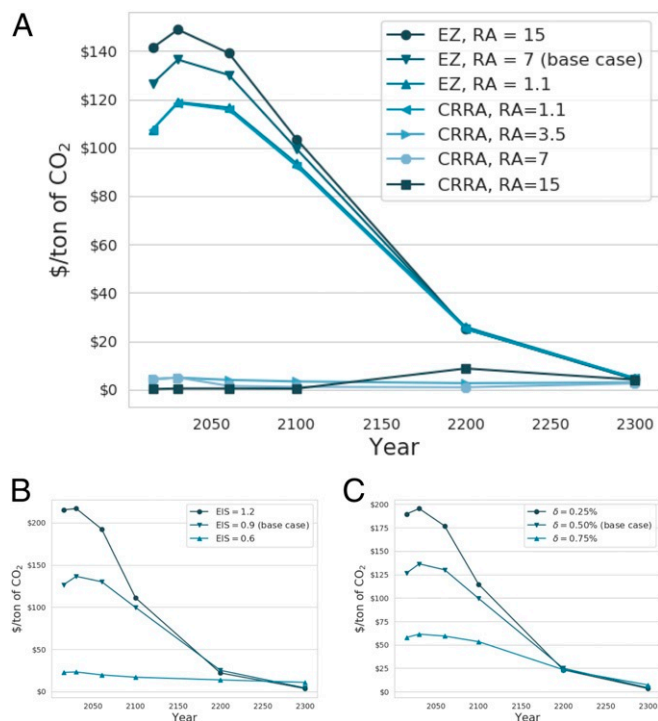


Fig. 3. Declining CO₂ price paths. *A* shows how EZ utility here leads to CO₂ prices that start high and decline over time, regardless of assumed RA, a feature mimicked only by unrealistically low $RA = \frac{10}{9} \sim 1.1$, when EZ and CRRA utilities coincide (Economic Parameters). *B* and *C* show the importance of EIS and the rate of pure time preference (δ), respectively. *B* varies real interest rates from 2.74% (EIS = 1.2) to 3.77% (EIS = 0.6) to keep $\bar{c} = 1.5\%$. *C* fixes EIS at 0.9, while δ varies from 0.25% to 0.75% (SI Appendix, Figs. S5 and S6).

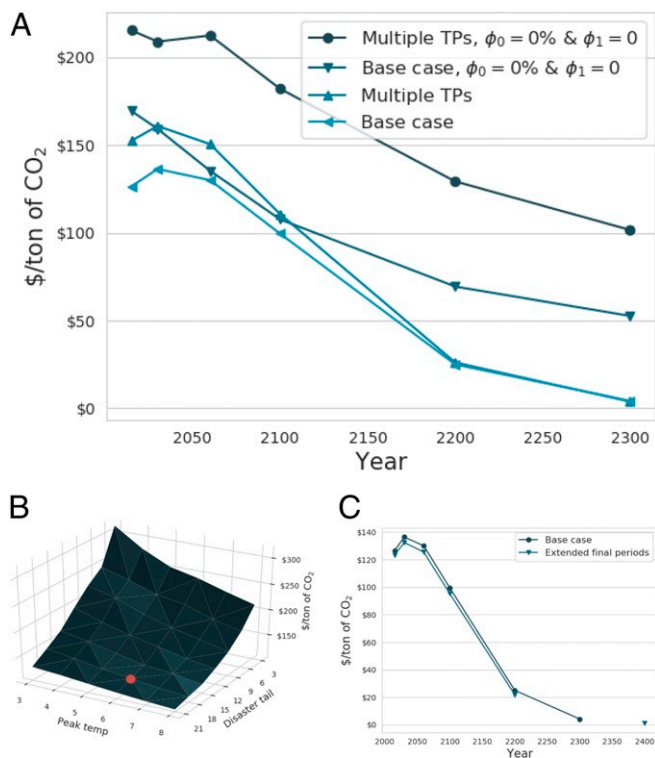


Fig. 4. CO₂ price sensitivities. *A* shows the implications of technological change and TP assumptions. Setting $\phi_0 = 0\%$ $\phi_1 = 0$ increases early-year and final-period prices, flattening the price path. Multiple TPs act akin to fattening the tail of the damage function, steepening the price path. They also interact with the no-technological-change assumption, increasing final-period prices. *B* shows that 2015 CO₂ prices depend crucially on “catastrophic” climate risk assumptions, set to $peakT = 6^\circ C$ and $disaster\ tail = 18$ in the base case (Climate Damages). *C*, by contrast, shows the minimal implications of extending the final period from 2300 to 2400 for $t = 6$ and 7, respectively, in the base case to 2400 and 2700.

going from CRRA to EZ preferences is not the only explanation, implied by the fact that CRRA price paths stay flat over time.

Others have pointed to reasons for declining CO₂ price paths including producer behavior (41), the need for directed technological change from “dirty” to “clean” sectors (42), or inertia (43). We here find 2 factors driving the declining CO₂ price paths: the resolution of uncertainty, combined with technological progress that makes mitigation significantly cheaper over time.

Our base case assumes exogenous technological progress $\phi_0 = 1.5\%$ and endogenous progress $\phi_1 = 0.015$ per year, linked to average mitigation efforts to date (Eq. 19). The combination makes mitigation costs diminishingly small hundreds of years out, helping to drive the declining price paths (Fig. 4*A* and SI Appendix, Fig. S3).

Another reason for declining price paths is the assumed nature of TPs in the base case. Each node has a certain probability of hitting a TP, given by Eq. 22. Once hit, there is no reversing the resulting damages. That structure increases prices in early years, decreasing them later, as it introduces a nonconvexity into the damage function (37). Allowing for multiple TPs exacerbates that result in the base case, as it fattens the tail of the damage function (Fig. 4*A*). Assuming no technological change, meanwhile, increases final-period prices, more so with multiple TPs.

While the declining CO₂ price path is a persistent feature across model specifications (SI Appendix, Figs. S3, S5, and S6),

Table 1. Social cost of delay by first-period length

First-period length, y	Annual consumption impact during first period, %
5	11
10	23
15	36

Absolute annual household consumption numbers are around \$40 trillion per year.

the absolute CO₂ price in early years depends crucially on a number of calibration choices. Fig. 3 shows the importance of economic parameters, chiefly, EIS and the pure rate of time preference (δ). Fig. 4B shows the sensitivity of the initial CO₂ price to assumptions around “catastrophic” climate risk. Our base case assumes 6 °C for the “peak temperature” (*peakT*) and 18 for the *disaster tail* calibrations. While there is seeming convergence around 6 °C as an upper bound for what could conceivably be quantified (see, for example, <https://helixclimate.eu/>), declaring it equivalent to a “global TP” is at best unduly conservative (11, 15, 40), at worst arbitrary. Much more work is needed to justify any one particular parameter value and, thus, any one CO₂ price (7). Our goal with EZ-Climate is to provide a simple, modular framework to think about climate risks, uncertainties, TPs, and their implications for CO₂ prices.

Social Cost of Delay. The optimal-control nature of EZ-Climate also allows for a calculation of the social cost of delay in implementing CO₂ prices. Unlike prior efforts (2, 7), we do not look to the CO₂ price for estimating that cost. In fact, doing so can be positively misleading. After constraining the price to \$0 in the first period, the price in the second period is lower than in the unconstrained case. The price reflects the marginal benefits of additional emissions reductions, which are now lower. We here instead quantify the cost of delay by constraining mitigation to zero in the first period and asking how much additional consumption would be required during that period in order to bring the utility of the representative agent to the level of the unconstrained solution.

Table 1 shows the annual consumption loss during the constrained first period. For a 10-y delay, the equivalent annual consumption loss over the first constrained period is ~23%. Each year of delay increases the annual consumption loss over the entire constrained period by ~2.3%. It also increases the time interval of the loss, thus leading to a slightly more than quadratic rate of increase in the deadweight loss of utility over time. In rough monetary terms, delaying implementation by only 1 y costs society approximately \$1 trillion. A 5-y delay creates the equivalent loss of approximately \$24 trillion, comparable to a severe global depression. A 10-y delay causes an equivalent loss in the order of \$10 trillion per year, approximately \$100 trillion in total.

Conclusion

Our conclusion could mimic that of DICE, introduced over 25 y ago (1), with one crucial difference: Like with DICE, and despite crucial recent advances (7, 35), “it should be emphasized that this analysis has a number of important qualifications,” especially, ironically, “the economic impact of climate change” (1). Unlike DICE, EZ-Climate does not “[abstract] from issues of uncertainty” (1). It embraces them, following a simple binomial-tree framework long used in the finance literature (36). The simple, modular framework also highlights the sensitivity of CO₂ prices to key inputs. There is no single, correct, “optimal” price path. One persistent feature, however, is declining price paths. That puts the focus on near-term action and on the large costs of delay.

Methods

Utility Specification. Eq. 1 represents a special case of Kreps–Porteus preferences (23), following EZ (24, 25) for $t \in \{0, 1, 2, \dots, T - 1\}$, with c_t given by Eqs. 2–4. In $t = T$ (=2400, in our base-case calibration), the representative agent receives utility from all present discounted consumption from T onward. Consumption grows at a constant rate r for $t \geq T$:

$$c_t = c_T(1 + r)^{t-T}, \tag{5}$$

with c_T given by Eq. 4. The resulting final-period utility is:

$$U_T = \left[\frac{1-\beta}{1-\beta(1+r)^P} \right]^{\frac{1}{\beta}} c_T. \tag{6}$$

Risk Decomposition. Fig. 1B shows the split between EDs and the RP in explaining 2015 CO₂ prices (44). To calculate the cost of an additional ton of CO₂ emissions, we sum over all consumption damages, in every state of nature s at every future time t , multiplied by the value of an additional unit of consumption for each s and t . The 2015 CO₂ price, thus, is:

$$\sum_{t=1}^T \sum_{s=1}^{S(t)} \pi_{s,t} m_{s,t} D_{s,t} = \sum_{t=1}^T E_0 [\tilde{m}_t \tilde{D}_t], \tag{7}$$

where $S(t)$ denotes the number of states at time t , $\pi_{s,t}$ the probability of state s at time t , and pricing kernel $m_{s,t} = \left(\frac{\partial U}{\partial c_{s,t}} \right) / \left(\frac{\partial U}{\partial c_0} \right)$. Eq. 7 can be further decomposed into ED and RP:

$$\underbrace{\sum_{t=1}^T E_0 [\tilde{m}_t] E_0 [\tilde{D}_t]}_{ED} + \underbrace{\sum_{t=1}^T cov_0 [\tilde{m}_t, \tilde{D}_t]}_{RP}. \tag{8}$$

Note that $E_0 [\tilde{m}_t] = 1/R^f(0, t)$, where $R^f(0, t)$ is the payoff, at time t , to a \$1 investment in a risk-free bond at t_0 . Alternatively, $E_0 [\tilde{m}_t]$ is the risk-free discount factor between today and t . ED, thus, is the sum of marginal climate damages, discounted back to the present at the risk-free rate:

$$ED = \sum_{t=1}^T E_0 [\tilde{D}_t] / R^f(0, t). \tag{9}$$

RP then is the difference between the CO₂ price and ED.

Mitigation Costs. Calibrating the mitigation cost requires specifying a relationship between the marginal cost of emissions reductions, equal to the per-ton tax rate τ , the resulting flow of emissions per year $g(\tau)$, and the fraction of emissions reduced $x(\tau)$.

Many modeling efforts have attempted to estimate the marginal abatement costs (MACs), often as part of integrated assessment models. See, for example, Stanford’s Energy Modeling Forum (<https://emf.stanford.edu/>). Perhaps the most influential, independent effort comes from McKinsey & Company in an attempt to estimate a bottom-up MAC curve (MACC) (45). McKinsey’s MACCs are, to a large extent, based on bottom-up “engineering” estimates. That makes them an easy target for critique by economists, who often focus on the large abatement opportunities with “negative” costs or the “energy-efficiency gap” (46, 47). We calibrate τ , $g(\tau)$, and $x(\tau)$ based on McKinsey’s global MACC effort (48), with one crucial modification: We set $x(\tau) = 0$ for $\tau \leq 0$; i.e., we assume no net-negative cost mitigation. *SI Appendix, Table S1* shows the resulting modified point estimates, which we fit to a power function for $x(\tau)$, yielding:

$$x(\tau) = 0.0923\tau^{0.414}. \tag{10}$$

The corresponding inverse function, solving for τ to achieve x , yields the marginal cost of abatement:

$$\tau(x) = 314.32x^{2.413}. \tag{11}$$

Ultimately, we are interested in the total cost to society $\kappa(\tau)$ for each particular tax τ . We calculate $\kappa(\tau)$ using the envelope theorem, assuming the representative agent chooses $g(\tau)$ so as to maximize consumption c given τ : $\frac{dc(\tau)}{d\tau} = -g(\tau)$. Consumption x associated with a particular τ , thus, equals:

$$c(\tau) = \bar{c} - \int_0^\tau g(s) ds. \tag{12}$$

However, Eq. 12 is only correct if the government were to collect τ and then waste 100% of the proceeds $g(\tau)\tau$. Here, we assume instead that the proceeds are refunded in lump sum (49). Refunding $g(\tau)\tau$ and rewriting Eq. 12 yields total mitigation costs of:

$$K(\tau) = \int_0^\tau g(s) ds - g(\tau)\tau. \quad [13]$$

The lump-sum refund does not allow for CO₂ tax proceeds to be used to decrease other distortionary taxes, which would make the total costs smaller still (50, 51). Rewriting $g(\tau) = g_0(1 - x(\tau))$, where g_0 is the emissions baseline, we can rewrite Eq. 13 as:

$$K(\tau) = g_0 [\tau x(\tau) - \int_0^\tau x(s) ds]. \quad [14]$$

Substituting Eqs. 10 and 11 into Eq. 14 and dividing by current aggregate consumption yields the societal cost of a given level of mitigation as a percentage of initial consumption c_0 :

$$\kappa(x) = (g_0 92.08 / C_0) x^{3.413}. \quad [15]$$

Our base-case calibration assumes $g_0 = 52$ billion tons of CO₂-equivalent emissions and $c_0 = \$31$ trillion/year representing current 2015 global consumption in 2015 US\$. Eq. 15 assumes no technological progress and no backstop technology.

Backstop Technology. We also allow for a backstop technology in form of CO₂ removal (38, 52) to become available at cost τ^* at x_0 and to be used exclusively for MACs $\geq \bar{\tau}$. The MACC including a backstop follows:

$$B(x) = \bar{\tau} - \left(\frac{k}{x}\right)^{\frac{1}{b}}. \quad [16]$$

We calibrate Eq. 16 to set $B(x_0) = \tau^*$, and we impose a smooth-pasting condition at x_0 , resulting in:

$$k = x_0 (\bar{\tau} - \tau^*) \frac{\bar{\tau} - \tau^*}{(\alpha - 1) \tau^{*\alpha}}. \quad [17]$$

Our base case assumes $\tau^* = \$2,000$ and $\bar{\tau} = \$2,500$ in 2015 dollars. Under the most aggressive backstop scenario presented in *Results and Discussion*, we assume $\tau^* = \$300$ and $\bar{\tau} = \$350$. *SI Appendix, Fig. S1* shows the resulting 2015 $\tau(x)$.

Technological Progress. *SI Appendix, Fig. S1* is calibrated to $t = 0$. In subsequent periods, we allow the MACC to decrease at a rate determined by a set of technological change parameters: a constant component φ_0 and an endogenous component linked to mitigation efforts to date $\varphi_1 X_t$, where X_t is the average mitigation up to time t defined by:

$$X_t = \sum_{s=0}^t g_s x_s / \sum_{s=0}^t g_s. \quad [18]$$

Mitigation costs at time t are:

$$\kappa_t(x) = \kappa(x) [1 - \varphi_0 - \varphi_1 X_t]^t. \quad [19]$$

This functional form allows for easy calibration. For example, if $\varphi_0 = 0.005$ and $\varphi_1 = 0.01$, and with average mitigation $X_t = 50\%$, $\kappa_t(x)$ decreases at a rate of 1% per year.

SI Appendix, Fig. S3 shows the implications of both backstop technology and technological progress on the CO₂ price.

Climate Damages. We derive $D_t(CRF_t, \theta_t)$ in 2 steps. Damages are a function of temperature changes ΔT , which, in turn, are a function of CRF. We calibrate ΔT over time based on the International Energy Agency's (IEA's) projections for ΔT_{100} using its "new policies scenario" (53), equating the values with $t = 100$. We fit a displaced gamma distribution (54) around the IEA's 2100 greenhouse gas (GHG) concentration projections for 450, 650, and 1,000 ppm. *SI Appendix, Table S2* shows the calibration results. We translate ΔT_{100} into ΔT_t for $t = 0$ (year 2015) through $t = 385$ (year 2400) using:

$$\Delta T_t = 2 \Delta T_{100} \left[1 - 0.5 \frac{t}{100} \right]. \quad [20]$$

SI Appendix, Fig. S2 shows the results for ΔT_t based on ΔT_{100} values. We then fit a log-normal distribution for equilibrium climate sensitivity (15, 55) around ΔT_{100} to generate distributions for ΔT_t based on emissions-reduction pathways x_t . One important possible extension is around timing, further probing our assumption of equating effects at ΔT_{100} with climate sensitivity (56).

We calibrate damages D_t based on ΔT considering 2 multiplicative components: a noncatastrophic loss function and a catastrophic hazard function (32). The noncatastrophic component is in the form of a displaced gamma distribution (54), resulting in the loss function:

$$L(\Delta T_t) = e^{-13.97 \gamma \Delta T_t^2}. \quad [21]$$

Parameter γ is drawn from a displaced gamma distribution with its 3 parameters $a = 4.5$, $d = 21,341$, and $p = -0.0000746$. That calibration, much like other econometrically based estimates extrapolating from past experience (57), all but rules out "catastrophic" damages and so-called climatic TPs or "tipping elements" (39, 40).

We augment this calibration with a "catastrophic" component, assuming a particular probability of hitting a climatic TP, $\text{Prob}(TP)$, in any given period, if temperature changes cross a $peakT$ threshold:

$$\text{Prob}(TP) = 1 - \left(1 - \left[\frac{\Delta T(t)}{\max[\Delta T(t), peakT]} \right]^2 \right)^{\frac{period}{30}}. \quad [22]$$

In the base-case calibration, we set $peakT = 6^\circ\text{C}$ (Fig. 3B). While ad hoc, the number is, if anything, unduly conservative (15, 39, 40, 55). *SI Appendix, Fig. S4A* shows the resulting probabilities for a 30-y period. Conditional on hitting a TP at time t^* , the level of consumption for each subsequent $t \geq t^*$ is reduced by the factor $e^{-TP_{damage}}$, where TP_{damage} is drawn from a gamma distribution with parameters $\alpha = 1$ and $\beta = \text{disaster tail}$. *SI Appendix, Fig. S4B* shows the resulting probability of economic damages exceeding a particular percentage of total output.

We then generate distributions for $D_t(CRF_t, \theta_t)$ for each period for each of 3 maximum GHG concentration levels—450, 650, and 1,000 ppm—based on 6 million draws each. These 3 scenarios correspond roughly to constant mitigation of slightly over 90% for 450 ppm, almost 60% for 650 ppm, and the IEA's "new policies scenario" (53) for 1,000 ppm. The mapping happens via CRF, interpolating and extrapolating across RCP scenarios (58). We fit a log-function, estimating radiative forcing based on GHG emissions in any 10-y interval as: $5.351 |\log(\text{GHG}) - \log(278)|$. Carbon absorption in any 10-y interval is given by $0.94835 |\text{GHG} - (285.6268 + 0.88414 \sum \text{absorption})|^{0.741547}$. We then interpolate between the 3 GHG levels to find a smooth damage function for any particular level of CRF_t. We assume a linear interpolation of damages between 650 and 1,000 ppm and a quadratic interpolation between 450 and 650 ppm, including a smooth pasting condition at 650 ppm. Below 450 ppm, we assume that climate damages exponentially decay toward zero, setting $S = dp / (l \ln(0.5))$, where d is the derivative of the quadratic damage interpolation function at 450 ppm and $p = 0.91667$ is the average mitigation in the 450 ppm simulation, with l as damage levels (*SI Appendix, Fig. S4C*).

The representative agent knows the distribution of possible final states θ_T . She does not know θ_t for $t < T$. In line with the recombining tree structure (36), climate damages are the probability-weighted average of the interpolated damage function over all final climate states reachable from any one node:

$$D_t(CRF_t, \theta_t) = \sum_{\theta_T} Pr(\theta_T | \theta_t) D_t(CRF_t, \theta_T). \quad [23]$$

Introducing CO₂ removal (38, 52) (Backstop Technology), combined with stochastic climate states θ_t , creates the possibility of GHG concentrations falling below preindustrial levels of 280 ppm. While we know of no analysis that estimates economic damages below 280 ppm, there clearly are costs, much like going (well) above 280 ppm. We introduce a penalty function of the form:

$$f(x) = [1 + e^{k(x-m)}]^{-1}. \quad [24]$$

We arbitrarily set $m = 200$ as the GHG level resulting in half the total penalty. Scalar $k = 0.05$ ensures a smooth penalty function. The combination ensures that $f(280)$ is almost zero, while still achieving a smooth surface. We further restrict $D_t \geq 0$ and $x_t \geq 0$.

Economic Parameters. Fig. 3A shows the CO₂ price sensitivity to RA. With EZ, the CO₂ declines regardless of RA. We choose RA = 7 in our base-case calibration, a value roughly in line with attitudes toward large income risk across wealthy countries (59): RA in the United States alone is often higher (> 8), while RA in European welfare states can be as low as 3. Similar patterns as those displayed in Fig. 3A (and in Fig. 4B for climate damage parameters) hold for mitigation cost parameters (*SI Appendix, Fig. S3*) and an exhaustive list of other economic parameters. The most important of these parameters

appear to be the EIS (Fig. 3B and *SI Appendix, Fig. S5*) and the pure rate of time preference $\delta = (1 - \beta)/\beta$ from Eq. 1 (Fig. 3C and *SI Appendix, Fig. S6*).

Our base case assumes an economic growth rate $\bar{c} = 1.5$ with the rate itself unaffected by climate change, an important assumption to probe in future work (60, 61). *SI Appendix, Fig. S5A* shows that varying it, while keeping EIS = 0.9, has little influence on prices. *SI Appendix, Fig. S5 B and C*, however, shows the large influence of EIS, regardless of assumed \bar{c} . Note that EIS calibrations have changed widely over time, dependent on the type of risk modeled (62, 63). Modern comparable estimates range as high as 1.5 in a model with EZ preferences and consumption shocks (64), with significant implications for CO₂ prices in early years. We choose a lower EIS = 0.9 for our base case, in part because the only shocks to consumption in our model stem from climate risk. Crucially, our model only captures societal

risk. Epstein himself, writing with 2 coauthors, has since offered a potent critique of EZ preferences as applied to individual preferences (65).

ACKNOWLEDGMENTS. For helpful comments and discussions, we thank Jeffrey Bohn, V. V. Chari, Don Fullerton, Ken Gillingham, Christian Gollier, William Hogan, Christos Karydas, Dana Kiku, Gib Metcalf, Robert Socolow, Adam Storeygard, Christian Traeger, Martin Weitzman, Richard Zeckhauser, Stanley Zin, and seminar participants at American Economic Association meetings; the Environmental Defense Fund; ETH Zürich; Global Risk Institute; Harvard; the *Journal of Investment Management* conference; New York University; Tufts; University of Illinois Urbana-Champaign; and University of Minnesota. We also thank Oscar Sjogren, Weiyu Wan, and Shu Ye for helping prepare our code for distribution via <https://gwagner.com/EZClimate>.

- W. D. Nordhaus, An optimal transition path for controlling greenhouse gases. *Science* **258**, 1315–1319 (1992).
- W. D. Nordhaus, Revisiting the social cost of carbon. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 1518–1523 (2017).
- W. Nordhaus, Evolution of modeling of the economics of global warming: Changes in the DICE model, 1992–2017. *Clim. Change* **148**, 623–640 (2018).
- N. H. Stern, *The Economics of Climate Change: The Stern Review* (Cambridge University Press, Cambridge, United Kingdom, 2007).
- R. S. Tol, The economic effects of climate change. *J. Econ. Perspect.* **23**, 29–51 (2009).
- W. D. Nordhaus, P. Sztorc, *DICE 2013R: Introduction and User's Manual* (Yale University, New Haven, CT, 2013).
- National Academy of Sciences, *Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide* (National Academy of Sciences, Washington, DC, 2017).
- M. Burke et al., Opportunities for advances in climate change economics. *Science* **352**, 292–293 (2016).
- Y. Cai, K. L. Judd, T. M. Lenton, T. S. Lontzek, D. Narita, Environmental tipping points significantly affect the cost-benefit assessment of climate policies. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 4606–4611 (2015).
- T. S. Lontzek, Y. Cai, K. L. Judd, T. M. Lenton, Stochastic integrated assessment of climate tipping points indicates the need for strict climate policy. *Nat. Clim. Chang.* **5**, 441–444 (2015).
- Y. Cai, T. M. Lenton, T. S. Lontzek, Risk of multiple interacting tipping points should encourage rapid CO₂ emission reduction. *Nat. Clim. Change* **6**, 520–525 (2016).
- Y. Cai, T. Lontzek, The social cost of carbon with economic and climate risks. *J. Polit. Econ.* <https://doi.org/10.1086/701890> (5 December 2018).
- R. S. Pindyck, Climate change policy: What do the models tell us? *J. Econ. Lit.* **51**, 860–872 (2013).
- N. H. Stern, The structure of economic modeling of the potential impacts of climate change: Grafting gross underestimation of risk onto already narrow science models. *J. Econ. Lit.* **51**, 838–859 (2013).
- G. Wagner, M. L. Weitzman, *Climate Shock: The Economic Consequences of a Hotter Planet* (Princeton University Press, Princeton, NJ, 2015).
- M. Golosov, J. Hassler, P. Krusell, A. Tsyvinski, Optimal taxes on fossil fuel in general equilibrium. *Econometrica* **82**, 41–88 (2014).
- R. Shiller, *Irrational Exuberance* (Princeton University Press, Princeton, NJ, 2000).
- R. Mehra, E. C. Prescott, The equity premium: A puzzle. *J. Monet. Econ.* **15**, 145–161 (1985).
- P. Weil, The equity premium puzzle and the risk-free rate puzzle. *J. Monet. Econ.* **24**, 401–421 (1989).
- R. J. Barro, Rare disasters and asset markets in the twentieth century. *Q. J. Econ.* **121**, 823–866 (2006).
- I. W. R. Martin, On the valuation of long-dated assets. *J. Polit. Econ.* **120**, 346–358 (2012).
- I. W. R. Martin, Disasters and the welfare cost of uncertainty. *Am. Econ. Rev.* **98**, 74–78 (2008).
- D. M. Kreps, E. L. Porteus, Temporal resolution of uncertainty and dynamic choice theory. *Econometrica* **46**, 185–200 (1978).
- L. G. Epstein, S. E. Zin, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica* **57**, 937–969 (1989).
- L. G. Epstein, S. E. Zin, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis. *J. Polit. Econ.* **99**, 263–286 (1991).
- P. Weil, Unexpected utility in macroeconomics. *Q. J. Econ.* **105**, 29–42 (1990).
- R. J. Barro, J. F. Ursua, Stock-market crashes and depressions. *Res. Econ.* **71**, 384–398 (2017).
- F. Ackerman, E. A. Stanton, R. Bueno, Epstein-Zin utility in DICE: Is risk aversion irrelevant to climate policy? *Environ. Resour. Econ.* **56**, 73–84 (2013).
- B. Crost, C. P. Traeger, Optimal CO₂ mitigation under damage risk valuation. *Nat. Clim. Change* **4**, 631–636 (2014).
- C. Hambel, H. Kraft, E. Schwartz, Optimal carbon abatement in a stochastic equilibrium model with climate change, (NBER Working Paper 21044, National Bureau of Economic Research, Cambridge, MA, 2018).
- S. Jensen, C. P. Traeger, Optimal climate change mitigation under long-term growth uncertainty: Stochastic integrated assessment and analytic findings. *Eur. Econ. Rev.* **69**, 104–125 (2014).
- D. Lemoine, C. Traeger, Watch your step: Optimal policy in a tipping climate. *Am. Econ. J. Econ. Policy* **6**, 137–166 (2014).
- R. S. Pindyck, N. Wang, The economic and policy consequences of catastrophes. *Am. Econ. J. Econ. Policy* **5**, 306–339 (2013).
- C. P. Traeger, *Analytic Integrated Assessment and Uncertainty* (Working paper, 2015). <https://ssrn.com/abstract=2667972>.
- D. Lemoine, I. Rudik, Managing climate change under uncertainty: Recursive integrated assessment at an inflection point. *Annu. Rev. Resour. Econ.* **9**, 117–142 (2017).
- J. C. Cox, S. A. Ross, M. Rubinstein, Option pricing: A simplified approach. *J. Financ. Econ.* **7**, 229–263 (1979).
- L. Summers, R. Zeckhauser, Policymaking for posterity. *J. Risk Uncertain.* **37**, 115–140 (2008).
- National Research Council, *Climate Intervention: Carbon Dioxide Removal and Reliable Sequestration* (National Academies Press, Washington, DC, 2015).
- T. M. Lenton et al., Tipping elements in the earth's climate system. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 1786–1793 (2008).
- R. E. Kopp, R. Shwom, G. Wagner, J. Yuan, Tipping elements and climate-economic shocks: Pathways toward integrated assessment. *Earth's Future* **4**, 346–372 (2016).
- A. Ulph, D. Ulph, The optimal time path of a carbon tax. *Oxf. Econ. Pap.* **46**, 857–868 (1994).
- D. Acemoglu, P. Aghion, L. Bursztyn, D. Hemous, The environment and directed technical change. *Am. Econ. Rev.* **102**, 131–166 (2012).
- D. Lemoine, I. Rudik, Steering the climate system: Using inertia to lower the cost of policy. *Am. Econ. Rev.* **107**, 2947–2957 (2017).
- D. Lemoine, The climate risk premium. (Working Paper 15-01, University of Arizona, Tucson, AZ, 2015).
- K. Gillingham, J. H. Stock, The cost of reducing greenhouse gas emissions. *J. Econ. Perspect.* **32**, 53–72 (2018).
- K. Gillingham, K. Palmer, Bridging the energy efficiency gap: Policy insights from economic theory and empirical evidence. *Rev. Environ. Econ. Policy* **8**, 18–38 (2014).
- H. Allcott, M. Greenstone, Is there an energy efficiency gap? *J. Econ. Perspect.* **26**, 3–28 (2012).
- McKinsey, Pathways to a low-carbon economy version 2 of the global greenhouse gas abatement cost curve (McKinsey & Company, Stockholm, 2009).
- N. G. Mankiw, M. Weinzierl, D. Yagan, Optimal taxation in theory and practice. *J. Econ. Perspect.* **23**, 147–174 (2009).
- L. H. Goulder, Environmental taxation and the double dividend: A reader's guide. *Int. Tax Publ. Financ.* **2**, 157–183 (1995).
- D. W. Jorgenson, R. J. Goettle, M. S. Ho, P. J. Wilcoxon, *Double Dividend: Environmental Taxes and Fiscal Reform in the United States* (MIT Press, Cambridge, MA, 2013).
- D. W. Keith, G. Holmes, D. St. Angelo, K. Heidel, A process for capturing CO₂ from the atmosphere. *Joule* **2**, 1573–1594 (2018).
- International Energy Agency, "World energy outlook" (Tech. Rep., International Energy Agency, Paris, 2013).
- R. S. Pindyck, Uncertain outcomes and climate change policy. *J. Environ. Econ. Manag.* **63**, 289–303 (2012).
- M. L. Weitzman, On modeling and interpreting the economics of catastrophic climate change. *Rev. Econ. Stat.* **91**, 1–19 (2009).
- G. H. Roe, Y. Bauman, Climate sensitivity: Should the climate tail wag the policy dog? *Clim. Change* **117**, 647–662 (2012).
- S. Hsiang et al., Estimating economic damage from climate change in the United States. *Science* **356**, 1362–1369 (2017).
- Intergovernmental Panel on Climate Change, "Fifth assessment report: Climate change" (Tech. Rep., Intergovernmental Panel on Climate Change, Geneva, 2013).
- F. Schroyen, K. O. Aarbu, Attitudes towards large income risk in welfare states: An international comparison. *Economica* **85**, 846–872 (2018).
- R. Bansal, M. Ochoa, D. Kiku, "Climate change and growth risks" (NBER Working Paper 23009, National Bureau of Economic Research, Cambridge, MA, 2016).
- G. Heal, J. Park, Temperature stress and the direct impact of climate change: A review of an emerging literature. *Rev. Environ. Econ. Policy* **10**, 1–17 (2016).
- J. Thimme, Intertemporal substitution in consumption: A literature review. *J. Econ. Surv.* **31**, 226–257 (2017).
- T. Havránek, Measuring intertemporal substitution: The importance of method choices and selective reporting. *J. Eur. Econ. Assoc.* **13**, 1180–1204 (2015).
- R. Bansal, A. Yaron, Risks for the long run: A potential resolution of asset pricing puzzles. *J. Financ.* **59**, 1481–1509 (2004).
- L. G. Epstein, E. Farhi, T. Strzalecki, How much would you pay to resolve long-run risk? *Am. Econ. Rev.* **104**, 2680–2697 (2014).