Applying Asset Pricing Theory to Calibrate the Price of Climate Risk

Kent D. Daniel, Robert B. Litterman, and Gernot Wagner*

10 October 2018

Pricing greenhouse gas emissions involves making trade-offs between consumption today and unknown damages in the (distant) future. This setup calls for an optimal control model to determine the carbon dioxide (CO_2) price. It also relies on society's willingness to substitute consumption across time and across uncertain states of nature, the forte of Epstein-Zin preference specifications.

We develop the EZ-Climate model, a simple discrete-time optimization model in which uncertainty about the effect of CO_2 emissions on global temperature and on eventual damages is gradually resolved over time. We embed a number of features including potential tail risk, exogenous and endogenous technological change, and backstop technologies.

The EZ-Climate model suggests a high optimal carbon price today that is expected to decline over time as uncertainty about the damages is resolved. It also points to the importance of backstop technologies and to very large deadweight costs of delay. We decompose the optimal carbon price into two components: expected discounted damages and the risk premium.

JEL code: D81, G11, Q54.

* Daniel: Columbia Business School, 3022 Broadway, New York, NY, 10027, and National Bureau of Economic Research (e-mail: kd2371@columbia.edu); Litterman: Kepos Capital, 620 Eighth Avenue 44th Floor, New York, NY 10018 (e-mail: blitterman@gmail.com); Wagner: Harvard John A. Paulson School of Engineering and Applied Sciences, and Harvard Kennedy School; Harvard University Center for the Environment, 26 Oxford Street, Cambridge, MA, 02138 (e-mail: gwagner@fas.harvard.edu). For helpful comments and discussions, and without any implications, we thank Mariia Belaia, Jeffrey Bohn, Ujjayant Chakravorty, V.V. Chari, Frank Convery, Don Fullerton, Ken Gillingham, Christian Gollier, Cameron Hepburn, William Hogan, Pete Irvine, David Keith, Dana Kiku, Gib Metcalf, Torben Mideksa, Thomas Sterner, Adam Storeygard, Christian Traeger, Martin Weitzman, Richard Zeckhauser, Stanley Zin, and seminar participants at: American Economic Association meetings; Environmental Defense Fund; Global Risk Institute; Harvard University; New York University; Journal of Investment Management (JOIM) conference on Long-run Risks, Returns and ESG Investing; Tufts University; University of Illinois Urbana-Champaign; and University of Minnesota. For excellent research assistance, we thank Oscar Sjogren, Weiyu Wan, and Shu Ye, who were instrumental in preparing our code for open source distribution on GitHub. For code and documentation see: github.com/litterman/EZClimate

1. Introduction

Each ton of carbon dioxide (CO_2) and other greenhouse gasses (GHGs) released into the atmosphere leads to global warming, ocean acidification, and other ecological degradation—all of which impacts societal well-being. The relationship between these damages and GHG emissions is uncertain. The problem will not solve itself without government intervention, as property rights to the atmosphere are poorly defined (Coase, 1960). However, following Pigou (1920), optimal usage of the atmosphere's capacity to absorb GHGs can be obtained, in both theory and practice, when individuals are charged the full social cost of each ton they emit into the atmosphere, or conversely the benefits that accrue to society with the reduction of GHG emissions by one ton. The cost of putting an additional ton of CO_2 into the atmosphere at any given time t, assuming an optimal emissions reductions pathway in the future, is commonly known as the optimal CO_2 price.[1] This paper builds on a long literature in addressing the determination of that price path.

With some notable exceptions, however, that literature largely operates in a deterministic framework. The most famous entrant by far, Nordhaus's (2017a, 2013, 1992) dynamic integrated climate-economy (DICE) model is not, in fact, an optimal control model. Neither are its many derivates. Cai et al's (2016, 2015, 2013) and Golosov et al.'s (2014) models are. We follow their lead, and advice by Lemoine and Rudik (2017a), among others, by moving beyond DICE toward an optimal control model in creating the "EZ-Climate" model.

We approach climate change as a standard asset pricing problem. CO_2 in the atmosphere is an 'asset'—albeit one with negative payoffs. The modern approach to asset pricing recognizes that the optimal CO_2 price is determined by appropriate discounting of the marginal benefits of reducing emissions by one ton at all future times and across all states of nature (Duffie, 2010; Hansen and Richard, 1987). In practice this can be done by discounting those future benefits by a stochastic discount factor appropriate to each possible outcome.

Our choice of the preference specification we use to calculate this stochastic discount factor is dictated by evidence from asset markets. The valuation assigned to different

[1] The assumption of an optimal emissions reductions pathway beginning at $t = 0$ is an important assumption in this and similar modeling exercises (Nordhaus, 2017a, 2013, 1992; Nordhaus and Sztorc, 2013). It is distinct from efforts that calculate the 'Social Cost of Carbon' (SCC), which typically assume no such path. Instead, the SCC focuses on pricing the marginal ton of emissions given the current trajectory (U.S. Government Interagency Working Group on Social Cost of Carbon, 2015). For a recent overview and assessment, see: National Academies of Sciences (2017) and Diaz and Moore (2017). Also note that, throughout this paper, we calculate the optimal price of a ton of CO_2 as opposed to the price of a ton of carbon (C). Given the molecular weights of carbon and oxygen, \$100 per ton of CO_2 is equal to about \$27 per ton of C (= (12/44) \$100). Lastly, while we discuss GHGs more broadly, the calibration itself is, in fact, mostly applicable for long-lived climate forcers, primarily CO_2. Short-lived climate forcers necessitate their own damage function calibration, have their own marginal abatement cost curves, and will, thus, have different optimal pricing pathways (e.g., Marten and Newbold, 2012; Shindell et al., 2017).

traded assets suggests that society is willing to pay only a small premium to substitute consumption across time, but a large premium to substitute across different states of nature. For example, between 1871 and 2012, a portfolio of U.S. bonds earned an average annual real return of 1.6 percent, and a diversified portfolio in U.S. stocks earned an average annual real return of 6.4 percent.

Conversely, society is willing to pay handsomely for the right pattern of cash flows across states: these numbers imply that a portfolio which was short US equities, providing insurance against bad economic outcomes, earned an annual return of *negative* 4.8% per year over this long period. Society presumably discounts equity payoffs at a far higher discount rate because equities earn large returns in good economic times (when marginal utility is low) but often perform poorly precisely when economic growth is low (and marginal utility is high).[2]

The *high* historical equity premium, combined with the low historical volatility of consumption growth, suggests that society is unwilling to substitute consumption across states of nature at some future point in time. In contrast, the *low* risk-free rate in combination with the high average consumption growth rate over the past 150 years suggests that agents are far more willing to substitute consumption across time. These two empirical regularities are inconsistent with a log-normal utility and constant-relative risk aversion (CRRA) preference specification.[3]

These irregularities feature prominently in the equity premium, risk-free rate, and equity volatility puzzles (e.g., Bansal and Yaron, 2004; Mehra and Prescott, 1985; Weil, 1989). There are broadly speaking two responses to these puzzles. One focuses on tail risks and uncertainty, the other on the preference structure itself. Rietz (1988), for example, focuses on extreme events as the driver of the puzzle. For further explorations following this line of thought, see, e.g., Barro (2009, 2006), Barro and Jin (2011), Martin (2008, 2012a), and Weitzman (2007a).

The other approach looks to a richer set of preferences in form of Epstein-Zin utility functions (Epstein and Zin, 1991, 1989; Kreps and Porteus, 1978; Weil, 1990) used throughout the asset pricing literature. Some, such as Barro and Ursua (2008) and Martin (2012b), use both Epstein-Zin preferences and extreme events. We here follow that lead, calibrating an Epstein-Zin utility function while also allowing for potentially large climate risks.

The climate-economic literature has increasingly recognized the importance of exploring a richer set of preferences in calibrating climate risk. Lemoine and Rudik (2017a) survey the early "recursive integrated assessment" literature. An important contribution by Ha-Duong and Treich (2004) explores the importance of Epstein-Zin preferences in general. Ackerman, Stanton, and Bueno (2013) extend the well-known DICE model to incorporate Epstein-Zin preferences and find a significant increase in the optimal CO_2 price as a

[2] Based on data collected by Shiller (2000) and since continuously updated: econ.yale.edu/~shiller/data.htm
[3] Note that CRRA preferences are invariably known as constant-elasticity of substitution (CES) preferences or also as power utility, a special case of which is when utility is a natural-log function of consumption.

result. Other important contributions include work by Christian Traeger, Derek Lemoine and co-authors[4] as well as Cai, Lenton, and Lontzek (2016) and Belaia, Funke, and Glanemann (2017). One broad—perhaps somewhat unfair—conclusion from that work is that the added modeling sophistication and computational power needed compared to, for example, the standard DICE model, may not be justified. Both Crost and Traeger (2014) and Cai *et al.* (2016) find a small role for climate risk in the final optimal CO_2 price figure, which Crost and Traeger (2014) hypothesize might be due to a failure to properly account for climatic disasters. Belaia *et al.* (2017) focus on the effect of incorporating the Atlantic Thermohaline Circulation (THC) tipping point. They, too find that risk aversion is of less importance compared with other factors: "the assumption of higher risk aversion neither changes the near-term policy in a qualitatively meaningful way nor does it affect the THC dynamics to any great extent" (Belaia et al., 2017). Models exist in both finance and climate economics in which the potential for future catastrophic events, though deemed highly unlikely and unseen to date, would lead to significant impacts on current valuations.[5]

DICE itself incorporates the costs of climatic disasters (Nordhaus, 2016, 2013). In particular, ever since the 1999 version of DICE, "catastrophic" damages have accounted for around two-thirds of total economic damages forecast by the model (Kopp et al., 2016). Nordhaus and Boyer (2000) incorporate risk based on an expert survey, asking scientists and economists to assess the probability of a "Great Depression"-size 25% loss of global income due to global average temperatures increasing by 3°C and 6°C (Nordhaus, 1994).[6] That said, DICE, and integrated assessment models like it are simply not equipped to reflect the costs of distant tipping points like the THC. The tipping point occurs too far in the future to matter to today's policy decisions, in part because of discounting, in part because of lack of inertia built into the model (Belaia et al., 2017).[7]

One important lesson to take from prior attempts to build on DICE or introduce alternative models, also echoed by the National Academy of Sciences (2017), is the importance of 'keeping it simple'.[8] DICE has long set the standard for climate-economy modeling because of its simplicity. The core of the model has fewer than 20 equations

---

[4] See e.g. Crost and Traeger (2014), Jensen and Traeger (2014), Lemoine and Traeger (2014), Lemoine and Traeger (2016), Lemoine (2015), and especially Traeger (2015).

[5] For the former, see references in footnote 3. For the latter, see, e.g., Barro (2015), Weitzman (2009, 2007b), and Wagner and Weitzman (2015).

[6] See Kopp *et al.* (2016) for a discussion and critique of this methodology. See e.g. Brock and Hansen (2017), Burke et al. (2016), Convery and Wagner (2015), Kaufman (2012), Pindyck (2013), Stern (2013), and Wagner and Weitzman (2015) for extensive discussions and critiques of prior treatments of risk and uncertainty in climate-economic models in general and U.S. SCC calculations in particular. See Weitzman (2015) for a direct response to Nordhaus's (2013) book-length entry into this debate, and, in turn, Nordhaus's (2015) review of Wagner and Weitzman (2015). Golosov *et al.* (2014), among others, pursues another extension of standard climate-economy models, employing a dynamic stochastic general-equilibrium (DSGE) model, while still relying on a CRRA utility function.

[7] See Nordhaus (1991) for an early discussion of the effects of inertia in climate economics, Lemoine and Rudik (2017b) for a theoretical exploration, and especially Mastrandrea and Schneider (2001) for a comprehensive discussion of the implications.

[8] National Academy of Sciences (2017) emphasizes the importance of modularity of the modeling effort to ensure natural points of entry for multi-disciplinary collaborations providing crucial inputs into the core model.

(Nordhaus and Sztorc, 2013). Implementations of Epstein-Zin utility functions typically do not lend themselves to this level of simplicity, ease of use, and transparency. It also means DICE needs to cut corners. In fact, Nordhaus and Sztorc (2013) cite Epstein and Zin (1991, 1989), but only while warning not to confuse a CRRA utility as representing risk aversion explicitly. The implementation of DICE reverts to the CRRA utility function.

In keeping with this missive for simplicity, we introduce an optimal control problem employing a discrete-time binomial tree model. This approach will be familiar to financial economists, who employ such trees in many financial-economic modeling applications (see Cox, Ross and Rubinstein (1979) for an early example). We build on a related approach in Summers and Zeckhauser (2008), who employ a simple three-period model—sans Epstein-Zin preferences—to isolate some key features of clearing up uncertainty over time.

While striving for simplicity in itself is an important goal, Figure I shows the implications of employing Epstein-Zin preferences: CRRA and Epstein-Zin preferences result in wildly divergent optimal CO_2 price pathways for different levels of risk calibration, holding everything else constant.[9]

Standard CRRA specifications embed the assumption that agents' willingness to substitute consumption across states of nature is the same as their willingness to substitute consumption over time. Thus, an increase in the coefficient of risk aversion (or, conversely, a decreased elasticity of substitution across states), is necessarily linked to a decreased elasticity of intertemporal substitution (EIS). Given the fact that consumption grows at a rate of about 1.5% per year, an unwillingness to substitute across time leads to a (counterfactually) high risk-free discount rate. Since consumption damages occur far into the future, a CRRA utility function with a high level of risk-aversion (and a reasonable rate of time preference) implies a high discount rate for these damages, and a low optimal CO_2 price—tending toward zero.

In contrast, Epstein-Zin utility allows for separation of the coefficient of risk-aversion and the EIS, consistent with the equity-premium/risk-free rate puzzle. With an Epstein-Zin specification, holding the EIS fixed at 0.9 and increasing the degree of risk aversion, the optimal CO_2 price increases, while the real interest rate remains at around 3.11%/year.[10]

---

[9] See part II for our EZ-Climate model setup and calibration. Our original calibration (Daniel et al., 2016), was anchored around the U.S. SCC of $40 for a ton of CO_2 released in 2015, in 2015 US$, the central value calculated by the U.S. Government Interagency Working Group on Social Cost of Carbon (2015). While further revising the paper, we introduced radiative forcing as a stock variable and added carbon cycle feedbacks explicitly (see section II.C). That step alone increased the 2015 number from $40 to over $100 in our base case. Even without any tail risks, the 2015 optimal CO_2 price never dips below $60, well above the U.S. SCC figure. Note also that our optimal CO_2 price is distinct from the U.S. SCC (see footnote 1).

[10] The exact interest rate in our Epstein-Zin calibration is almost independent of the risk aversion (RA) coefficient. Bansal and Yaron (2004) are able to match the equity premium with a far lower coefficient of risk-aversion, owing to the presence of shocks to the long-term growth rate of consumption in their model, which are correlated with equity returns. Similarly, in the EZ-Climate model presented here, a link between higher climate fragility and lower consumption growth rates would lead to a higher optimal CO_2 price with a lower coefficient of risk-aversion. In general, climate damages hitting growth rates rather than levels of GDP can have a significant effect on the optimal CO_2 price (Bansal and Ochoa, 2011; Dell et al., 2012; Diaz and Moore, 2017; Heal and Park, 2016; Moore and Diaz, 2015; Wagner and Weitzman, 2015).

As the level of risk aversion is raised from a very low level to a level consistent with the historically observed equity-risk premium, the optimal CO_2 price increases by around 50% (Figure I).[11]

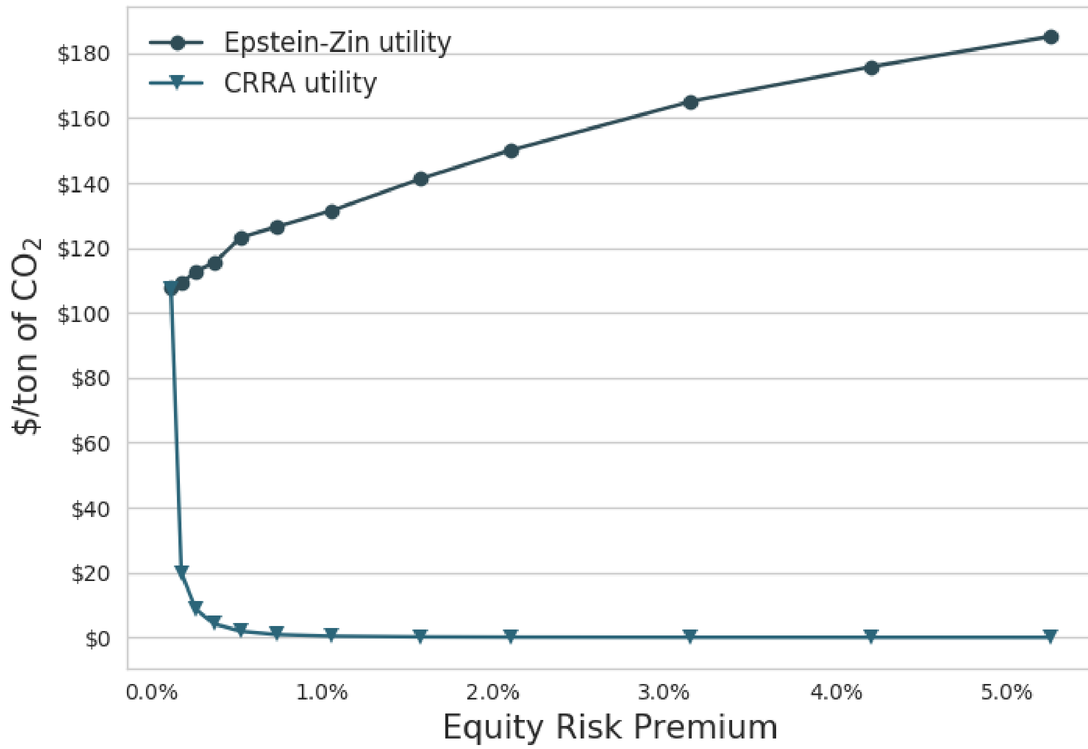

Figure I—Using Epstein-Zin utility functions results in increasing optimal 2015 CO_2 prices, in 2015 US$, with increasing risk aversion, translated into the implied equity risk premium using Weil's (1989) conversion, while holding implied market interest rates stable at 3.11%

We first introduce the EZ-Climate model and its calibration (Section 2) before presenting results, risk decomposition, and sensitivity analyses (Section 3), and discussing future research and extensions (Section 4). We conclude with an analogy (Section 5), grounded in the result presented by Figure I. Climate policy treated as an asset pricing problem is, after all, fundamentally about risk mitigation.

2. The Model

Our representative agent solves the optimization problem of trading off the (known) costs of climate mitigation against the uncertain future benefits associated with mitigation. She maximizes lifetime expected utility at each time and for each state of nature by choosing the optimal mitigation at time t, $x_t^*(\theta_t)$, dependent on the current estimate of the Earth's fragility, θ_t, and on the future evolution of θ_t. Fragility θ_t evolves stochastically as described in section 2.3.

[11] Figure VIII in section III shows the optimal CO_2 price over time for our base case calibration of an EIS of 0.9 and $RRA = 7$.

Mitigating emissions is costly to any individual, but the resulting future benefits are dispersed across society. Hence, assuming no government action to price carbon, atomistic agents do zero mitigation. We calculate the optimal price on carbon emissions, as the price which would induce atomistic agents to reduce emissions to the level that would be chosen by the representative agent at each time and in each state.

As GHGs build up in the atmosphere, temperatures rise. As a result, a fraction of the baseline consumption is lost to damages. The damages as a function of mitigation are not known ex-ante. They are, in turn, a function of θ_t. Each period of the model, agents learn more about the level of fragility, but they only know the actual fragility in the final two periods of the model.

These assumptions simplify reality in two important ways: As θ_t is the only unknown in EZ-Climate model, we do not allow for interactions of shocks to fragility with those to other state variables (e.g., productivity). The second simplification is the assumption of full knowledge of θ in period $T-1$ (in the year 2300 in our base case). Important aspects of climate science are deeply and persistently uncertain, and science may not learn the true θ at a time scale relevant to policy (Wagner and Zeckhauser, 2017; Zeckhauser, 2006). We compromise by having the complete resolution of uncertainty delayed until 2300, when we might be able to expect to know the all-important climate sensitivity parameter, what happens to global average temperatures, in equilibrium, as concentrations of GHGs double from pre-industrial levels.[12]

The setting for the EZ-Climate model is a discrete time, endowment economy with a single representative agent. In each period $t\epsilon\{0,1,2,\dots,T\}$, the representative agent is endowed with a certain amount of the consumption good, $\bar{c}_t$. However, she is not able to consume the full endowed consumption for two reasons: climate change and climate policy.

In periods $t\epsilon\{1,2,\dots,T\}$, some of the endowed consumption may be lost due to climate change damages. In periods $t\epsilon\{0,1,2,\dots,T-1\}$, the agent may elect to spend some of the endowed consumption to reduce her impact on the climate. The resulting consumption c_t, after damages D_t and mitigation costs κ_t are taken into account, is given by:

(1) $$c_0 = \bar{c}_0 \cdot \big(1-\kappa_0(x_0)\big),$$
(2) $$c_t = \bar{c}_t \cdot \big(1-\kappa_t(x_t)\big) \cdot \big(1-D_t(CRF_t,\theta_t)\big), \text{ for } t\epsilon\{1,2,\dots,T-1\}, \text{ and}$$
(3) $$c_T = \bar{c}_T \cdot (1-D_T(CRF_T,\theta_T)).$$

In equations (2) and (3), the climate damage function $D_t(CRF_t,\theta_t)$ captures the fraction of endowed consumption that is lost due to damages from climate change. If

[12] The ‘likely’ range for climate sensitivity has been 1.5–4.5°C ever since Charney *et al.* (1979)—with one short exception. In its Fourth Assessment Report, IPCC (2007) narrowed the range to 2–4.5°C, only to be expanded back to the prior range in the Fifth Assessment Report (IPCC, 2013). The “equilibrium” climate sensitivity range is the so-called “fast” equilibrium, as distinct from Earth system sensitivity. The latter includes broader changes, which could, in turn, result in still larger changes. Previdi *et al.* (2013) present a range of 6–8°C for a doubling of atmospheric CO_2 concentrations. (See section II.C.ii.)

$D_t(CRF, \theta_t) = 0$, the agent would receive the full consumption endowment. However, damages from climate change can push D_t above zero. D_t, in turn, depends on two variables: CRF_t, which we define as the cumulative radiative forcing up to time t, which determines global average temperature, and θ_t, the Earth's fragility which, as discussed earlier, characterizes the uncertain relation between GHG concentrations and consumption damages.

Cumulative radiative forcing, CRF_t, in turn, depends on the level of mitigation in each period $x_s$from $s = 0$ to t. An $x_s = 1$ would mean that, in period s, GHG emissions are reduced to zero. Mitigation of $x_s = 0$ is "business as usual," meaning that individuals and business do not face any taxes or other restrictions on GHG emissions. Mitigation is further discussed in Section 2.3.3.

Mitigation reduces the stock of GHGs in the atmosphere and leads to lower climate damages, and, hence, to higher future consumption. But mitigating GHG emissions is costly. Mitigating a fraction x_t of emissions costs a fraction $\kappa_t(x_t)$ of the endowed consumption. We describe the details of the cost function, and our calibration, in Section 2.2.

Cumulative mitigation between periods 0 and t is given by:

(4) $$\mathrm{X_t} = \frac{\sum_{s=0}^{t} g_s \cdot x_s}{\sum_{s=0}^{t} g_s},$$

where g_s is the flow of GHG emissions into the atmosphere in period s, for each period up to t, absent any mitigation.[13] Cumulative mitigation, $\mathrm{X_t}$, enters the determination of the rate of technological change, discussed in Section 2.2.2.

Rather than restricting mitigation x_t to be below 1, in our baseline analysis we allow for the use of a *backstop technology* to pull CO_2 directly out of the atmosphere, potentially leading to $x_t > 1$. Backstop technologies are typically labeled carbon dioxide removal (CDR) or direct carbon removal (DCR). See the discussion in sections 2.2.1 on backstop technologies and 2.3.4 on the resulting possibility of having concentrations fall below 280 ppm.

To make the solution tractable, the EZ-Climate model employs a binominal tree for the resolution of uncertainty about climate damage, discussed in detail in Section 2.4. The baseline analysis uses a 7-period tree, beginning in 2015. An initial mitigation decision is made in 2015, and subsequent mitigation decisions are made after information is revealed about climate fragility and the resulting damages in years 2030, 2060, 2100, 2200, and 2300. The final period, in which consumption simply grows at a constant rate, begins in 2400 and lasts forever. At each node of the tree, more information about the consumption damage function is revealed (as reflected in the fragility parameter θ_t), but uncertainty is

[13] The cumulative GHG emissions that must be absorbed into the atmosphere or oceans is $G_t(1 - X_t)$, where $G_t = \sum_{s=0}^{t} g_s$ denotes the cumulative emissions under the BAU scenario.

not fully resolved until the beginning of the next-to-last period in 2300.[14] The agent's utility in each state is calculated based on interpolated consumption flows at five-year sub-periods. We solve for mitigation levels over time that maximize expected utility, looking forward, at the start of each period (except the final period), and in each fragility state θ_t. The resulting optimal CO_2 price in each period and state is the price that implements this level of mitigation.

In the next section, we describe the agent's preferences, and provide some additional motivation for the preferences specification we employ. In Sections 2.2 and 2.3, we lay out the cost and damage functions, respectively, and describe their calibration. In keeping with the National Academy of Science's (2017) call for modularity in climate-economy modeling efforts, these calibrations are largely independent of each other and could easily be swapped for different, better calibrations. Section 2.4 describes EZ-Climate's tree structure in more detail. It, too, can be readily adjusted, in keeping with the modularity of EZ-Climate. Section 3 presents the results based on our calibrations. Section 5 concludes.

2.1. Preferences

As noted in the introduction, the CRRA specification typically used in climate-economy models like DICE embeds the assumption that agents' willingness to substitute consumption across states of nature is the same as their willingness to substitute consumption over time. This is inconsistent with the observed low risk-free rate and high equity premium (Mehra and Prescott, 1985; Weil, 1989). To resolve this puzzle, financial economists have begun to employ the preference specification suggested by Epstein and Zin (1991, 1989) and Weil (1990) that allows for different rates of substitution across time and states.[15] This is the specification underlying EZ-Climate.

---

[14] Possibly unique among climate-economy models, our setup of allowing for negative emissions creates the possibility of optimal GHG concentrations going below 280 ppm—*if*, in the final stage of our tree structure fragility, θ_t, turns out to be worse than expected. We, thus, include another damage component with increasing damage function for GHG concentrations below a certain level in EZ-Climate, also perhaps unique among climate-economy models. See Section II.C for a more detailed description of climate damages.

[15] See Bansal and Yaron (2004) and Hansen, Heaton, and Li (2008) for more detailed discussions. Bansal and Ochoa (2011, 2009) and Bansal, Ochoa and Kiku (2016) use this preference specification in combination with a framework in which temperature shocks affect future consumption growth. For example, Ackerman, Stanton, and Bueno (2013), Crost and Traeger (2014), and, most recently, Cai, Lenton, and Lontzek (2016) use this utility function in DICE. Making no further adjustments, doing so poses oft-significant computational challenges. Ha-Duong and Treich (2004) appear to have been the first to encounter those challenges. They explore the importance of Epstein-Zin preferences in calibrating climate risk in general and conclude, somewhat too cautiously, that setting $\rho = \alpha$ "may misinterpret the sensitivity of the climate policy to risk-aversion." Importantly, Ha-Duong and Treich (2004) explore a more general implementation of Epstein-Zin preferences, turning c_t also into a function equal to the certainty-equivalent of future lifetime income. Their equivalent of our equation (5) could have been written more generally as $U_t = \left[(1-\beta)[\mu_t(\tilde{c}_t)]^\rho + \beta\left[\mu_t(\widetilde{U}_{t+1})\right]^\rho\right]^{1/\rho}$, with the definition of $\mu_t(\tilde{c}_{t+1})$ mirroring our equation (6), for the current period t: $\mu_t(\tilde{c}_t) = (E_t[\tilde{c}_t^\alpha])^{1/\alpha}$. The difference is subtle but potentially important. Ha-Duong and Treich's (2004) extension allows for consumption to be uncertain within each period. Our and others' implementation of Epstein-Zin preferences in form of (5) and (6) implies full knowledge of each period's

In an Epstein-Zin utility framework, the agent maximizes at each time t:

(5) $$U_t = \left[(1-\beta)c_t^{\rho} + \beta\left[\mu_t\left(\tilde{U}_{t+1}\right)\right]^{\rho}\right]^{1/\rho},$$

where $\mu_t\left(\tilde{U}_{t+1}\right)$ is the certainty-equivalent of future lifetime utility, based on the agent's information at time t, and is given by:

(6) $$\mu_t\left(\tilde{U}_{t+1}\right) = \left(E_t[U_{t+1}^{\alpha}]\right)^{1/\alpha}.$$

In this specification, $(1-\beta)/\beta$ is the pure rate of time preference, commonly denoted by δ. The parameter ρ measures the agent's willingness to substitute consumption across time. The higher is ρ, the more willing the agent is to substitute consumption across time. The elasticity of intertemporal substitution is given by $\sigma = 1/(1-\rho)$.

Finally, α captures the agent's willingness to substitute consumption across (uncertain) future consumption streams. The higher is α, the more willing the agent is to substitute consumption across states of nature at a given point in time. The coefficient of relative risk aversion at a given point in time is $\gamma = (1-\alpha)$. This added flexibility allows for calibration across states of nature and time. With $\rho = \alpha$, equations (5) and (6) are equivalent to the standard CRRA utility specification.

Plugging (6) into (5) results in EZ-Climate's utility specification:

(7) $$U_0 = \left[(1-\beta)c_0^{\rho} + \beta\left(E_0\left[\tilde{U}_1^{\alpha}\right]\right)^{\rho/\alpha}\right]^{1/\rho}$$

(8) $$U_t = \left[(1-\beta)c_t^{\rho} + \beta\left(E_t[U_{t+1}^{\alpha}]\right)^{\rho/\alpha}\right]^{1/\rho}, \text{ for } t\epsilon\{1,2,\ldots,T-1\}.$$

with c_0 and c_t, respectively, given by equations (1) and (2).

In the final period, which, in our base case, is the period starting in 2400, the agent receives the utility from all consumption from time T forward. Given our assumption that all uncertainty has been resolved at this point, consumption grows at a constant rate r from T through infinity:

(9) $$c_t = c_T(1+r)^{t-T} \text{ for } t \geq T.$$

The resulting final-period utility is:

(10) $$U_T = \left[\frac{1-\beta}{1-\beta(1+r)^{\rho}}\right]^{1/\rho} c_T,$$

with c_T given by equation (3).

consumption at the time of the mitigation decision. This is a reasonable assumption for short time intervals. As those intervals become larger, within-period uncertainty might become more important.

2.2. Mitigation Cost

Calibrating the economic cost side of EZ-Climate requires specifying a relationship between the marginal cost of emissions reductions or per-ton tax rate, τ, the resulting flow of emissions in gigatonnes of CO_2-equivalent emissions per year (Gt CO_2e), $g(\tau)$, and the fraction of emissions reduced, $x(\tau)$.

Many modeling efforts have attempted to estimate the marginal abatement costs, often as part of integrated assessment models and based on a number of assumption.[16] Perhaps the most influential, independent effort comes from McKinsey & Company in an attempt to estimate a bottom-up marginal abatement cost curve (MACC). McKinsey's MACCs are, to a large extent, based on bottom-up 'engineering' estimates. That makes them an easy target for critique by economists, which often focuses on the large abatement opportunities with 'negative' costs—positive net present value. McKinsey's report on energy efficiency in the United States identifies energy efficiency savings opportunities with positive net present value commensurate with emissions reductions of over 1 Gt CO_2e per year (McKinsey, 2009a). Lots of effort has gone into assessing (Gerarden et al., 2015) and helping to bridge (Gillingham and Palmer, 2014) that potential energy efficiency gap, without conclusive evidence (Allcott and Greenstone, 2012). These critiques notwithstanding, McKinsey's effort stands as a unique, bottom-up, data-driven attempt at estimating abatement costs. We calibrate τ, $g(\tau)$, and $x(\tau)$ in EZ-Climate based on McKinsey's global MACC effort (McKinsey, 2009b), with one crucial modification: We assume no mitigation ($x(\tau) = 0$) at $\tau \leq 0$; i.e. no net-negative or zero-cost mitigation. Table I shows the resulting calibration.[17]

Table I—Marginal abatement cost curve for 2030, based on McKinsey (2009b), modified to have $x(\tau)=0$ for $\tau\leq 0$.

GHG taxation rate τ	GHG emissions flow $g(\tau)$	Fractional GHG reduction $x(\tau)$
€0/ton	70 Gt CO_2e/year	0
€60/ton	32 Gt CO_2e/year	0.543
€100/ton	23 Gt CO_2e/year	0.671

Fitting McKinsey's modified point estimates (in $US using an average 2005 exchange rate of 1.206 $ per €) from Table I to a power function for $x(\tau)$ yields:

(11) $$x(\tau) = 0.0923 \cdot \tau^{0.414}.$$

The corresponding inverse function, solving for the appropriate tax rate to achieve x is:

---

[16] See Stanford's Energy Modeling Forum as one such major effort, working with a number of different models: https://emf.stanford.edu. See Huntington (2011) as an overview of Stanford EMF's work focused on the 'energy efficiency gap'.

[17] We have emissions stabilize at 57% above current levels. In our 'unmitigated' baseline scenario, following the IEA's New Policies Scenario (which does, in fact, have what we would describe as modest mitigation), GHG concentrations reach approximately 1,000 ppm by 2200.

(12) $$\tau(x) = 314.32 \cdot x^{2.413}.$$

Equation (12) shows the marginal cost of abatement. 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. Intuitively, GHG emissions are an input to the production process that generates consumption goods. Assuming the agent chooses the level of GHG emissions $g(\tau)$ so as to maximize consumption c given τ, the marginal cost of increasing the tax rate must be the quantity of emissions at that tax rate, that is:

(13) $$\frac{dc(\tau)}{d\tau} = -g(\tau),$$

Thus, to calculate the consumption associated with a particular tax rate of τ, we integrate (13), resulting in:

(14) $$c(\tau) = \bar{c} - \int_0^\tau g(s)\, ds,$$

where $\bar{c}$ is the endowed level of consumption (assuming zero climate damages). However, this equation is correct only if the GHG tax is purely dissipative—that is, if the government were to collect the tax and then waste 100% of the proceeds. In our analysis, we make the opposite assumption: the proceeds of the tax $(g(\tau) \cdot \tau)$ are refunded lump-sum.[18] That makes the decrease in consumption equal to the distortionary effect of the tax (in dollars):

(15) $$K(\tau) = \int_0^\tau g(s)\, ds - g(\tau) \cdot \tau.$$

Writing $g(\tau) = g_0\big(1 - x(\tau)\big)$, where g_0 is the baseline level of GHG emissions, we can rewrite (15) as:

$$\begin{aligned} K(\tau) &= g_0 \int_0^\tau \big(1 - x(s)\big) ds - \tau g_0\big(1 - x(\tau)\big) \\ &= g_0 \left[\tau - \int_0^\tau x(s) ds\right] - \tau g_0 + \tau g_0 x(\tau) \\ &= g_0\left[\tau x(\tau) - \int_0^\tau x(s) ds\right] \end{aligned} \tag{16}$$

Substituting (11) into (16) and simplifying gives the total cost K as a function of the tax rate τ:

$$\begin{aligned} K(\tau) &= g_0[0.09230 \cdot \tau^{1.414} - 0.06526 \cdot \tau^{1.414}] \\ &= g_0 \cdot 0.02704 \cdot \tau^{1.414}, \end{aligned} \tag{17}$$

Substituting (12) into (17) gives K as a function of fractional-mitigation x:

[18] Note that were the proceeds from the (Pigouvian) GHG tax used to reduce other distortionary taxes, the effective cost of the carbon tax would be still lower than what we calculate here, and thus would justify a higher optimal τ. For a summary of this "double-dividend" argument, see Goulder (1995).

$$K(x) = g_0 92.08 \cdot x^{3.413}, \quad (18)$$

where total cost $K(x)$ is expressed in dollars. Finally, we divide by current (2015) aggregate consumption to determine the cost as a fraction of baseline consumption:

$$\kappa(x) = \left(\frac{g_0 \cdot 92.08}{c_0}\right) \cdot x^{3.413}, \quad (19)$$

where $g_0 = 52$ Gt CO_2e represents the current level of global annual emissions, and $c_0 =$ \$31 trillion/year is current (2015) global consumption in 2015 dollars. Equation (19) expresses the societal cost of a given level of mitigation as a percentage of consumption. We assume that, absent technological change, the function $\kappa(x)$ is time invariant. See 2.2.2 for our discussion of technological change. First, we explore the impact of backstop technologies.

2.2.1. Backstop Technology

The McKinsey estimates on which our total cost function, $\kappa(x)$, are based reflect the cost of traditional abatement technologies. In addition, carbon dioxide removal (CDR) can pull CO_2, and potentially GHGs, directly from the atmosphere (National Research Council, 2015a). We label these *backstop technologies*.

We assume our backstop technology is available at a marginal cost of τ^* for the first ton of carbon that is removed from the atmosphere and that unlimited amounts of CO_2 can be removed as the marginal cost approaches $\tilde{\tau} \geq \tau^*$. In fitting the marginal cost curve to τ^* and $\tilde{\tau}$ we build a marginal cost function for the backstop technology of the form:

$$B(x) = \tilde{\tau} - \left({k}/{x}\right)^{1/b}. \quad (20)$$

The upper bound of the cost function is, thus, $\tilde{\tau}$. We calibrate (20), such that:

$$B(x_0) = \tilde{\tau} - \left({k}/{x_0}\right)^{1/b} = \tau^*, \quad (21)$$

which allows us to express:

$$k = x_0(\tilde{\tau} - \tau^*)^b, \quad (22)$$

where x_0 is the point at which the backstop technology begins to be used. We also impose a smooth-pasting condition at x_0; i.e. the derivative of the marginal cost curve is continuous at x_0. This allows us to solve for parameter b:

$$b = \frac{\tilde{\tau} - \tau^*}{(3.413 - 1)\tau^*}. \quad (23)$$

Our base case assumes $\tau^* = \$2{,}000$ and $\tilde{\tau} = \$2{,}500$ in 2015 dollars. Under the most aggressive backstop scenario presented in the results section, we assume $\tau^* = \$300$ and $\tilde{\tau} = \$350$ in 2015 dollars. These aggressive values imply that the backstop technology kicks in at mitigation levels above around 100%, whereas our base case all but assures that backstop technologies do not get used for a considerable period of time. The $350 is on the low end of possible assumptions, and the extent to which there can be a true backstop technology remains uncertain (Socolow et al., 2011). The need for one is clear, and much work remains to be done to demonstrate possible technologies and refine price estimates (Keith, 2000; National Research Council, 2015a).[19] Figure II shows the marginal cost, $\tau(x)$, with both base-case and aggressive backstop technology assumptions.

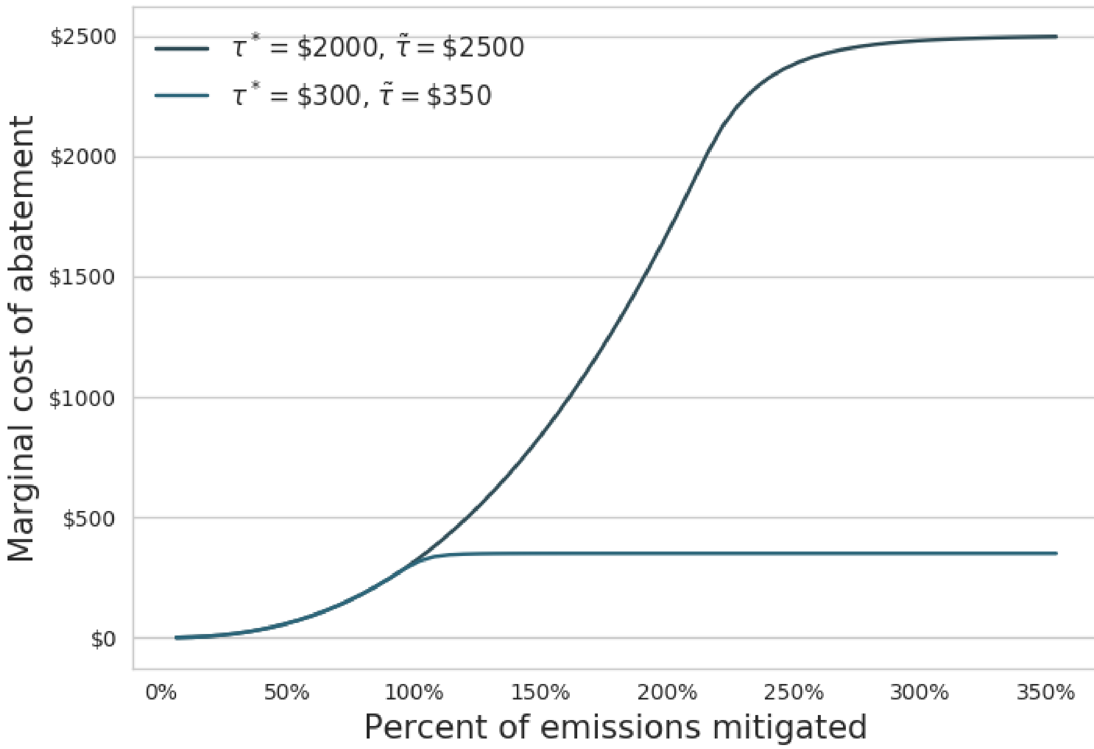

Figure II—Marginal cost of abatement, $\tau(x)$, in 2015 (in 2015 \$/ton) under base-case ($\tau^* = \$2{,}000$, $\tilde{\tau} = \$2{,}500$) and aggressive backstop technology ($\tau^* = \$300$, $\tilde{\tau} = \$350$) assumptions

2.2.2. Technological Change

These cost curves are calibrated to $t = 0$. In subsequent periods, we allow the marginal cost curve to decrease at a rate determined by a set of technological change parameters: a constant component, φ_0, and a component linked to mitigation efforts to date, $\varphi_1 X_t$, where X_t is the average mitigation up to time t, defined by equation (4). Thus, at time t, the total cost curve is given by:

[19] Note that CDR as a backstop technology is entirely distinct from solar geoengineering, also known as ‘albedo modification’ or ‘solar radiation management’ (National Research Council, 2015b). The one commonality is that both ‘carbon geoengineering’, CDR, and the potential availability of solar geoengineering likely reduce the optimal CO_2 price. EZ-Climate models the former, not the latter.

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

This functional form allows for easy calibration. For example, if $\varphi_0 = 0.005$ and $\varphi_1 = 0.01$, and with average mitigation of 50%, marginal costs decrease as a percentage of consumption at a rate of 1% per year.

2.3. Damage Function Specification

We next specify the climate damage function $D_t(CRF_t, \theta_t)$.[20] Damages are a function of temperature changes, which, in turn, are a function of cumulative solar radiative forcing CRF_t, which, in our setting, are determined by the mitigation path up to that point in time. We then compare CRF_t to three baseline emissions paths, g_t for which we have created associated damage simulations. The one way, then, to affect the level of damages is to change mitigation across time, x_t. The specification of damages has two components: a non-catastrophic component and an additional catastrophic component triggered by crossing a particular threshold. The hazard rate associated with hitting that threshold increases with temperature. If the threshold is crossed at any time, additional damages decrease consumption in all future periods.

We calculate the overall damage function $D_t(CRF_t, \theta_t)$ for the baseline emissions paths, g_t, using Monte-Carlo simulation. As we describe in detail below, we run a set of simulations for each of three constant mitigation levels X_t, which determine cumulative radiative forcing at each point in time. In each run of the simulation, we draw a set of random variables: [1] global average temperature change; [2] the parameter characterizing non-catastrophic damages as a function of temperature; [3] an indicator variable that determines whether or not the atmosphere hits a tipping point at any particular time and state, and [4] the tipping point damage parameter. The state variable θ_t indexes the distribution resulting from these sets of simulations, and interpolation across the three mitigation levels gives us a continuous function D_t across cumulative radiative forcing levels CRF_t.

2.3.1. Temperature as a Function of GHG Levels

The distribution of temperature outcomes as a function of mitigation strategies is calibrated to three carbon scenarios, indexed by a maximum level of CO_2 in the atmosphere. For the subsequent base case calibration, we follow Weitzman (2009) and Wagner and Weitzman (2015) in calibrating a log-normal distribution for equilibrium climate sensitivity—the eventual temperature rise as atmospheric concentrations of CO_2

[20] Our damage function calibration follows the basic logic of Pindyck (2012), with one crucial exception: Pindyck (2012) assumes gamma distributions for temperature levels given greenhouse gas concentrations, and for economic damages given temperature levels. We explore other functional forms for both. See sections II.C.i and II.C.ii.

double.[21] Specifically, Wagner and Weitzman (2015) calibrate a log-normal function assuming a 78% probability of climate sensitivity in the 1.5-4.5°C "likely" range.[22] Moreover, the Intergovernmental Panel on Climate Change (IPCC)'s Fifth Assessment Report (IPCC, 2013) judges climate sensitivity above 6°C to be "very unlikely," giving it a 0-10% probability. We again follow Wagner and Weitzman (2015) in assigning it a roughly 5% chance.

Wagner and Weitzman (2015) then use this calibration to translate the International Energy Agency's (IEA) projections for concentrations of CO_2-equivalent tons into final temperature outcomes. Under the assumptions of the its "new policies scenario," IEA (2013) projects that atmospheric concentrations will reach 700 ppm CO_2e by 2100. That concentration would result in a projected, eventual median temperature increase of 3.6°C. Wagner and Weitzman (2015) present eventual median temperature outcomes for concentrations of between 400 and 800 ppm. We take their calibration and extrapolate to 1000 ppm, which we assume to be the zero-mitigation scenario, marking an upper bound of sorts. We similarly assume that 100% mitigation over time leads to a maximum GHG level of 400 ppm. Other fixed levels of mitigation are assumed to lead to damages associated with GHG levels linearly interpolated between those levels. Thus, mitigation of 50% through any point in time leads to the interpolated damages at that time along a path associated with a maximum GHG level of 700 ppm.

Table II gives the probability of different levels of ΔT_{100}—the temperature change over the next 100 years—for given maximum levels of GHGs in atmosphere. The 450 ppm, 650 ppm, and 1000 ppm maximum levels of CO_2 equivalents in the atmosphere reflect, respectively, a strict, a modest, and an ineffective mitigation scenario.

Table II—Probability of $\Delta T_{100} > T$

	Maximum GHG Level (ppm of CO_2)		
T	450	650	1000
2°C	**0.40**	**0.85**	**0.99**
3°C	**0.13**	**0.54**	**0.86**
4°C	**0.04**	**0.30**	**0.66**
5°C	**0.02**	**0.15**	**0.46**
6°C	**0.00**	**0.07**	**0.30**

We then use assumptions akin to Pindyck (2012) to fit a displaced gamma distribution around final GHG concentrations, while setting levels of GHG 100 years in the future equal to equilibrium levels. Table III gives the parameters for these distributions, and the probabilities from the fitted displaced gamma distributions, which line up well with the

[21] This log-normal calibration results in similar CO_2 price estimates as a distribution calibrated by Roe and Baker (2007). It results in higher CO_2 price estimates compared with Pindyck's (2012) gamma distribution calibration.
[22] The IPCC says that range is "likely," which it defines as having at least a 66% probability. The IPCC's "very likely" designation implies at least a 90% probability. We follow Wagner and Weitzman (2015) in splitting the difference to arrive at 78%.

numbers in Table II, especially for scenarios closer to 450 and 650 ppm than the 1,000 ppm zero-mitigation case.

Table III—Fitted values of $Prob(\Delta T_{100} > T)$ for three specified gamma distributions

	Maximum GHG Level (ppm of CO2)		
T	450	650	1,000
2°C	**0.40**	**0.87**	**0.99**
3°C	**0.14**	**0.57**	**0.91**
4°C	**0.04**	**0.29**	**0.70**
5°C	**0.01**	**0.12**	**0.44**
6°C	**0.00**	**0.05**	**0.24**
Gamma distribution parameters			
$\boldsymbol{\alpha}$	2.810	4.630	6.100
$\boldsymbol{\beta}$	0.600	0.630	0.670
Displacement	-0.25	-0.5	-0.9

To obtain the temperature distribution at other times, we again follow Pindyck (2012), and specify that the time path for the temperature change at time t (in years) is given by:

(25) $$\Delta T(t) = 2\,\Delta T_{100}\left[1 - 0.5^{\frac{t}{100}}\right].$$

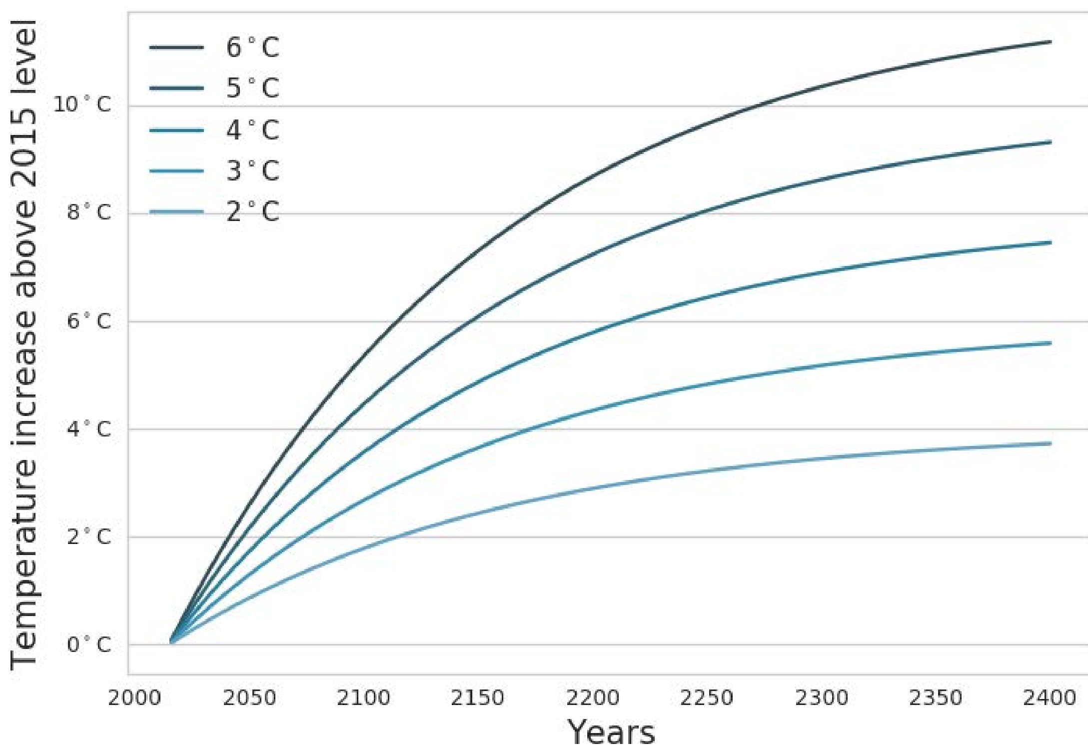

Figure III—Calibrated time path for temperature increases given assumed temperature increases within a century

Figure III plots temperature paths for different levels of ΔT_{100}. As time increases, the temperature change asymptotes to double the value of ΔT_{100}. Even though these calibrations are, by now, 'established' in the climate-economic literature, both the distribution of ΔT_{100} and the functional form for the path in equation (25) clearly merit further scientific scrutiny. Both are likely on the conservative side of actual projections.

2.3.2. *Damages as a Function of Temperature*

The next step is to translate average global surface warming into global mean economic losses via the damage function D_t. There are two components to D_t: a *non-catastrophic* and a *catastrophic* one. The functional form of each component is known to the agent. However, as with the GHG-ΔT_{100} relationship discussed in the prior section, the functional form for each damage function component contains a parameter that characterizes the uncertainty in our present understanding of this relationship. In EZ-Climate the agent knows the form of the distribution of this parameter at the initial date, and in each period she learns more about the distribution of the parameter. However, the final realization of the parameter is not known until the next-to-last period.

The non-catastrophic component of our damages follows Pindyck (2012), who fits a functional form to data from the IPCC's Fourth Assessment Report (IPCC, 2007), and obtains a loss function of the form:

$$\text{(26)} \qquad \mathrm{L}(\Delta T(t)) = e^{-13.97 \cdot \gamma \cdot \Delta T(t)^2},$$

where γ is drawn from a displaced gamma distribution with parameters $r = 4.5$, $\lambda = 21341$, and $\theta = -0.0000746$.

Based on non-catastrophic damages, consumption at any time t is reduced as follows:

$$\text{(27)} \qquad \mathrm{CD}_t = \bar{c}_t \cdot \mathrm{L}(\Delta T(t)).$$

A major concern with this damage function is that it effectively rules out catastrophic risks, even at high temperature changes. Take an 8°C temperature change, well outside the range typically assumed to be 'safe'. If per capita consumption is assumed to grow in real terms by 2% annually, then such damage applied to consumption 50 years hence would reduce the average consumption from 2.7 times today's value to 2.2 times, a significant reduction, but hardly a catastrophe of significant concern today. Even the 1% point in the outcome distribution conditional on an 8°C average temperature change is assumed here to be a reduction in consumption of only 32% which implies the representative agent is still 1.8 times wealthier than today. We hence augment Pindyck's (2012) damage function with the possibility of catastrophic events after reaching a particular temperature threshold, which itself creates at least the potential for a much larger impact on consumption, once thus calibrated.

While the possibility of climate tipping elements is receiving considerable attention in the scientific community, there is no single right specification (Kopp et al., 2016). There is, however, seeming convergence around global average warming of 6°C representing something akin to an upper bound for what could conceivably be quantified. The figure is a recurring theme in the literature, from the EU's High-End cLimate Impacts and eXtremes (HELIX) research project, which ends at 6°C, to Mark Lynas's popular book *Six Degrees*, which does the same (Lynas, 2008).[23] Its sixth chapter, "Six Degrees," begins with a vivid reference to Dante Alighieri's Sixth Circle of Hell. We take 6°C as our base level calibration for a parameter we label "$peakT$," above which we can expect to have hit a climatic 'tipping point' of sorts.

Specifically, we use Prob(TP) to denote the probability of hitting a 'Tipping Point' over a given interval of length "$period$" as a function of the global temperature change as of that time ($\Delta T(t)$), and of the parameter $peakT$:

$$(28) \qquad \text{Prob(TP)} = 1 - \left(1 - \left[\frac{\Delta T(t)}{\max[\Delta T(t), peakT]}\right]^2\right)^{\frac{period}{30}}.$$

Figure IV plots Prob(TP) as a function of $\Delta T(t)$ for a 30-year $period$ and a set of values of $peakT$. As $peakT$ increases, the probability of reaching a climatic tipping point decreases for a given $\Delta T(t)$.

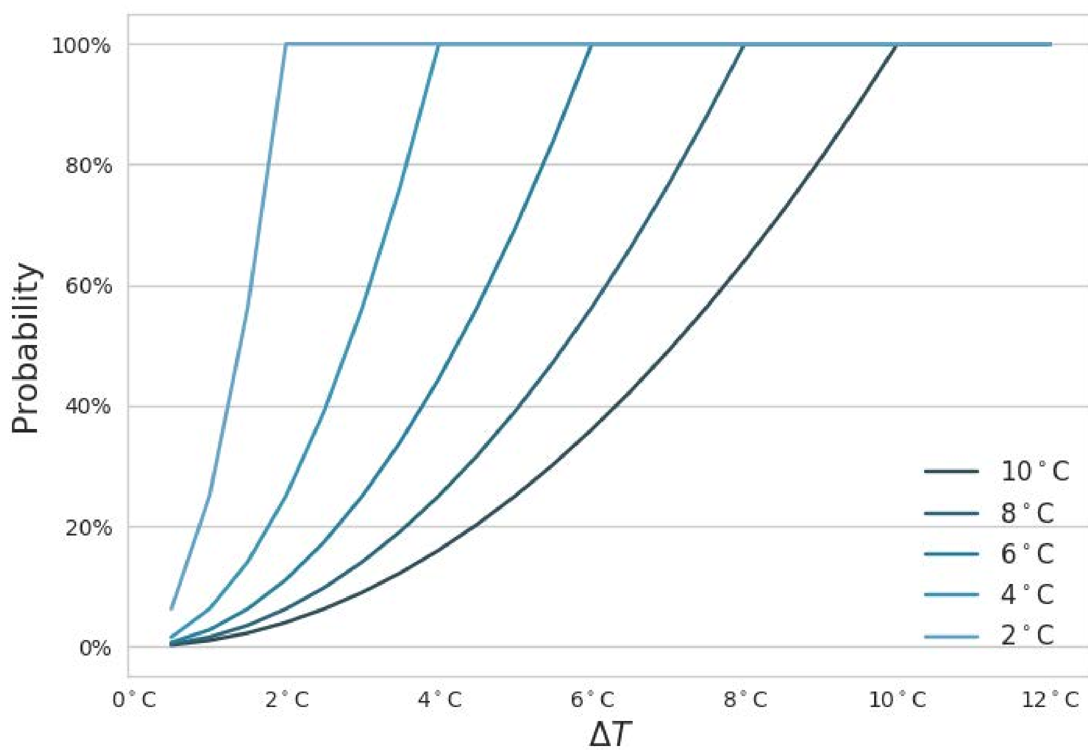

Figure IV—Probability of reaching a climatic tipping point as a function of *peakT*

[23] See Wagner and Weitzman (2015) for more context on the 6°C threshold. See helixclimate.eu for more on the EU's HELIX project.

In each period and for each state, there is a probability Prob(TP) that a tipping point will be hit, given $\Delta T(t)$ and $peakT$. Conditional on hitting a tipping point at time t^*, the level of consumption for each period $t \geq t^*$ is then at a level of:

$$\text{(29)} \qquad \mathrm{CDTP}_t = \mathrm{CD}_t \cdot \mathrm{e}^{-TP\_damage} = \bar{c}_t \cdot \mathrm{L}(\Delta T(t)) \cdot \mathrm{e}^{-TP\_damage} \text{ for } t \geq t^*,$$

where TP_damage is a random variable drawn from a gamma distribution with parameters $\alpha = 1$ and $\beta = disaster\_tail$. Figure V shows the cumulative distribution for tipping point damage (i.e., $(1 - \mathrm{e}^{-TP\_damage})$) for values of $disaster\_tail$ ranging from 6 to 30. Our admittedly *ad hoc* base case calibration uses $disaster\_tail = 18$.

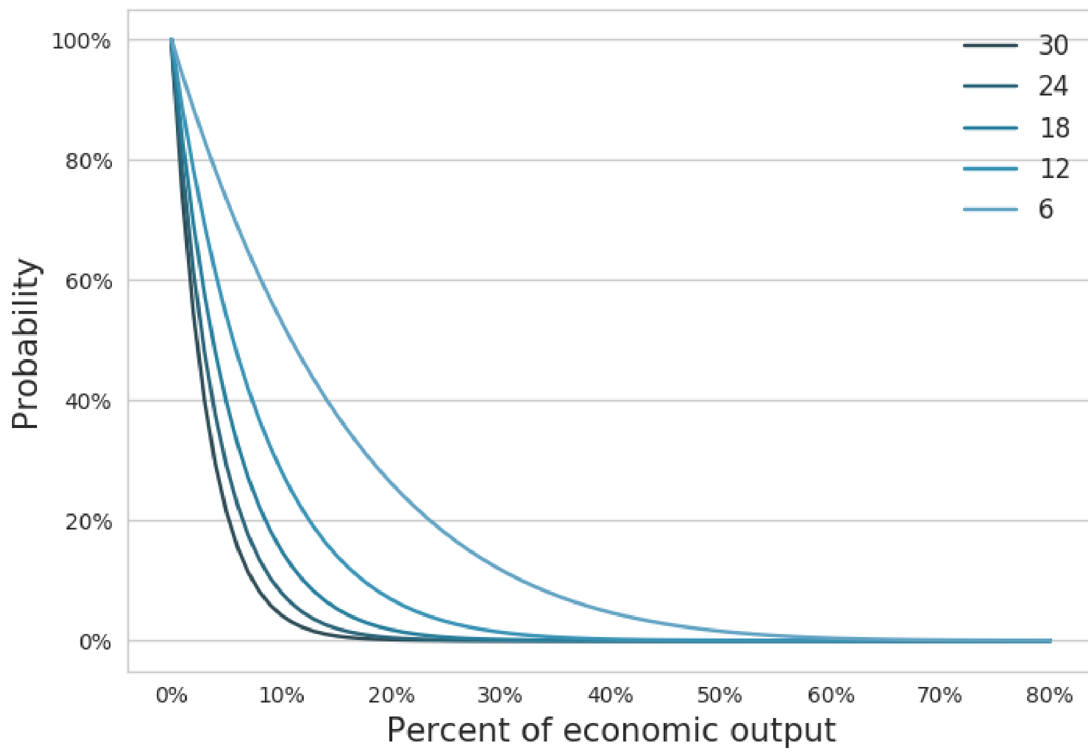

Figure V—Probability of damages greater than a particular percentage of output, given different levels of *disaster_tail*

2.3.3. Damage Function Uncertainty

The mapping from mitigation policy to damages over time, D_t, goes via cumulative radiative forcing, which determines the excess energy created by GHGs in the atmosphere. The damage distribution associated with a given level of radiative forcing is interpolated, or extrapolated, relative to the radiative forcing of damage distributions estimated from the three baseline scenarios. The first is based on the IEA's (2013) reference New Policies Scenario and leads to eventual atmospheric CO_2 levels of around 1,000 ppm. The second assumes constant mitigation leading to eventual levels of 650 ppm, equivalent to reducing emissions by almost 60% relative to the 1000 ppm scenario.

The third scenario assumes a constant mitigation of over 90%, leading to eventual CO_2 concentrations of 450 ppm.

For each of the three maximum GHG concentration levels—450, 650, and 1,000 ppm—we run a set of 6,000,000 random scenarios to generate a distribution of D_t for each period. We order the scenarios based on D_T, the damage to consumption in the final period. We then choose states of nature with specified probabilities to represent different percentiles of this distribution. For example, if the first state of nature is the worst 1% of outcomes, then we assume the damage coefficient at time t for the given level of mitigation is the average damage at time t for the worst 1% of values for D_t.

More generally, if the k^{th} state of nature represents the simulation outcomes in the range $[prob(k-1), prob(k)]$, then the damage coefficient for the k^{th} state of nature is the average damage in that range of scenarios in which the distribution for D_t lies within those percentiles.

The simulations are used to calculate damages in each period for any particular state of nature, θ_t, and any chosen time path for mitigation actions. We do this by first calculating the radiative forcing associated with each simulation at the end of each period, and then interpolating the damage smoothly between the three different simulations with respect to their levels of radiative forcing. Functional forms for both GHG levels and climate forcing as a function of GHG emissions are fitted to the Representative Concentration Pathway (RCP) scenarios adopted by the IPCC for its Fifth Assessment Report (IPCC, 2013). In the IPCC report emissions, GHG concentrations, and radiative forcing are given for each of three RCP scenarios. The radiative forcing is assumed to be given by a log-function fitted to these RCP scenarios.[24] The carbon absorption itself is similarly fit to the RCP scenarios, and is assumed to be proportional to the difference between the GHG level in the atmosphere and the cumulative carbon absorption up to that point in time, raised to a power.[25]

Our task now is to calculate an interpolated damage function using our three simulations where we have damage coefficients (for a given state and period) to find a smooth function that gives damages for any particular level of radiative forcing up to each point in time. To do so, we assume a linear interpolation of damages between the 650 and 1,000 ppm scenarios, and a quadratic interpolation between 450 and 650 ppm. In addition, we impose a smooth pasting condition at 650 ppm, having the level and derivative of the interpolation below 650 ppm match the level and slope of the line above.

Below 450 ppm, we assume climate damages exponentially decay toward zero. Mathematically, we let $S = d \cdot p \;/\; (l \cdot \ln(0.5))$, where d is the derivative of the quadratic damage interpolation function at 450 ppm, $p = 0.91667$ is the average mitigation in the

---

[24] Radiative forcing in a ten-year interval is given by: $5.351 \cdot [log(\mathrm{GHG}) - log(278.063)]$, where GHG is the average level of atmospheric CO_2. We estimated the constants from the three IPCC RCP scenarios.

[25] The carbon absorption in a ten-year interval is given by: $0.94835 \cdot |\mathrm{GHG} - (285.6268 + 0.88414 \cdot \sum absorption)|^{0.741547}$, where the sum is over absorption in previous periods. We again estimated the constants from the three IPCC RCP scenarios.

450 ppm simulation, and the level of damages is l. Radiative forcing at any point below 450 ppm then is x percent below that of the 450 ppm simulation, with $x = \frac{R-r}{R}$, where R is the radiative forcing in the 450 ppm simulation and r is the radiative forcing given the mitigation policy. Letting $\sigma = 60$, the extension of the damage function for $x > 0$ is defined as: $Damage(x) = l \cdot .5^{(x \cdot S)} e^{-[(x \cdot p)^2/\sigma]}$, which has the desired properties.

Figure VI shows the simulated distribution of the resulting damage functions in our base case, using $peakT = 6$ and $disaster\_tail = 18$, and assuming constant mitigation.

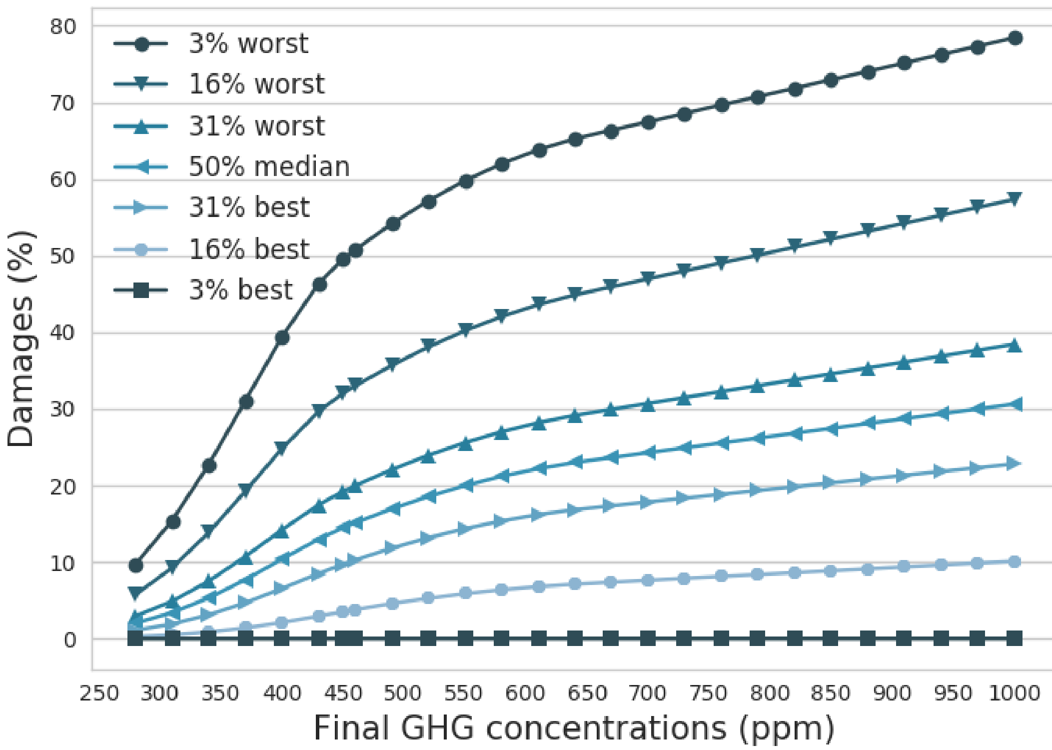

Figure VI—Interpolated final period damage functions

The climate sensitivity—summarized by state of nature θ_t—is not known prior to the final period ($t = T$). Rather, what the representative agent knows is the distribution of possible final states, θ_T. We specify that the damage in period t, given a cumulative radiative forcing, CRF_t, up to time t, is the probability weighted average of the interpolated damage function over all final states of nature reachable from that node. Specifically, the damage function at time t, for the node indexed by θ_t is assumed to be:

(30) $$D_t(CRF_{t,}\theta_t) = \sum_{\theta_T} \Pr(\theta_T | \theta_t) \cdot D_t(CRF_{t,}\theta_T),$$

where the sum is taken over all states that are possible from the node indexed by θ_t (i.e., for which $\Pr(\theta_T | \theta_t) > 0$).

2.3.4. *Damages for Concentrations Below Pre-Industrial Levels*

Introducing carbon dioxide removal (CDR) backstop technologies, combined with stochastic fragility θ_t creates a unique possibility: that, in some states of the world, GHG concentrations may fall below pre-industrial levels of 280 ppm. There is nothing magical about 280 ppm—in an absolute sense, it may not be the 'optimal' climate to begin with—but it does serve as the baseline for damage calculations based on global warming above pre-industrial levels. It is clear that going (well) below 280 ppm would lead to climate damages, much like going (well) above 280 ppm does. We introduce a simple penalty function of the form:

$$\text{(31)} \qquad \mathrm{f}(x) = \left[1 + e^{k\,(x-m)}\right]^{-1},$$

where m is the level of GHG concentrations where calibrated at half the total penalty, and k is a simple scalar. For our base case calibration, we use $m = 200$ and $k = 0.05$. The benefit of a low k and, thus, a smooth penalty function is largely computational. More importantly, the calibration ensures that the penalty (31) at 280 ppm is close to zero. In our optimization, we also restrict climate damages and mitigation, ${x_t}^*$, to be nonnegative.

2.4. Tree structure

Figure VII illustrates the tree structure employed in EZ-Climate's baseline analysis. Beginning with the first node, in 2015, the agent is assumed to know the structure of the decision tree, the state probabilities, and the damage function in each future state of the world. Period zero runs from 2015 through 2030. In 2030, the agent learns whether the world is in state up ('u') or state down ('d'). There is a 50% probability of each of the two states. Similarly, at the end of period one (in year 2060) she learns whether the world is in state 'uu', 'ud', or 'dd', etc. Notice that at the end of period four, all uncertainty is resolved, in that the agent will learn which of the six final states the world she is in, and what the true damage function is. Following this point, in period five, she has one final period in which she can make a mitigation decision. In period six, which in our base case runs from 2400 on to infinity, the agent can no longer mitigate. Consumption continues to grow deterministically from this point forward at a rate r. Consumption, thus, is given by $c_t = c_T(1+r)^{t-T}$ after $t = 2400$, and period six utility is given by equation (10).

In the baseline model, where a move up or down in each period is equally likely, the probabilities of the final states are given by a binomial distribution, the simplest possible probability representation.

Another feature evident in Figure VII is particularly important, given our use of the Epstein-Zin preference specification: the recombining tree structure. This implies two features: For one, the damage function in each state (after period zero) is independent of the way in which information was revealed at the end of each period. For example, the damage function in state 'uud' (the blue path in Figure VII) is identical to that in 'udu' (green) and 'duu' (red).

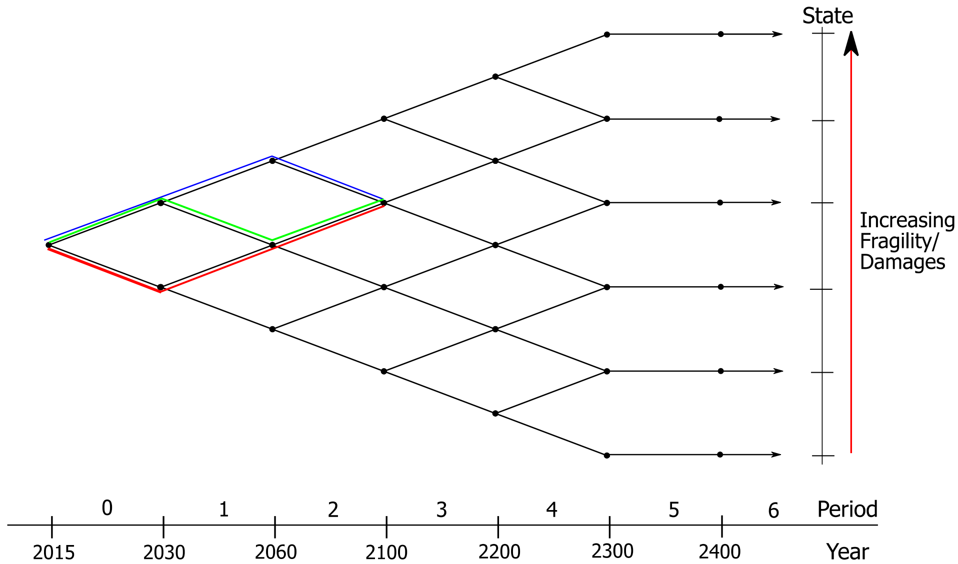

Figure VII—Diagram of tree structure used in solving the model for each state of nature across time

Second, the agent's utility is path-dependent. The history of mitigation depends on the process by which the agent learns the state. Thus, consumption, and mitigation, will depend upon the path. Consequently, in solving for the agent's utility along each of these paths, we need to keep track of the path by which the agent learned about the damage function. Consumption decisions depend on it.

For example, following equation (2), the consumption flow at the start of period 1 is given by $c_1 = \bar{c}_0 \cdot e^{0.02\times 15}\big(1 - D_1(CRF_1, \theta_1) - \kappa_1(x_1)\big)$. That is, the consumption at the start of period one (in 2030), c_1, is equal to endowed consumption (2015 consumption plus 1.5% growth for 15 years), minus the fractional cost of damages and of mitigation chosen at the beginning of period one. Mitigation is optimally chosen by the agent, and is therefore a function of the state—mitigation will be lower if the agent learns that the world is in state 'd' rather than state 'u'.

This analysis gives us consumption levels c_0 and c_1, for the two states in 2015 and in 2030, respectively. To interpolate between c_0 and c_1, we fit an exponential growth function to consumption levels, using 5-year intervals. Note that this is equivalent to assuming that immediately after choosing the mitigation level in period zero, the agent's consumption starts to reflect climate damages from the first revealed state ('u' or 'd'). However, she is not allowed to change the period zero mitigation to reflect this knowledge until the next period.

This interpolation ensures that the agent's consumption path is relatively smooth. It also introduces approximation errors. However, adding more periods at which the agent can choose a new level of mitigation would result in far higher computational costs. With T periods, we have a "$2^{T+1} - 1$"-dimensional optimization problem: Each model run requires choosing $2^{T+1} - 1$ optimal mitigation levels. In the 'easy' spirit of EZ-Climate, simplifying makes the solution both tractable and doable in the first place. While similar attempts at integrating Epstein-Zin preferences into climate-economy models require supercomputers (see, e.g., Cai et al., 2016, 2015, 2013), EZ-Climate is solvable on a standard personal computer within minutes.

3. Results

EZ-Climate's main output is the optimal price of one ton of CO_2 today, in year 2015, and at the beginning of each of the subsequent five periods—years 2030, 2060, 2100, 2200, and 2300. These are the times in the model when mitigation decisions are made. Figure VIII shows the results for the CRRA model run and our base-case model calibration, using $peakT = 6$ and $disaster\_tail = 18$, as well as an EIS of 0.9 and $RA = 7$, calibrated to observed financial asset prices.

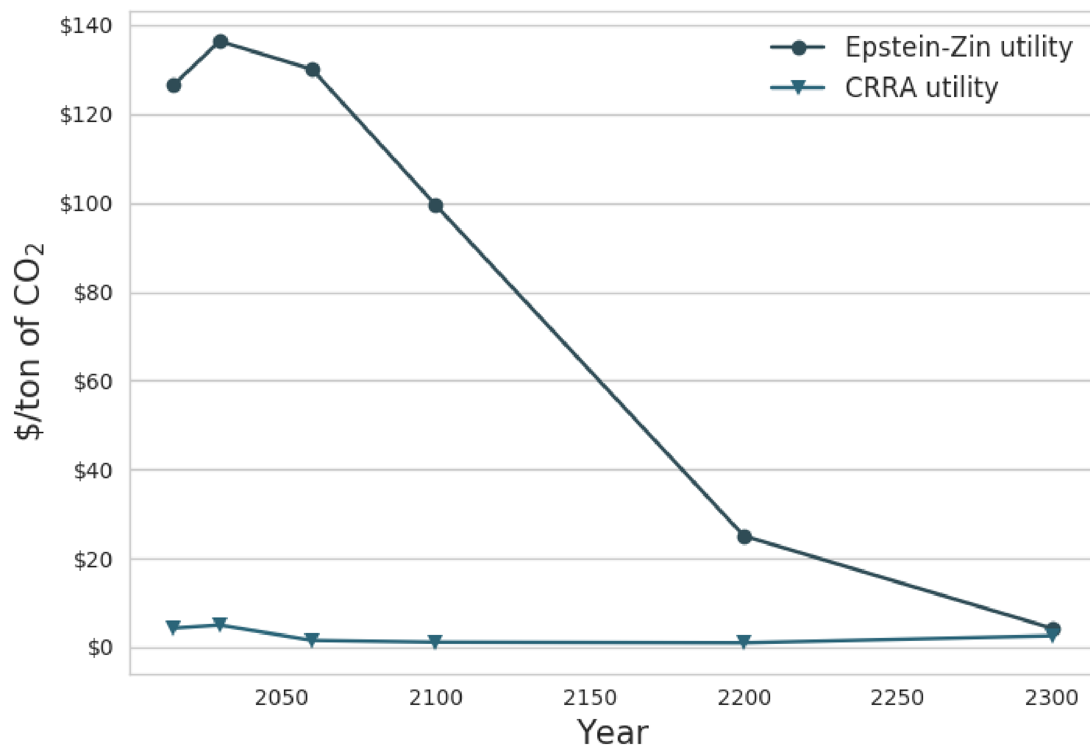

Figure VIII—Optimal price per ton of CO_2 under constant relative risk aversion (CRRA) and our base case, employing $RA = 7$ in both cases, and $peakT = 6$, $disaster\_tail = 18$, and $EIS = 0.9$ in the base case

The difference in optimal price paths is striking: Under the CRRA case, the expected optimal price is consistently low. Under the Epstein-Zin base case, the price first increases slightly and then decreases to a value almost as low as in the CRRA case. Significantly, an RA = 7, calibrated to observed equity risk premia, decreases the CO_2 price under CRRA assumptions to below $1 in 2015, while the price in EZ-Climate's base case is well over $100/ton CO_2.

It is difficult to overemphasize the probabilistic nature of this model. The precise 2015 price represented in Figure VIII for Epstein-Zin utility is \$124.94, which we round to \$125 throughout our discussion. Even when leaving the damage simulation, explained in section 2.3, constant optimal prices in a random sampling of 30 optimization runs range from \$124.89 to a high of \$126.50. The \$124.94, thus, is toward the lower, conservative end of this particular sample. Anything more precise than saying "around \$125" would amount to false precision.

Note that the 2015 price comes from a single node in the tree. In each subsequent year, that price is set in expectation over all possible states of nature in that given year. The left panel of Figure IX shows optimal base-case CO_2 prices across time and states for one model run. All grouped nodes at a given time have the same degree of fragility and, thus, the same damage for a given amount of atmospheric greenhouse-gas concentrations. The lines connecting the boxes indicate the paths that information about the Earth's fragility, θ, has taken, a feature explored in Section 2.4 above.

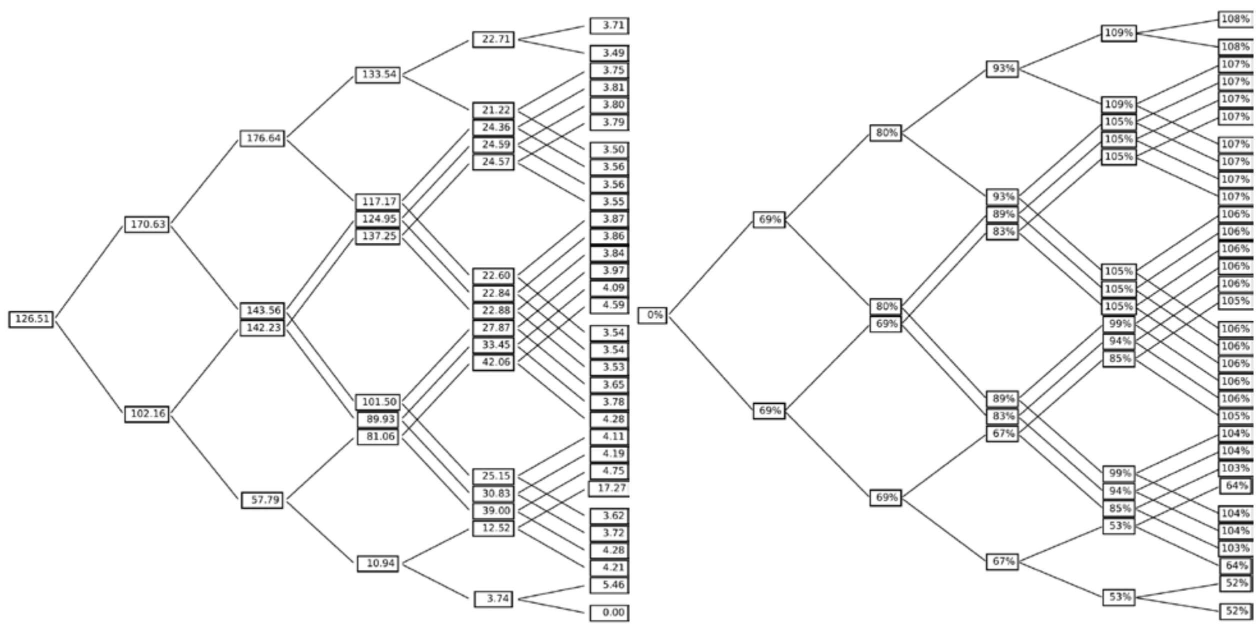

Figure IX—Optimal price per ton of CO_2 across time and states (left) and average mitigation up to a particular time and state (right) in the base-case calibration ($peakT = 6$, $disaster\_tail = 18$, and $EIS = 0.9$)

The right panel of Figure IX shows the fractional, average mitigation up to each time and state. It reveals the reason for sometimes wildly different prices at the same node. In general, the greater is θ revealed to be, the higher is the optimal average mitigation effort, ranging from slightly over 50% to over 100% in the final period.

Optimal average mitigation levels, in turn, are closely linked to climate damages at each node, which similarly depend on the path chosen (Figure X, left panel). Both Figure IX and Figure X also show the large costs associated with negative θ draws in latter periods. Repeated good ('*d*' for down) draws in early periods, followed by a bad ('*u*' for up) draw in the final period results in half the amount of average mitigation up to that point (52%)

compared with reaching the same node via an early bad ('*u*') draw followed by good draws (104%). Associated climate damages range from 5.5% ('*ddddu*') to 1.4% ('*udddd*') , respectively (Figure X, left panel).

Lastly, the right panel in Figure X plots GHG concentrations along the optimal base-line pathway. The relatively small changes along most paths across time reveal the inherent inertia in the climate system. The large differences across nodes in the final period, from close to pre-industrial levels of 280 ppm to well over 800 ppm, reveal the enormous costs of bad θ draws. Looking at GHG levels also confirms our prior conclusions around the importance of path dependency: the '*ddddu*' path results in ultimate GHG levels above 800 ppm, whereas the '*udddd*' path results in levels close to 300 ppm. Bad news is costly. Bad news received late is extremely costly.

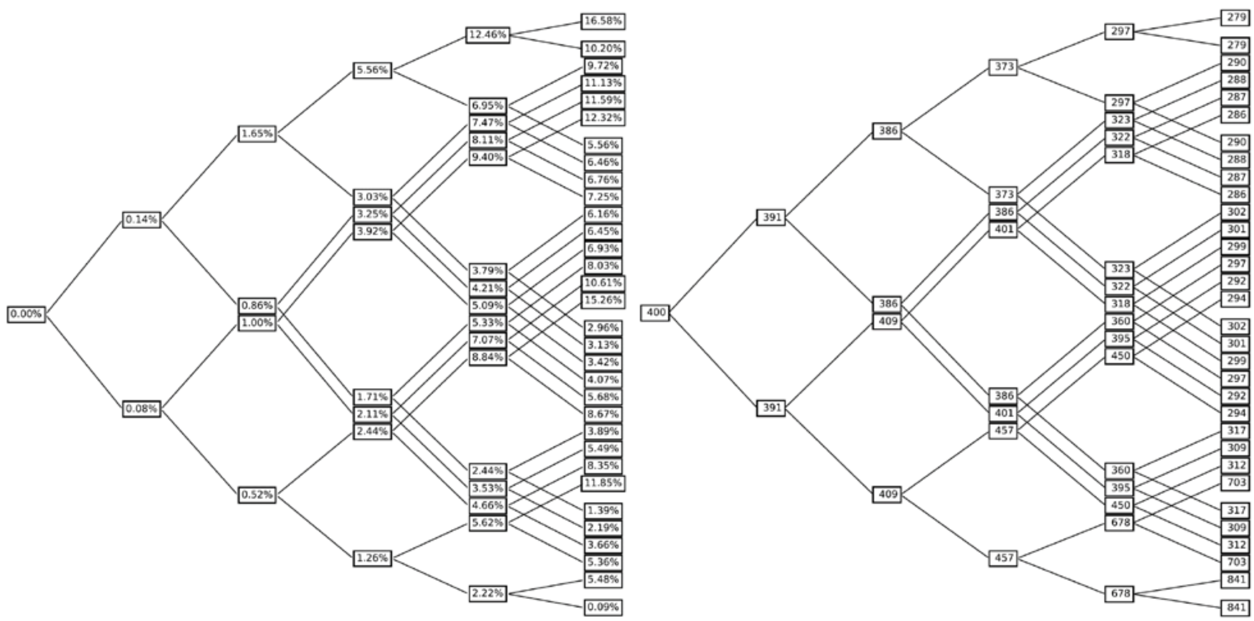

Figure X—Climate damages (left) and GHG concentrations (in parts per million, ppm, right) along the optimal CO_2 price path across time and states in the base-case calibration

Following GHG levels closely across time, especially in early periods, also demonstrates the positive environmental impact of moving onto an optimal CO_2 price path. Despite inertia, from 2015 to 2030 alone, optimal GHG concentrations decline from 400 to around 390 ppm.[26]

Finally, we present both consumption of our representative agent in each node of the tree along the optimal path in the base case, and costs of mitigation as a percentage of economic output (Figure XI). The latter, in part, corroborates earlier observations around how costly bad news is—however, not necessarily the fact that bad news received late is

[26] See Section III.E on the large social costs of delayed implementation of the optimal CO_2 price path. See, e.g., Le Quéré et al. (2016) for rough corroborative evidence for such a relatively rapid decline in concentrations under aggressive mitigation scenarios, relying on the net flow of atmospheric carbon into land and ocean sinks.

extremely costly. A '*ddduu*' path, for example, results in costs larger than those following any other path of getting to the same final node.

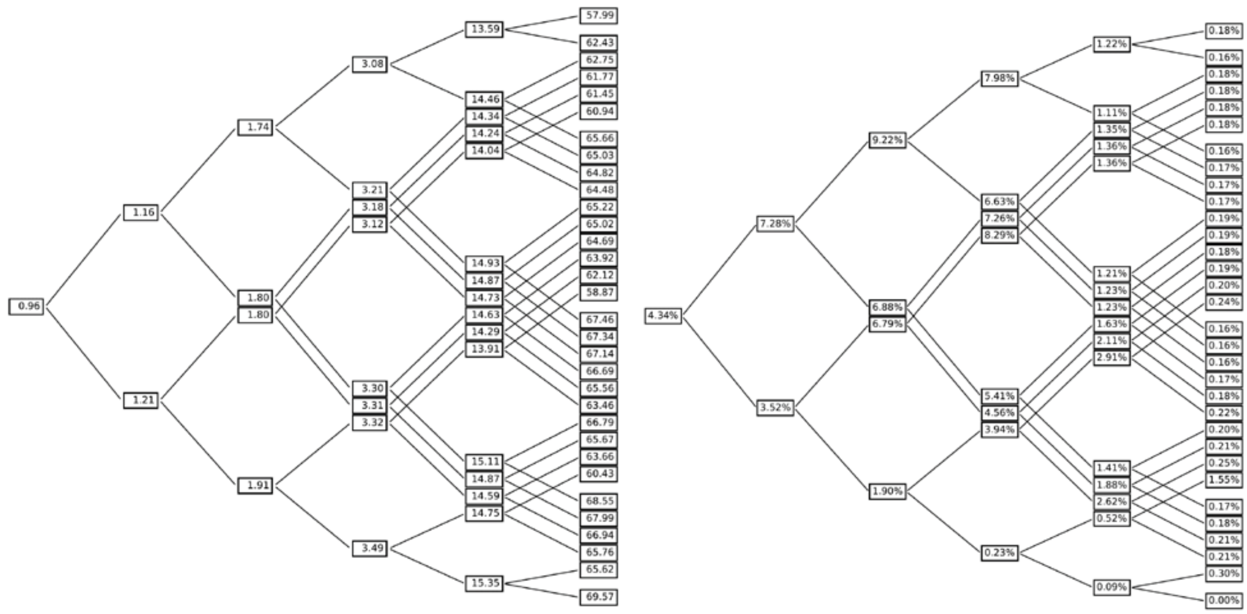

Figure XI—Consumption (left) and cost of emissions reductions (right) along the optimal CO_2 price path across time and states in the base-case calibration

EZ-Climate's optimal carbon price depends on a number of inputs. Figure I reveals the importance of *RA* for calibrating economic variables to capture observed equity risk premia, and its influence on the optimal CO_2 price.

Given the apparent importance of the use of Epstein-Zin preferences to capture climate risk, it might then be surprising to see how little of the overall optimal CO_2 price is explained by risk aversion as opposed to expected climate damages (Figure XII). There, too, the importance of moving to Epstein-Zin preferences in the first place is apparent, but once done, most of the impact comes from the expected damage component of the damage distribution being discounted at lower rates, rather than the higher curvature of the utility function across states of nature.

Crost and Traeger (2014) and Cai *et al.* (2016) support this conclusion, though their explanations differ. Crost and Traeger (2014) suggests it is because of a failure to account for disasters. Cai *et al.* (2016) attempt to model such climate disasters and still find a similarly small role for risk. One explanation might be that even the type of disasters modeled by Cai *et al.* (2016) and represented here in EZ-Climate do not yet capture true uncertainty (Brock and Hansen, 2017; Wagner and Zeckhauser, 2017). For example, although we include tipping points in the simulation of our damage function, these events are averaged with others to create an average loss in a given state. Moreover, although the state of nature is not known in advance, the probability of each state and the average loss in any given state and for any given degree of mitigation is known in advance. That information is used optimally to mitigate against those potential damages. Our tipping points do not in any sense catch our agent by surprise. All this ascribes yet more importance to the calibration of the damage function itself, and to considerations of model mis-specification (Brock and Hansen, 2017).

3.1. Risk decomposition

Figure I presents the optimal CO_2 price as a function of the assumed equity risk premium for both Epstein-Zin and CRRA utility and points to the importance of using the former. We further decompose the optimal CO_2 price into a risk aversion and an expected damages component.

Let $D_{s,t}$ denote the marginal damage, that is the loss of consumption in state s in future period t that results from putting one more ton of carbon into the atmosphere today (at time 0). The optimal CO_2 price then is given by:

$$(32) \qquad \sum_{t=1}^{T} \sum_{s=1}^{S(t)} \pi_{s,t} m_{s,t} D_{s,t} \left(= \sum_{t=1}^{T} E_0\left[\widetilde{m}_t \widetilde{D}_t\right]\right),$$

where $m_{s,t}$ is the pricing kernel in state s at time t, which is the marginal value today of one additional unit of consumption in state s at time t,[27] $\pi_{s,t}$ denotes the probability of state s at time t, and $S(t)$ denotes the number of states at time t. That is, to calculate the cost to the representative agent of an additional ton of carbon emissions, we sum over all consumption damages, in every state of nature at every future time, multiplied by the value of an additional unit of consumption in that state at that time. Equation (32) can be decomposed to equal:

$$(33) \qquad \underbrace{\sum_{t=1}^{T} E_0[\widetilde{m}_t] \cdot E_0\left[\widetilde{D}_t\right]}_{ED} + \underbrace{\sum_{t=1}^{T} cov_0\left(\widetilde{m}_t, \widetilde{D}_t\right)}_{RP}.$$

Note that $E_o[\widetilde{m}_t] = \frac{1}{R^f(0,t)}$, where $R^f(0,t)$ is the payoff, at time t, to a \$1 investment in a risk-free bond at time 0. Alternatively, $E_0[\widetilde{m}_t]$ is the risk-free discount factor between today and t. We can, thus, rewrite the first component of (33) as the sum of the marginal damages, discounted back to the present at the risk-free rate and label it expected damages, $ED = \sum_{t=1}^{T} \frac{E_0[\widetilde{D}_t]}{R^f(0,t)}$. The second component is the risk premium (RP) over the expected damages that society is willing to pay, defined as the covariance of the marginal damages with marginal utility.

Rewriting (33) then gives the risk premium as the difference between the optimal CO_2 price and the expected-damages:

$$(34) \qquad RP = \text{optimal CO}_2 \text{ price} - ED,$$

[27] Equivalently, $m_{s,t}$ is defined as the ratio of marginal utility with respect to current consumption in that state to the marginal utility today, that is $m_{s,t} = \left(\frac{\partial U}{\partial c_{s,t}}\right) / \left(\frac{\partial U}{\partial c_0}\right)$, where $c_{s,t}$ denotes the agent's consumption in state s at time t.

both of which are readily calculated by EZ-Climate. Figure XII shows the result for our base-case calibration ($peakT = 6$ and $disaster\_tail = 18$).

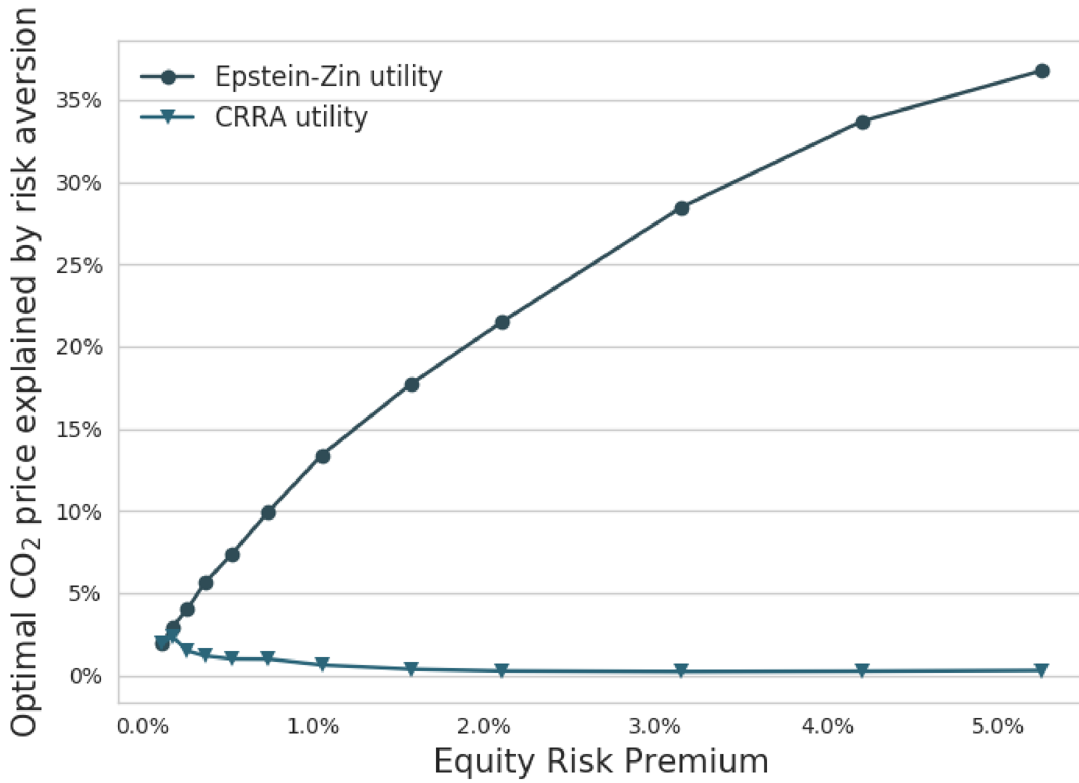

Figure XII—Percentage of the optimal 2015 CO_2 price explained by risk aversion, as opposed to expected damages, under Epstein-Zin and CRRA preferences, translated into the implied equity risk premium using Weil (1989)'s conversion, while holding implied market interest rates stable at 3.11%

Like the optimal CO_2 price in Figure I, this decomposition varies widely with the assumed equity risk premium—and it crucially depends on the distinction between using Epstein-Zin versus CRRA preferences. For low risk premia, less than 5% of the optimal CO_2 price is explained by risk aversion. For our base case calibration with an EIS of 0.9 and $RA = 7$, around 10% is explained by risk aversion.

3.2. Sensitivity to climate damage parameters

Two key damage-function parameters are *peakT* and *disaster_tail*, assumed to equal 6 and 18, respectively, in our base-case calibration. Figure XIII shows the sensitivity of the optimal CO_2 price to both parameters: the optimal CO_2 price increases with decreasing *peakT* and *disaster_tail*.

There is no single 'right' combination of *peakT* and *disaster_tail*, at least given our current state of knowledge of the underlying climate science. The base-case calibration in

EZ-Climate employs $peakT = 6$ and $disaster\_tail = 18$. Lower values of each increases the optimal CO_2 price along their respective dimension.[28]

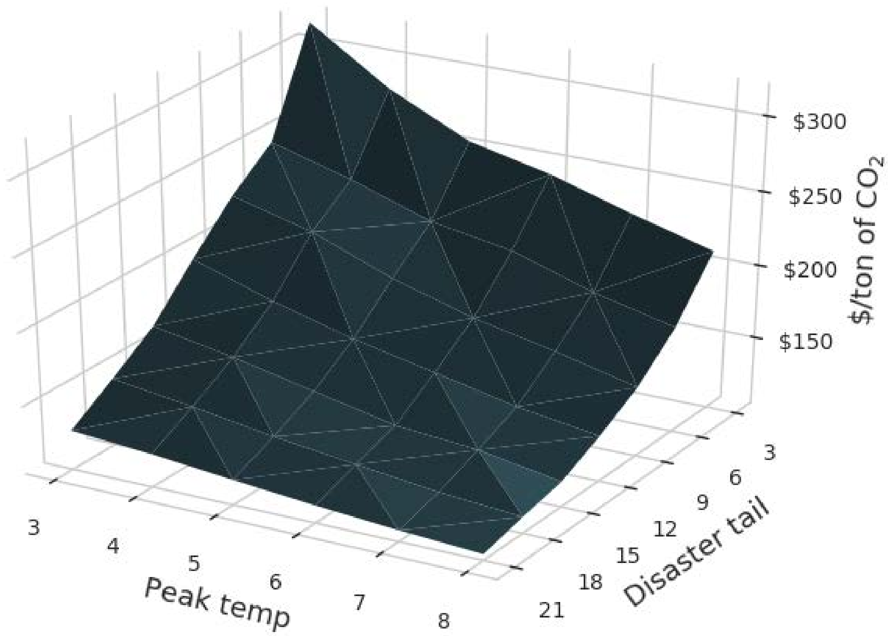

Figure XIII—Optimal 2015 CO_2 price increases with decreasing *peakT* and *disaster_tail*

3.3. Sensitivity to mitigation cost parameters

On the mitigation cost side, two important factors are the rate of technological change and the potential for an outright backstop technology. Climate-EZ includes parameters capturing both exogenous and endogenous technical change, represented by coefficients φ_0 and φ_1, respectively, in equation (24).

Figure XIV shows how the optimal CO_2 price in early periods first increases with φ_0 moving from 0 to 1.5% before decreasing again with increased exogenous technical change. The relatively low optimal price of CO_2 in early years with low φ_0 is explained by the fact that the larger future GHG concentrations will lead to large damages in bad states, decreasing the marginal benefit of mitigation today.

[28] See Section II.C.ii, in particular Figure IV and Figure V, on more on the damage-function calibration, incorporating *peakT* and *disaster_tail*.

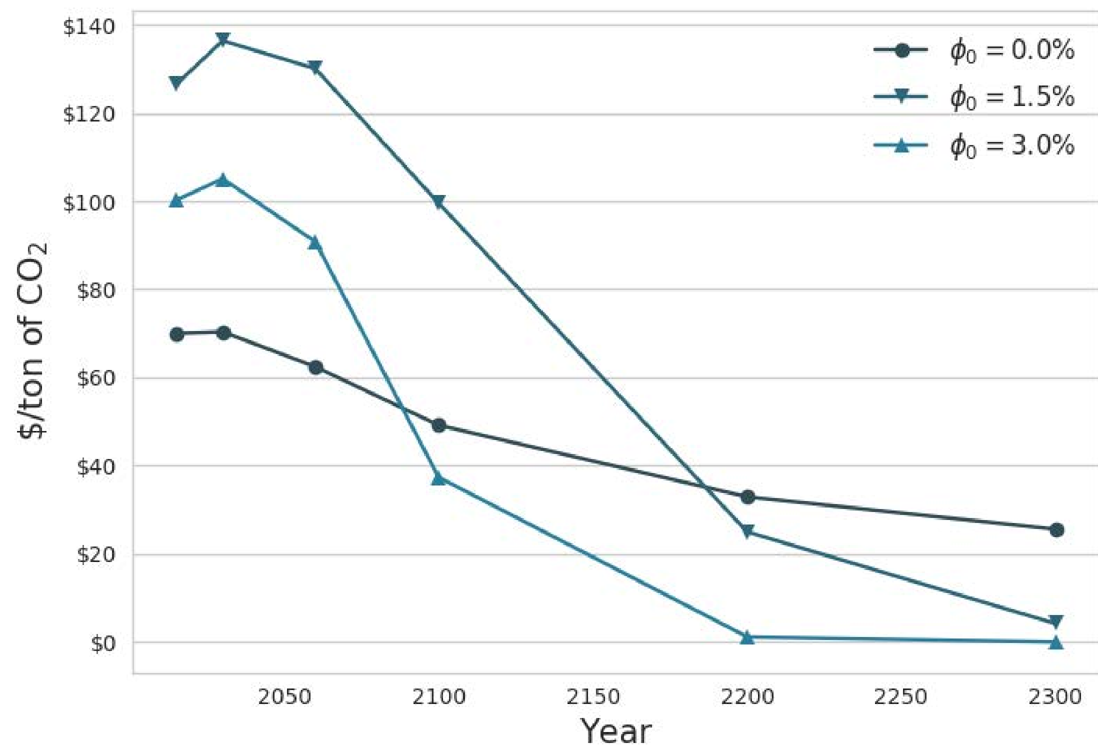

Figure XIV—Optimal CO_2 price in early years first increases then decreases with higher exogenous technical change, φ_0

Endogenous technical change is assumed to be proportional to average mitigation up to time t. Figure XV shows the sensitivity of the optimal CO_2 price to φ_1 ranging from 0 to 1.5%, which can best be seen as multiples of average mitigation affecting learning-by-doing.[29] The optimal CO_2 price increases at first slightly and then decreases with increased endogenous technical change in later years.

[29] In all cases, exogenous technical change, φ_0, is 1.5% per year, plus the multiple φ_1 of average mitigation to date (see Section II.B.ii). For example, with $\varphi_1 = 1$ and average mitigation for the first 15 years equal to 0.50, then the rate of technical change would be φ_0, or 1.5%, plus $\varphi_1 \cdot X_{15}$, or $1 \cdot 0.50\%$, summing to 2.0% per year. Costs after 15 years, $\kappa_{15}(x)$, will have declined to $(1-2\%)^{15} = 73.8\%$ of today's costs, given by equation (24). For simplicity, our base case assumes $\varphi_1 = 0\%$, in addition to $\varphi_0 = 1.5\%$.

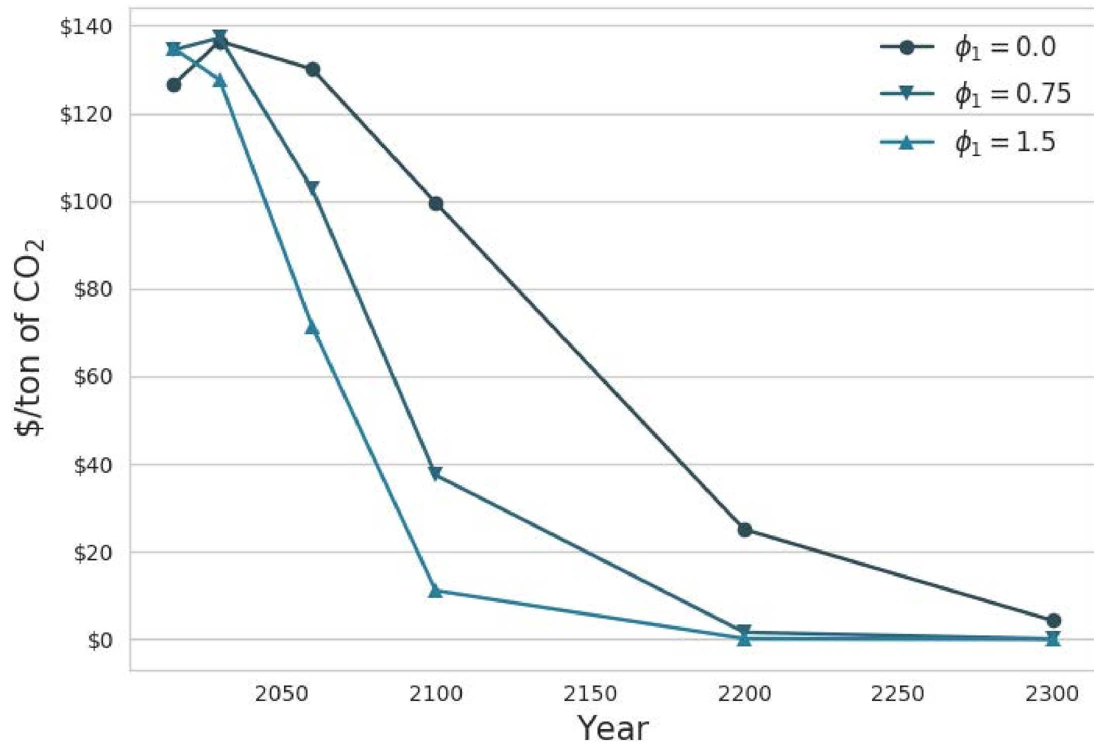

Figure XV—Optimal CO_2 price decreases with increased endogenous technical change in later years

Endogenous technical change, in turn, interacts with the availability of a backstop technology (Figure XVI). The EZ-Climate base case assumes essentially no backstop technology—i.e. a 'join' price of $\tau^* = \$2{,}000$ and a full-on backstop at $\tilde{\tau} = \$2{,}500$.[30] The "$350 backstop" path assumes $\tau^* = \$300$ and $\tilde{\tau} = \$350$, with no endogenous technical progress, $\varphi_1 = 0$. The third scenario assumes $\tau^* = \$300$, $\tilde{\tau} = \$350$, and $\varphi_1 = 0.67$. This $350 backstop assumption plus endogenous technical change at first leads to a slightly higher optimal CO_2 price, though over time endogenous technical change helps decrease the optimal CO_2 price as well.

[30] See Section II.B.i for an explanation of the integration of backstop technologies. Figure II shows this baseline backstop assumption, and the more aggressive scenario assuming $\tau^* = \$300$ and $\tilde{\tau} = \$350$.

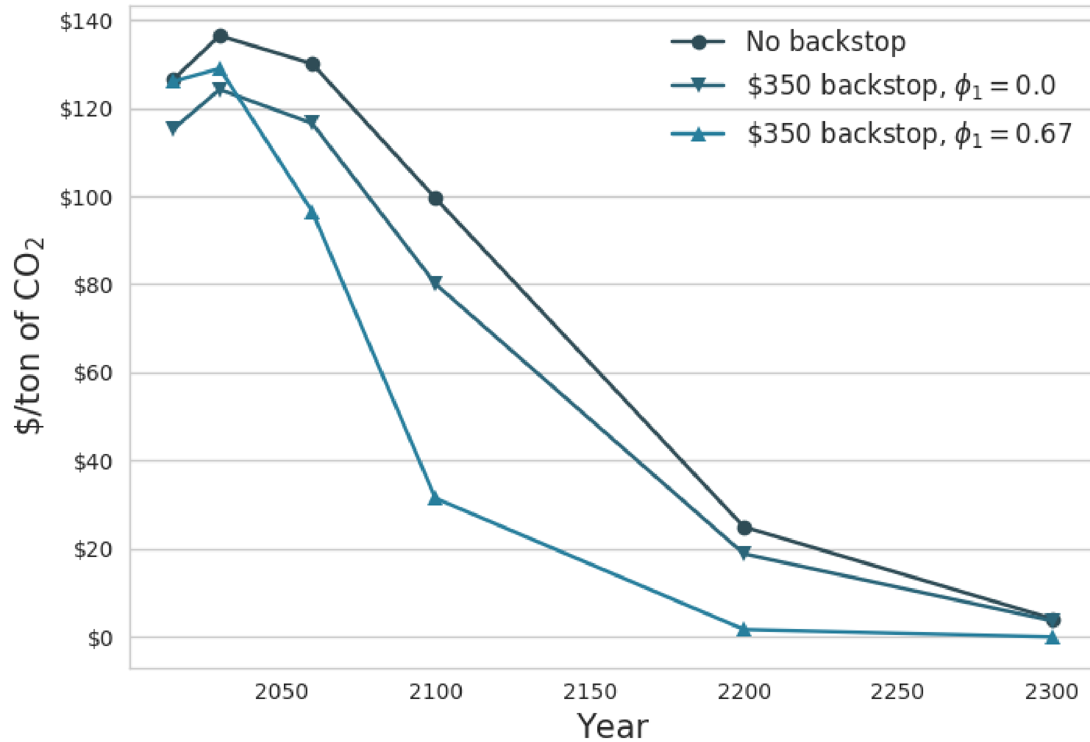

Figure XVI—Optimal CO_2 price decreases with backstop, with or without endogenous technological change

3.4. Sensitivity to economic and preference parameters

As important as climate risk and mitigation cost parameters are in EZ-Climate, a proper calibration of fundamental economic parameters has at least as much influence on the optimal CO_2 price. One key input is the assumed rate of economic growth.

A large fraction of climate damages is likely to occur in the distant future. Even in the worst states of nature modeled here, the economy in 2100 or 2300 will be a multiple of the size of today's. As a result of this large disparity in current and future consumption levels, future damages are discounted at a high rate[31], posing significant ethical challenges that can cut both ways.[32] On the one hand, a much richer society will be better equipped to deal with the challenges of climate change. On the other, a much richer society will be willing to pay comparatively more for a healthy planet—and, in the case of oft-irreversible damages brought about by climate change, would have wanted earlier generations to mitigate more (Sterner and Persson, 2008).

---

[31] In the EZ-Climate base case, with a growth rate of 1.5%/year, the risk-free rate is around 3.11%/year.

[32] The broader ethical and discounting debate has spawned hundreds of contributions, sometimes with well-founded opposing conclusions (e.g., Arrow et al., 2013; Broome, 2008; Dasgupta, 2008; Gollier, 2012; Gollier and Weitzman, 2010; Heal, 2017; Summers and Zeckhauser, 2008).

EZ-Climate, like most other climate-economy models such as, for example, DICE, takes economic growth as exogenous. Our base case assumption, before climate damages and mitigation costs, is an exogenous rate or growth of 1.5% per year.[33] This assumption is consistent with empirically observed growth rates over the past century. However, a number of scholars now argue that we should anticipate far lower growth going forward than we have observed in the past century (Gordon, 2016, 2012). At the same time, we cannot exclude the possibility of higher growth rates.

Figure XVII shows the effect of changing growth rates, while also adjusting the EIS, to hold real interest rates constant at the assumed rate of around 3.11%.[34] Optimal CO2 prices in early periods are highly sensitive to these simultaneous changes in economic growth rates and EIS.

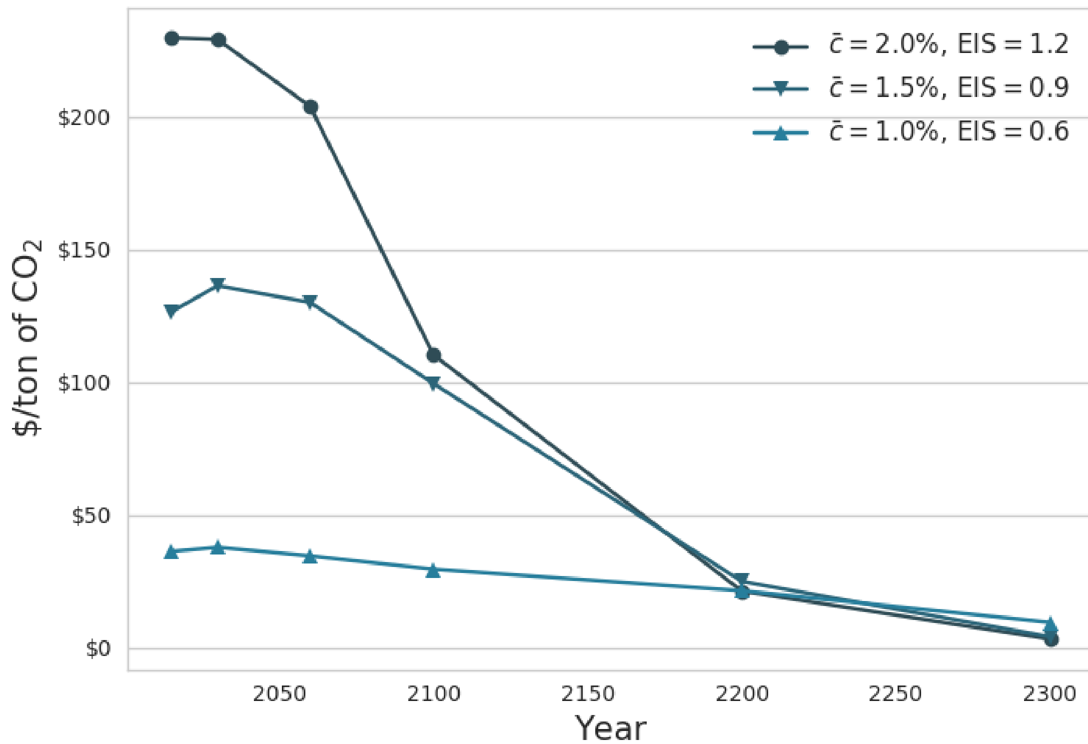

Figure XVII—Increasing (decreasing) economic growth rates, while changing EIS to keep real interest rates constant at 3.11% per year, increases (decreases) optimal CO_2 prices dramatically in early periods

Figure XVIII, by contrast, shows the impact of changing growth rates while holding the EIS constant. The impact on optimal CO_2 prices is comparatively small.

[33] See, for example, Bansal, Ochoa and Kiku (2016), who relax the assumption of an exogenous growth rate and investigate directly the impact of climate change on growth rates.
[34] See footnote 10 for a discussion of real interest rates in our base case.

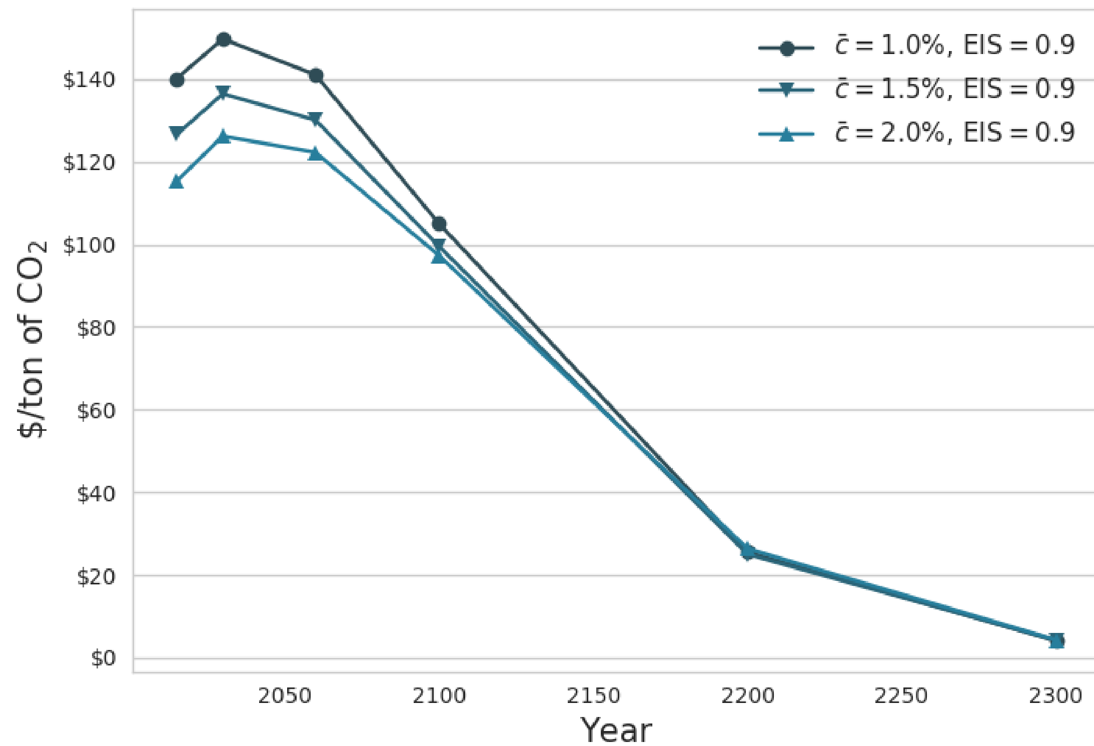

Figure XVIII—Changing economic growth rates, while keeping EIS constant at 0.9, has little impact on optimal CO_2 prices (note the y-axis scale compared with Figure XVII)

Given the considerably smaller discount rate, why is it that the SCC does not increase more in response to this change? The answer lies, in part, in our damage specification in equation (2), where damages, for a given level of fragility and cumulative mitigation, are proportional to the baseline level of consumption, given any specific economic growth rate r: $\overline{c_t} = \overline{c_0} \cdot (1+r)^t$. Thus, in an alternative specification with a 1% as opposed to 1.5% annual growth rate, and consumption in year 2300 that is lower by a factor of 4, we are also implicitly specifying that damages are lower by a factor of 4. Whether it is reasonable to assume that, if people are poorer in the future, they will be hurt proportionally less by climate change is subject to vigorous debate, and worthy of future research (e.g., Convery and Wagner, 2015; Heal and Park, 2016).

The comparison of Figure XVII and Figure XVIII also makes the importance of the EIS clear. Ours is calibrated to a risk-free bond yield rate of around 3.11%, implying an EIS of 0.9.[35] Figure XIX shows the optimal CO_2 price path for various EIS levels. Note that this figure does not hold real interest rates constant. That connection with bond yields implies the EIS is not a free parameter to be chosen but one that needs to be calibrated based on

[35] Note that calibrations of EIS have changed widely over time. Earlier estimates, most notably by Hall (1988), are much closer to zero. Bansal and Yaron (2004), in a model with Epstein-Zin preferences and consumption shocks, estimate an EIS of 1.5. Other estimates vary further (Havránek, 2015; Thimme, 2017), while Epstein *et al.* (2014) offers a potent critique by pointing to the sizeable magnitude of implied time premia by standard calibrations of Epstein-Zin utility specifications. Note that our calibrated EIS of 0.9 is below Bansal and Yaron's (2004), largely because their model assumes persistent shocks to consumption growth rates. The only shocks to consumption on Climate-EZ stem from climate risk.

assumed bond yields. Our base-case EIS of 0.9 corresponds with a bond yield of 3.11%. The higher EIS of 1.2 corresponds with a yield of 2.74%, and the lower EIS of 0.6 a yield of 3.77.

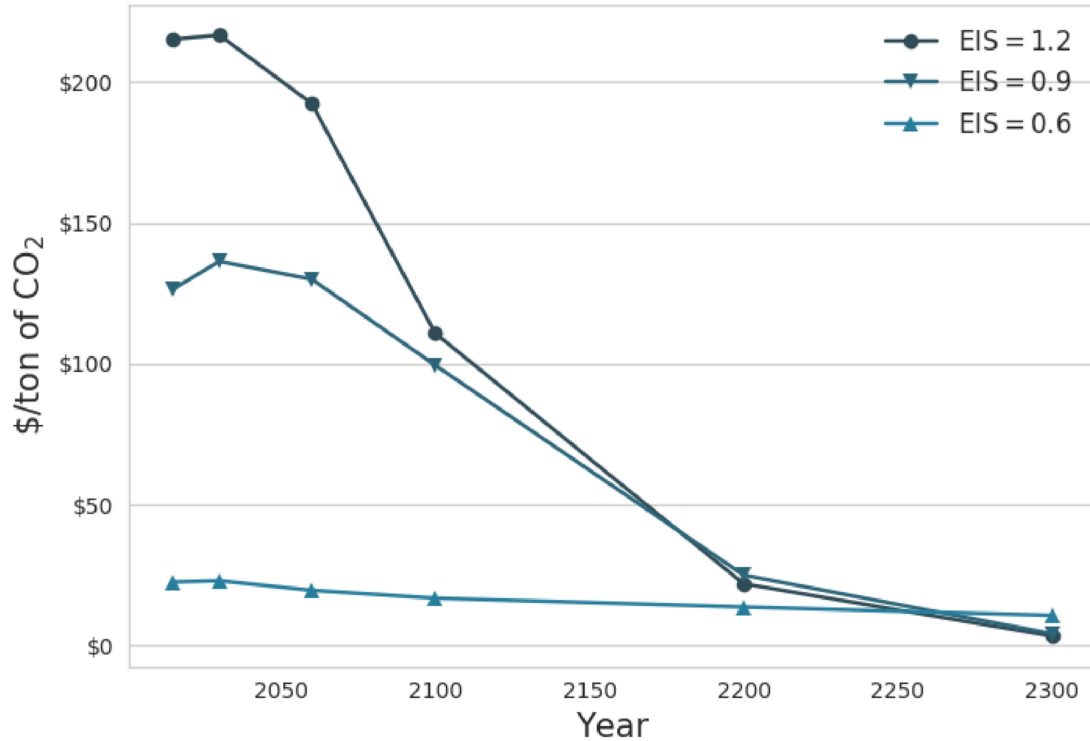

Figure XIX—A higher (lower) EIS goes hand-in-hand with a higher (lower) optimal CO_2 price in early years

Lastly, we explore the impact of the pure rate of time preference, $\delta = (1-\beta)/\beta$. First, we hold the EIS fixed at 0.9, while δ varies from 0.25% to 0.75% (Figure XX). Holding EIS fixed, a higher δ implies a higher real interest rate and, thus, lower optimal CO_2 prices, and vice versa. Corresponding bond yields vary from 2.89% to 3.32% for δ from 0.25% to 0.75%, respectively. Note the large changes in early optimal CO_2 prices, despite holding EIS constant at 0.9.

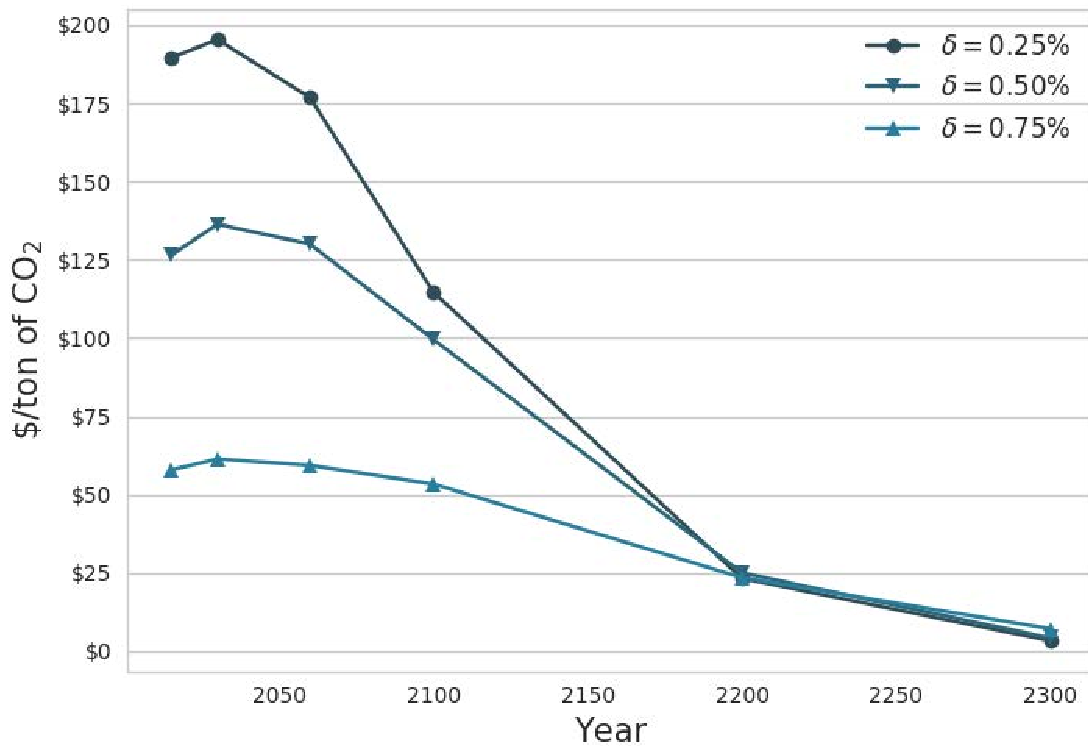

Figure XX—Optimal CO_2 prices increase (in early years) with decreasing pure rate of time preference, δ, holding EIS fixed at 0.90

Second, we vary EIS in order to hold real interest rates fixed at around 3.11% while adjusting δ (Figure XXI). The two effects from EIS and δ counteract each other. On net, a lower EIS, corresponding with a lower δ, implies a higher optimal CO_2 prices.

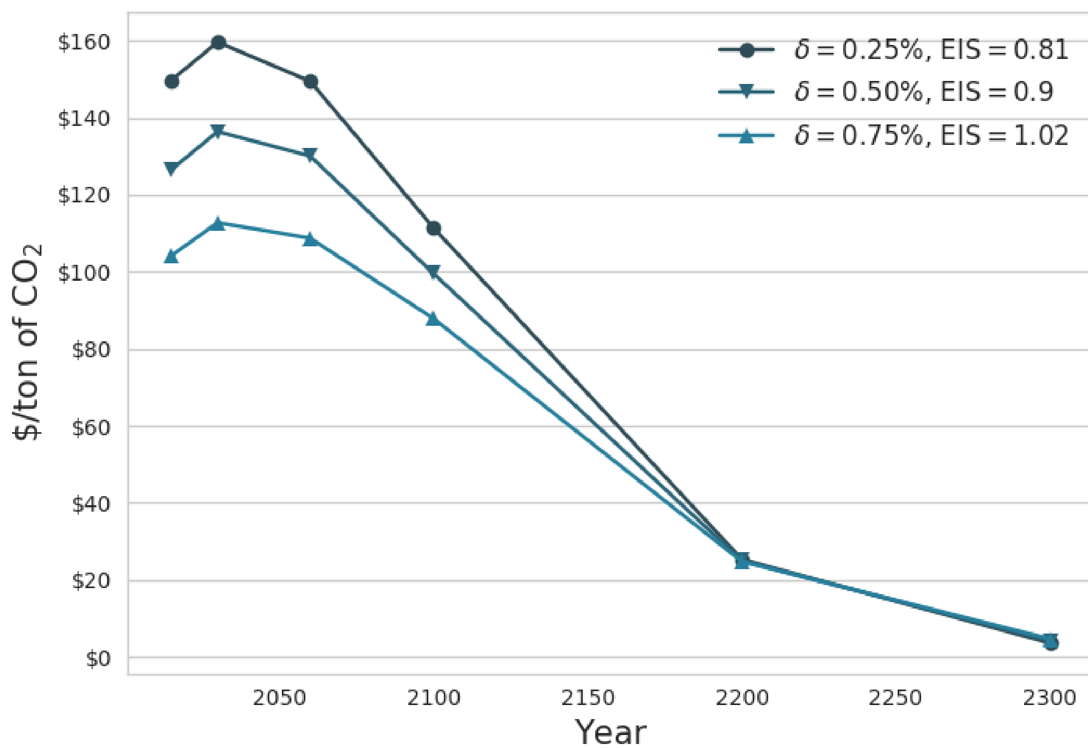

Figure XXI—Optimal CO_2 prices increase with decreasing pure rate of time preference, δ, holding real interest rates fixed at close to 3.11%,while adjusting EIS accordingly

Note that as EZ-Climate is calibrated to observed real interest rates, δ is a less important parameter than, for example, in a Ramsey-style model like DICE.

Lastly, Figure XXII shows the importance of the RA calibration. Optimal CO_2 prices decline under any EZ calibration. CRRA preferences lead to a declining CO_2 price only when risk aversion equals 10/9 ~ 1.1, when CRRA and EZ preference specifications coincide. We choose RA = 7 in our base-case calibration. Calibrations differ widely in the literature. See, e.g., Schroyen and Aarbu (2017) for a review of RA calibrations, showing significantly lower values in "welfare states" like Norway but higher values (>9) for countries like the United States. OECD averages for RA toward large income risks are close to 7 (Schroyen and Aarbu, 2017).

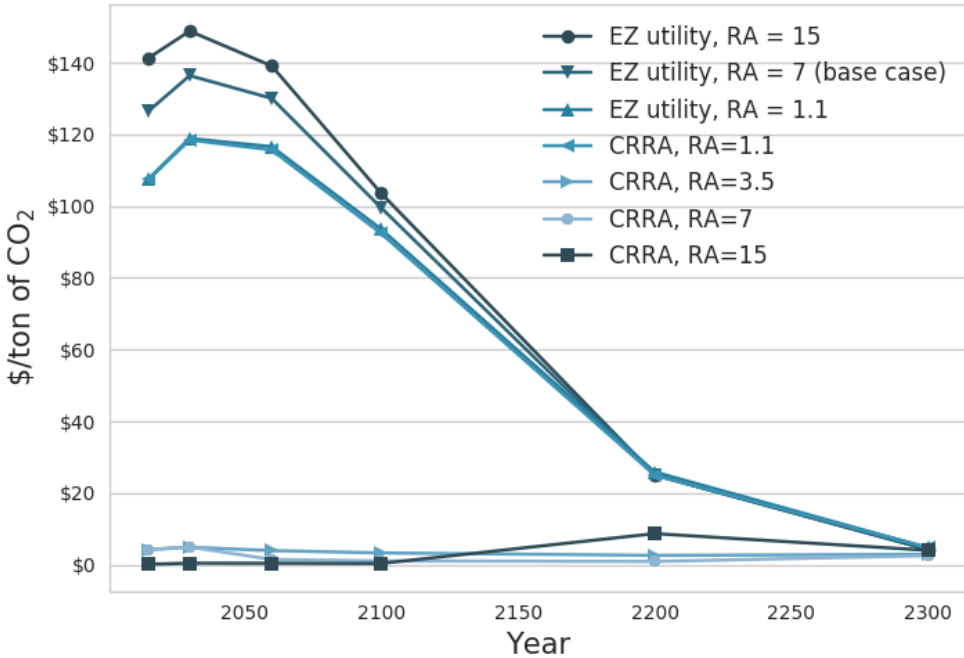

Figure XXII—Unlike with CRRA preferences, optimal CO_2 prices decline over time with EZ utility. The sole exception for CRRA preferences is with risk aversion (RA) = 10/9 ~ 1.1, when CRRA and EZ preferences coincide

3.5. Social cost of delay

Up to this point in the paper we have focused on analyzing the optimal CO_2 price, given that climate policy is assumed to be optimal over time. It is not. What then is the cost of a delay in pricing emissions at their optimal level?

We are hardly the first to point to the potential costs of delay (e.g., Nordhaus, 2016; Stern, 2015).[36] Nordhaus (2017b) suggests that the large increases in the optimal price in his preferred calibration of the DICE model, from around $5 per ton of CO_2 emitted in 2015 in early iterations published in the 1990s to around $30 per ton in the latest iteration, is largely a reflection of those costs of delay in implementing an optimal CO_2 price path.[37] However, this argument also points to the limits of models like DICE and the need for an optimal control model like EZ-Climate. In fact, looking at the price response to delay is positively misleading. The optimal CO_2 price in the second period, after it has been constrained to be $0 in the first period, is lower than in the unconstrained case. The price reflects the marginal benefits of additional emissions reductions, but the price alone does not show the enormous costs to society of not following the optimal path. That necessitates a look at the deadweight cost of delay.

[36] This discussion of the cost of delay is distinct from debates about the proper incorporation of climate risk, reflected, for example, in debates among Nordhaus (2015, 2013), Pindyck (2013, 2012), Stern (2013), and Weitzman (2011, 2009). See the discussion in the introductory section and references in footnotes 4 and 6.
[37] Nordhaus (2017a), meanwhile, decomposes the change further and counts "revisions in the treatment of the carbon cycle" as the largest single contributor, followed by revisions to the damage and utility functions.

We find that the avoidable deadweight cost of delay in pricing emissions is shockingly high. Employing EZ-Climate, we 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 utility of the unconstrained optimal solution. With our 15-year-long first period in the base case, consumption throughout the first period would have to increase by around 36% (Table IV), or by over $10 trillion each year, given current annual global consumption of around $30 trillion.

The cost of delay increases at a slightly faster rate than with the square of time. For example, decreasing the first-period length to 5 years, and, thus, having a 5-year instead of 15-year delay results in a total utility loss equivalent to increased consumption of over 11% over the 5 years. For a 10-year delay, the equivalent annual consumption loss over the period is almost 23%. Each year of delay causes the equivalent annual consumption loss *over the entire first period* to increase by roughly 2.3%, and it also increases the time interval of the loss, thus leading to a slightly more than quadratic rate of increase in deadweight loss of utility over time (Table IV). Given global consumption of around $30 trillion growing at 1.5% per year, a one-year delay in pricing emissions in our base case would cost the equivalent of about 2.3% or approximately $700 billion. A five-year delay would cost the equivalent of $17 trillion, and a 15 year delay would cost roughly $180 trillion, about six times current annual global consumption.

Table IV—Social cost of delay by first-period length

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

The marginal damage of emissions is much greater if mitigation is delayed. In our base case, the marginal benefit of reductions equals the carbon price of around $125 per ton of CO_2 in 2015 (see e.g. Figure VIII). If, instead, mitigation is constrained to zero in the first 15-year period, the marginal benefit of each ton of CO_2 reduced increases to over $340 per ton.

4. Future research and extensions

We have emphasized throughout the discussion here that the EZ-Climate model and its parameterization are *ad hoc* and merit considerable future scrutiny by the scientific community. What we believe is unique and innovative is the methodology and solution approach outlined here. The EZ-Climate model better takes account of climatic risks, without undue computational burdens.

There are many possible avenues for future work and research. First and foremost is a better calibration of some of the key parameters, perhaps led by climate damages in general and the probability and impact of tail risks in particular. The modular format of EZ-Climate allows for other, more sophisticated climate damages efforts to replace our current calibrations, while maintaining the fundamental model structure. Hsiang *et al.* (2017) is but the latest and perhaps most sophisticated entry into this literature, in this case focused on climate impacts in the United States. Our base case climate damages calibration assumes a *peakT* of 6 and *disaster_tail* of 18. While we explore various specifications, none is linked to actual probabilities and impacts of low-probability, high-consequence climatic events. The impacts of both parameters, and their interplay, is potentially large (Figure XIII). Moreover, our current modeling structure, while allowing for potentially large climate damages does not allow for true surprises, virtually eliminating catastrophic outcomes.

Kopp *et al.* (2016) suggests a starting point for further analysis: First, start with candidate 'climatic tipping elements' and take the growing body of empirical, econometric analyses to estimate their effects. Second, employ experiments with empirical, process-based impact models to assess the relative importance of different tipping elements. Third, examine social tipping elements such as civil conflict and their associated costs. Projects like the Climate Impacts Lab are beginning to quantify the impacts of some of these climatic and social tipping elements (Houser et al., 2015; Hsiang et al., 2017). Finally, where data are scarce, conduct structured expert elicitations to generate probability functions around tipping elements and their possible economic impacts (e.g., Cooke, 2013).[38] These calibrations are much needed to derive defensible estimates of the optimal CO_2 price.

A further extension is a reconsideration of how climatic risk and its interplay with economic variables is modeled. Climate damages, right now, are modeled to affect levels of consumption. What if they were to affect factor productivity and, thus, economic growth rates directly? While EZ-Climate, in its current iteration, relies on deterministic economic growth assumptions, a move to a stochastic economic growth framework would allow for exploring the interplay between climate and economic shocks (e.g., Bansal et al., 2016).

Another possible extension is around climate impacts on poor versus rich. EZ-Climate's representative agent framework models the average, globally representative consumer. One could easily imagine different climate impacts based on wealth and income. While the rich may have more capital at risk, it would likely be the poor who will be hurt proportionally more, as they are less able to adapt.

Lastly, any model like EZ-Climate needs to be seen in an appropriate context. Brock and Hansen (2017), among others, emphasize how climate-economic assessments are beset with at least three different types of uncertainty: climate risk in the traditional, Knightian sense of the word (Knight, 1921); ambiguity around model selection; and misspecification

---

[38] See Mach *et al.* (2017) on the broader need to consider expert judgement in environmental assessment, and how to do so.

of any one particular model.[39] Further improvements of EZ-Climate can hope to address the third kind of uncertainty. Addressing the first two will require an informed observer to be cautious about drawing conclusions. Some observers might point to the importance of calculating the optimal CO_2 price in the first place. Some others might conclude that any such model conveys false precision to policy-makers, who may make too much of any one single number. A model estimate of the optimal CO_2 price, in the end, is an important input into policy-makers' decisions, but it is just that. Good climate policy goes much beyond simply getting the price right. It considers the uncertainty of any such estimate, the potential consequences of model misspecifications (Brock and Hansen, 2017), as well as behavioral aspects of climate policy making (Wagner and Zeckhauser, 2011). Moreover, it takes into account political economy considerations of implementing alternative climate policies (Keohane et al., 1998; Meckling et al., 2017; Wagner et al., 2015).

5. Conclusion

An oft-told analogy in climate economics represents the climate system as a hard-to-navigate ocean liner. This image is often used to argue for early action through a slow and gradually increasing carbon price. Too strong a policy early on would be overly costly; a small course correction now will save us from hitting the far-off proverbial iceberg. There are indeed real costs of action. Tradeoffs abound.[40] But as EZ-Climate's base case optimal CO_2 price paths show, once we include a proper accounting of risk aversion and extreme events, this standard logic gets turned on its head: The optimal carbon price may, in fact, be high today, declining over time (e.g. Figure VIII).[41]

In EZ-Climate, that decline in optimal CO_2 prices over time reflects the rate at which information is revealed going forward, the degree of risk aversion, and the potential for technological progress and backstop technologies. Either way, the initial 'ocean liner' logic does not hold. Or perhaps it gets completed: for turning a large ship long down the line takes bearing off decisively and early, especially in a world of uncertain obstacles. The less certain we are about the climate risks facing us in future states of the world, the higher the optimal price on carbon today.

---

[39] See Heal (2017), Pindyck (2013, 2012), and Wagner and Zeckhauser (2017) for broader discussions of uncertainty in climate-economy models. See Lemoine and Rudik (2017a) for further guidance for recursive integrated assessment models.

[40] Typically modeled GHG emissions price paths, such as those feeding into the official U.S. SCC calculations, including perhaps most prominently Nordhaus's DICE model, follow this logic (U.S. Government Interagency Working Group on Social Cost of Carbon, 2015). Cf. footnote 1.

[41] Ulph and Ulph (1994) similarly derive conditions for a declining carbon price, though, in their case, driven by optimal resource extraction in the context of climate policy, invoking notions of the so-called "Green Paradox" (Jensen et al., 2015). Acemoglu *et al.* (2012) derives a high, "temporary" optimal price, together with a high subsidy for "clean" technologies, which eventually replace the "dirty" technology altogether and, thus, remove the need to price GHGs. Lemoine and Rudik (2017b), in turn, find that an efficient optimal CO_2 price first rises and then falls steeply over time.

Bibliography

Acemoglu, D., Aghion, P., Bursztyn, L., Hemous, D., 2012. The Environment and Directed Technical Change. Am. Econ. Rev. 102, 131–166. https://doi.org/10.1257/aer.102.1.131

Ackerman, F., Stanton, E.A., Bueno, R., 2013. Epstein–Zin Utility in DICE: Is Risk Aversion Irrelevant to Climate Policy? Environ. Resour. Econ. 56, 73–84. https://doi.org/10.1007/s10640-013-9645-z

Allcott, H., Greenstone, M., 2012. Is There an Energy Efficiency Gap? J. Econ. Perspect. 26, 3–28. https://doi.org/10.1257/jep.26.1.3

Arrow, K., Cropper, M., Gollier, C., Groom, B., Heal, G., Newell, R., Nordhaus, W., Pindyck, R., Pizer, W., Portney, P., Sterner, T., Tol, R.S.J., Weitzman, M., 2013. Determining Benefits and Costs for Future Generations. Science 341, 349–350. https://doi.org/10.1126/science.1235665

Bansal, R., Ochoa, M., 2011. Temperature, Aggregate Risk, and Expected Returns (Working Paper No. 17575). National Bureau of Economic Research.

Bansal, R., Ochoa, M., 2009. Temperature, Growth, and Asset Prices.

Bansal, R., Ochoa, M., Kiku, D., 2016. Climate Change and Growth Risks (Working Paper No. 23009). National Bureau of Economic Research. https://doi.org/10.3386/w23009

Bansal, R., Yaron, A., 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. J. Finance 59, 1481–1509. https://doi.org/10.1111/j.1540-6261.2004.00670.x

Barro, R.J., 2015. Environmental Protection, Rare Disasters and Discount Rates. Economica 82, 1–23. https://doi.org/10.1111/ecca.12117

Barro, R.J., 2009. Rare disasters, asset prices, and welfare costs. Am. Econ. Rev. 99, 243–264.

Barro, R.J., 2006. Rare Disasters and Asset Markets in the Twentieth Century. Q. J. Econ. 121, 823–866. https://doi.org/10.1162/qjec.121.3.823

Barro, R.J., Jin, T., 2011. On the Size Distribution of Macroeconomic Disasters. Econometrica 79, 1567–1589. https://doi.org/10.3982/ECTA8827

Barro, R.J., Ursua, J.F., 2008. Consumption disasters in the twentieth century. Am. Econ. Rev. 98, 58–63.

Belaia, M., Funke, M., Glanemann, N., 2017. Global Warming and a Potential Tipping Point in the Atlantic Thermohaline Circulation: The Role of Risk Aversion. Environ. Resour. Econ. 67, 93–125. https://doi.org/10.1007/s10640-015-9978-x

Brock, W.A., Hansen, L.P., 2017. Wrestling with Uncertainty in Climate Economic Models.

Broome, J., 2008. The ethics of climate change. Sci. Am. 298, 96–102.

Burke, M., Craxton, M., Kolstad, C.D., Onda, C., Allcott, H., Baker, E., Barrage, L., Carson, R., Gillingham, K., Graff-Zivin, J., Greenstone, M., Hallegatte, S., Hanemann, W.M., Heal, G., Hsiang, S., Jones, B., Kelly, D.L., Kopp, R., Kotchen, M., Mendelsohn, R., Meng, K., Metcalf, G., Moreno-Cruz, J., Pindyck, R., Rose, S., Rudik, I., Stock, J., Tol, R.S.J., 2016. Opportunities for advances in climate change economics. Science 352, 292–293. https://doi.org/10.1126/science.aad9634

Cai, Y., Judd, K.L., Lontzek, T.S., 2015. The social cost of carbon with economic and climate risks (No. 1504.06909v2). arXiv.
Cai, Y., Judd, K.L., Lontzek, T.S., 2013. The social cost of stochastic and irreversible climate change. National Bureau of Economic Research.
Cai, Y., Lenton, T.M., Lontzek, T.S., 2016. Risk of multiple interacting tipping points should encourage rapid CO2 emission reduction. Nat. Clim. Change 6, 520–525. https://doi.org/10.1038/nclimate2964
Charney, J.G., Arakawa, A., Baker, D.J., Bolin, B., Dickinson, R.E., Goody, R.M., Leith, C.E., Stommel, H.M., Wunsch, C.I., 1979. Carbon dioxide and climate: a scientific assessment. National Academy of Sciences, Washington, DC.
Coase, R.H., 1960. The Problem of Social Cost, in: Gopalakrishnan, C. (Ed.), Classic Papers in Natural Resource Economics. Palgrave Macmillan UK, pp. 87–137. https://doi.org/10.1057/9780230523210_6
Convery, F.J., Wagner, G., 2015. Reflections–Managing Uncertain Climates: Some Guidance for Policy Makers and Researchers. Rev. Environ. Econ. Policy 9, 304–320. https://doi.org/10.1093/reep/rev003
Cooke, R.M., 2013. Expert judgement assessment: Quantifying uncertainty on thin ice. Nat. Clim. Change 3, 311–312. https://doi.org/10.1038/nclimate1860
Cox, J.C., Ross, S.A., Rubinstein, M., 1979. Option pricing: A simplified approach. J. Financ. Econ. 7, 229–263. https://doi.org/10.1016/0304-405X(79)90015-1
Crost, B., Traeger, C.P., 2014. Optimal CO2 mitigation under damage risk valuation. Nat. Clim. Change 4, 631–636. https://doi.org/10.1038/nclimate2249
Daniel, K.D., Litterman, R.B., Wagner, G., 2016. Applying Asset Pricing Theory to Calibrate the Price of Climate Risk (Working Paper No. 22795). National Bureau of Economic Research. https://doi.org/10.3386/w22795
Dasgupta, P., 2008. Discounting climate change. J. Risk Uncertain. 37, 141–169. https://doi.org/10.1007/s11166-008-9049-6
Dell, M., Jones, B.F., Olken, B.A., 2012. Temperature Shocks and Economic Growth: Evidence from the Last Half Century. Am. Econ. J. Macroecon. 4, 66–95. https://doi.org/10.1257/mac.4.3.66
Diaz, D., Moore, F., 2017. Quantifying the economic risks of climate change. Nat. Clim. Change 7, 774–782. https://doi.org/10.1038/nclimate3411
Duffie, D., 2010. Dynamic Asset Pricing Theory. Princeton University Press.
Epstein, L.G., Farhi, E., Strzalecki, T., 2014. How Much Would You Pay to Resolve Long-Run Risk? Am. Econ. Rev. 104, 2680–2697.
Epstein, L.G., Zin, S.E., 1991. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis. J. Polit. Econ. 99, 263–286.
Epstein, L.G., Zin, S.E., 1989. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. Econometrica 57, 937–969. https://doi.org/10.2307/1913778
Gerarden, T.D., Newell, R.G., Stavins, R.N., 2015. Assessing the Energy-Efficiency Gap (Working Paper No. 20904). National Bureau of Economic Research. https://doi.org/10.3386/w20904
Gillingham, K., Palmer, K., 2014. Bridging the Energy Efficiency Gap: Policy Insights from Economic Theory and Empirical Evidence. Rev. Environ. Econ. Policy 8, 18–38. https://doi.org/10.1093/reep/ret021

Gollier, C., 2012. Pricing the Planet's Future: The Economics of Discounting in an Uncertain World. Princeton University Press.
Gollier, C., Weitzman, M.L., 2010. How should the distant future be discounted when discount rates are uncertain? Econ. Lett. 107, 350–353. https://doi.org/10.1016/j.econlet.2010.03.001
Golosov, M., Hassler, J., Krusell, P., Tsyvinski, A., 2014. Optimal Taxes on Fossil Fuel in General Equilibrium. Econometrica 82, 41–88. https://doi.org/10.3982/ECTA10217
Gordon, R.J., 2016. The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War. Princeton University Press.
Gordon, R.J., 2012. Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds (Working Paper No. 18315). National Bureau of Economic Research.
Goulder, L.H., 1995. Environmental taxation and the double dividend: A reader's guide. Int. Tax Public Finance 2, 157–183. https://doi.org/10.1007/BF00877495
Ha-Duong, M., Treich, N., 2004. Risk Aversion, Intergenerational Equity and Climate Change. Environ. Resour. Econ. 28, 195–207. https://doi.org/10.1023/B:EARE.0000029915.04325.25
Hall, R.E., 1988. Intertemporal Substitution in Consumption. J. Polit. Econ. 96, 339–357. https://doi.org/10.1086/261539
Hansen, L.P., Heaton, J.C., Li, N., 2008. Consumption Strikes Back? Measuring Long-Run Risk. J. Polit. Econ. 116, 260–302. https://doi.org/10.1086/588200
Hansen, L.P., Richard, S.F., 1987. The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models. Econometrica 55, 587–613. https://doi.org/10.2307/1913601
Havránek, T., 2015. Measuring Intertemporal Substitution: The Importance of Method Choices and Selective Reporting. J. Eur. Econ. Assoc. 13, 1180–1204. https://doi.org/10.1111/jeea.12133
Heal, G., 2017. The Economics of the Climate. J. Econ. Lit. 55, 1046–1063. https://doi.org/10.1257/jel.20151335
Heal, G., Park, J., 2016. Temperature Stress and the Direct Impact of Climate Change: A Review of an Emerging Literature. Rev. Environ. Econ. Policy rew007. https://doi.org/10.1093/reep/rew007
Houser, T., Hsiang, S., Kopp, R., Larsen, K., 2015. Economic Risks of Climate Change: An American Prospectus. Columbia University Press.
Hsiang, S., Kopp, R., Jina, A., Rising, J., Delgado, M., Mohan, S., Rasmussen, D.J., Muir-Wood, R., Wilson, P., Oppenheimer, M., Larsen, K., Houser, T., 2017. Estimating economic damage from climate change in the United States. Science 356, 1362–1369. https://doi.org/10.1126/science.aal4369
Huntington, H., 2011. EMF 25: Energy Efficiency and Climate Change Mitigation. Energy Model. Forum.
International Energy Agency, 2013. World Energy Outlook.
IPCC, 2013. Fifth Assessment Report: Climate Change.
IPCC, 2007. Fourth Assessment Report.
Jensen, S., Mohlin, K., Pittel, K., Sterner, T., 2015. An Introduction to the Green Paradox: The Unintended Consequences of Climate Policies. Rev. Environ. Econ. Policy 9, 246–265. https://doi.org/10.1093/reep/rev010

Jensen, S., Traeger, C.P., 2014. Optimal climate change mitigation under long-term growth uncertainty: Stochastic integrated assessment and analytic findings. Eur. Econ. Rev., Sustainability and Climate Change: From Theory to Pragmatic Policy 69, 104–125. https://doi.org/10.1016/j.euroecorev.2014.01.008
Kaufman, N., 2012. The bias of integrated assessment models that ignore climate catastrophes. Clim. Change 110, 575–595. https://doi.org/10.1007/s10584-011-0140-7
Keith, D.W., 2000. Geoengineering the climate: history and prospect. Annu. Rev. Energy Environ. 25, 245–284. https://doi.org/10.1146/annurev.energy.25.1.245
Keohane, N.O., Revesz, R.L., Stavins, R.N., 1998. The Choice of Regulatory Instruments in Environmental Policy. Harv. Environ. Law Rev. 22, 313.
Knight, F.H., 1921. Risk, uncertainty and profit. N. Y. Hart Schaffner Marx.
Kopp, R.E., Shwom, R., Wagner, G., Yuan, J., 2016. Tipping elements and climate-economic shocks: Pathways toward integrated assessment. Earths Future 4, 346–372. https://doi.org/10.1002/2016EF000362
Kreps, D.M., Porteus, E.L., 1978. Temporal Resolution of Uncertainty and Dynamic Choice Theory. Econometrica 46, 185–200. https://doi.org/10.2307/1913656
Le Quéré, C., Andrew, R.M., Canadell, J.G., Sitch, S., Korsbakken, J.I., Peters, G.P., Manning, A.C., Boden, T.A., Tans, P.P., Houghton, R.A., Keeling, R.F., Alin, S., Andrews, O.D., Anthoni, P., Barbero, L., Bopp, L., Chevallier, F., Chini, L.P., Ciais, P., Currie, K., Delire, C., Doney, S.C., Friedlingstein, P., Gkritzalis, T., Harris, I., Hauck, J., Haverd, V., Hoppema, M., Klein Goldewijk, K., Jain, A.K., Kato, E., Körtzinger, A., Landschützer, P., Lefèvre, N., Lenton, A., Lienert, S., Lombardozzi, D., Melton, J.R., Metzl, N., Millero, F., Monteiro, P.M.S., Munro, D.R., Nabel, J.E.M.S., Nakaoka, S., O'Brien, K., Olsen, A., Omar, A.M., Ono, T., Pierrot, D., Poulter, B., Rödenbeck, C., Salisbury, J., Schuster, U., Schwinger, J., Séférian, R., Skjelvan, I., Stocker, B.D., Sutton, A.J., Takahashi, T., Tian, H., Tilbrook, B., Laan-Luijkx, I.T. van der, Werf, G.R. van der, Viovy, N., Walker, A.P., Wiltshire, A.J., Zaehle, S., 2016. Global Carbon Budget 2016. Earth Syst. Sci. Data 8, 605–649. https://doi.org/10.5194/essd-8-605-2016
Lemoine, D., 2015. The Climate Risk Premium. Univ. Ariz. Work. Pap. 15-01.
Lemoine, D., Rudik, I., 2017a. Managing Climate Change Under Uncertainty: Recursive Integrated Assessment at an Inflection Point. Annu. Rev. Resour. Econ. 9, 117–142. https://doi.org/10.1146/annurev-resource-100516-053516
Lemoine, D., Rudik, I., 2017b. Steering the Climate System: Using Inertia to Lower the Cost of Policy. Am. Econ. Rev. 107, 2947–2957. https://doi.org/10.1257/aer.20150986
Lemoine, D., Traeger, C., 2014. Watch Your Step: Optimal Policy in a Tipping Climate. Am. Econ. J. Econ. Policy 6, 137–166. https://doi.org/10.1257/pol.6.1.137
Lemoine, D., Traeger, C.P., 2016. Economics of tipping the climate dominoes. Nat. Clim. Change 6, 514–519. https://doi.org/10.1038/nclimate2902
Lynas, M., 2008. Six Degrees: Our Future on a Hotter Planet. National Geographic.
Mach, K.J., Mastrandrea, M.D., Freeman, P.T., Field, C.B., 2017. Unleashing expert judgment in assessment. Glob. Environ. Change 44, 1–14. https://doi.org/10.1016/j.gloenvcha.2017.02.005

Marten, A.L., Newbold, S.C., 2012. Estimating the social cost of non-CO2 GHG emissions: Methane and nitrous oxide. Energy Policy, Renewable Energy in China 51, 957–972. https://doi.org/10.1016/j.enpol.2012.09.073
Martin, I.W.R., 2012a. On the Valuation of Long-Dated Assets. J. Polit. Econ. 120, 346–358. https://doi.org/10.1086/666527
Martin, I.W.R., 2012b. Consumption-based asset pricing with higher cumulants. Rev. Econ. Stud. 80, 745–773.
Martin, I.W.R., 2008. Disasters and the Welfare Cost of Uncertainty. Am. Econ. Rev. 98, 74–78. https://doi.org/10.1257/aer.98.2.74
Mastrandrea, M.D., Schneider, S.H., 2001. Integrated assessment of abrupt climatic changes. Clim. Policy 1, 433–449. https://doi.org/10.3763/cpol.2001.0146
McKinsey, 2009a. Unlocking energy efficiency in the US Economy. McKinsey & Company, London.
McKinsey, 2009b. Pathways to a Low-Carbon Economy Version 2 of the Global Greenhouse Gas Abatement Cost Curve. McKinsey & Company, Stockholm.
Meckling, J., Sterner, T., Wagner, G., 2017. Policy sequencing toward decarbonization. Nat. Energy. https://doi.org/10.1038/s41560-017-0025-8
Mehra, R., Prescott, E.C., 1985. The equity premium: A puzzle. J. Monet. Econ. 15, 145–161. https://doi.org/10.1016/0304-3932(85)90061-3
Moore, F.C., Diaz, D.B., 2015. Temperature impacts on economic growth warrant stringent mitigation policy. Nat. Clim. Change 5, 127–131. https://doi.org/10.1038/nclimate2481
National Academy of Sciences, 2017. Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide.
National Research Council, 2015a. Climate Intervention: Carbon Dioxide Removal and Reliable Sequestration. National Academies Press, Washington, D.C.
National Research Council, 2015b. Climate Intervention: Reflecting Sunlight to Cool Earth. National Academies Press.
Nordhaus, W.D., 2017a. Evolution of Assessments of the Economics of Global Warming: Changes in the DICE model, 1992 – 2017 (Working Paper No. 23319). National Bureau of Economic Research. https://doi.org/10.3386/w23319
Nordhaus, W.D., 2017b. Revisiting the social cost of carbon. Proc. Natl. Acad. Sci. 114, 1518–1523. https://doi.org/10.1073/pnas.1609244114
Nordhaus, W.D., 2016. Projections and Uncertainties About Climate Change in an Era of Minimal Climate Policies (Working Paper No. 22933). National Bureau of Economic Research.
Nordhaus, W.D., 2015. A New Solution: The Climate Club. N. Y. Rev. Books.
Nordhaus, W.D., 2013. The Climate Casino: Risk, Uncertainty, and Economics for a Warming World. Yale University Press.
Nordhaus, W.D., 1994. Expert opinion on climatic change. Am. Sci. 82, 45–51.
Nordhaus, W.D., 1992. An Optimal Transition Path for Controlling Greenhouse Gases. Science 258, 1315–1319. https://doi.org/10.1126/science.258.5086.1315
Nordhaus, W.D., 1991. To Slow or Not to Slow: The Economics of The Greenhouse Effect. Econ. J. 101, 920–937. https://doi.org/10.2307/2233864
Nordhaus, W.D., Boyer, J., 2000. Warming the World: Economic Models of Global Warming. MIT Press.
Nordhaus, W.D., Sztorc, P., 2013. DICE 2013R: Introduction and User's Manual.

Pigou, A.C., 1920. The Economics of Welfare. Palgrave Macmillan.
Pindyck, R.S., 2013. The Climate Policy Dilemma. Rev. Environ. Econ. Policy 7, 219–237.
Pindyck, R.S., 2012. Uncertain outcomes and climate change policy. J. Environ. Econ. Manag. 63, 289–303. https://doi.org/10.1016/j.jeem.2011.12.001
Previdi, M., Liepert, B.G., Peteet, D., Hansen, J., Beerling, D.J., Broccoli, A.J., Frolking, S., Galloway, J.N., Heimann, M., Le Quéré, C., Levitus, S., Ramaswamy, V., 2013. Climate sensitivity in the Anthropocene. Q. J. R. Meteorol. Soc. 139, 1121–1131. https://doi.org/10.1002/qj.2165
Rietz, T.A., 1988. The equity risk premium: a solution. J. Monet. Econ. 22, 117–131. https://doi.org/10.1016/0304-3932(88)90172-9
Roe, G.H., Baker, M.B., 2007. Why Is Climate Sensitivity So Unpredictable? Science 318, 629–632. https://doi.org/10.1126/science.1144735
Schroyen, F., Aarbu, K.O., 2017. Attitudes Towards Large Income Risk in Welfare States: An International Comparison (SSRN Scholarly Paper No. ID 3089690). Social Science Research Network, Rochester, NY.
Shiller, R., 2000. Irrational Exuberance. Princeton University Press.
Shindell, D., Borgford-Parnell, N., Brauer, M., Haines, A., Kuylenstierna, J.C.I., Leonard, S.A., Ramanathan, V., Ravishankara, A., Amann, M., Srivastava, L., 2017. A climate policy pathway for near- and long-term benefits. Science 356, 493–494. https://doi.org/10.1126/science.aak9521
Socolow, R., Desmond, M., Aines, R., Blackstock, J., Bolland, O., Kaarsberg, T., Lewis, N., Mazzotti, M., Pfeffer, A., Sawyer, K., et al., 2011. Direct air capture of CO2 with chemicals: a technology assessment for the APS Panel on Public Affairs. American Physical Society.
Stern, N., 2015. Why Are We Waiting?: The Logic, Urgency, and Promise of Tackling Climate Change. MIT Press.
Stern, N., 2013. 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. https://doi.org/10.1257/jel.51.3.838
Sterner, T., Persson, U.M., 2008. An Even Sterner Review: Introducing Relative Prices into the Discounting Debate. Rev. Environ. Econ. Policy 2, 61–76. https://doi.org/10.1093/reep/rem024
Summers, L., Zeckhauser, R., 2008. Policymaking for posterity. J. Risk Uncertain. 37, 115–140. https://doi.org/10.1007/s11166-008-9052-y
Thimme, J., 2017. Intertemporal Substitution in Consumption: A Literature Review. J. Econ. Surv. 31, 226–257. https://doi.org/10.1111/joes.12142
Traeger, C.P., 2015. Analytic Integrated Assessment and Uncertainty.
Ulph, A., Ulph, D., 1994. The Optimal Time Path of a Carbon Tax. Oxf. Econ. Pap. 46, 857–868.
U.S. Government Interagency Working Group on Social Cost of Carbon, 2015. Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866.
Wagner, G., Kåberger, T., Olai, S., Oppenheimer, M., Rittenhouse, K., Sterner, T., 2015. Energy policy: Push renewables to spur carbon pricing. Nature 525, 27–29. https://doi.org/10.1038/525027a

Wagner, G., Weitzman, M.L., 2015. Climate Shock: The Economic Consequences of a Hotter Planet. Princeton University Press, Princeton, NJ.
Wagner, G., Zeckhauser, R.J., 2017. Confronting Deep and Persistent Climate Uncertainty. Harvard Kennedy School, Faculty Research Working Paper Series, No. 16-025.
Wagner, G., Zeckhauser, R.J., 2011. Climate policy: hard problem, soft thinking. Clim. Change 110, 507–521. https://doi.org/10.1007/s10584-011-0067-z
Weil, P., 1990. Nonexpected Utility in Macroeconomics. Q. J. Econ. 105, 29–42. https://doi.org/10.2307/2937817
Weil, P., 1989. The equity premium puzzle and the risk-free rate puzzle. J. Monet. Econ. 24, 401–421. https://doi.org/10.1016/0304-3932(89)90028-7
Weitzman, M.L., 2015. Book Review—A Review of William Nordhaus' The Climate Casino: Risk, Uncertainty, and Economics for a Warming World. Rev. Environ. Econ. Policy 9, 145–156. https://doi.org/10.1093/reep/reu019
Weitzman, M.L., 2011. Fat-Tailed Uncertainty in the Economics of Catastrophic Climate Change. Rev. Environ. Econ. Policy 5, 275–292. https://doi.org/10.1093/reep/rer006
Weitzman, M.L., 2009. On Modeling and Interpreting the Economics of Catastrophic Climate Change. Rev. Econ. Stat. 91, 1–19. https://doi.org/10.1162/rest.91.1.1
Weitzman, M.L., 2007a. Subjective expectations and asset-return puzzles. Am. Econ. Rev. 97, 1102–1130.
Weitzman, M.L., 2007b. A Review of the *Stern Review on the Economics of Climate Change.* J. Econ. Lit. 45, 703–724. https://doi.org/10.1257/jel.45.3.703
Zeckhauser, R.J., 2006. Investing in the Unknown and Unknowable. Capital. Soc. 1/2.