Carbon Dioxide as a Risky Asset

Adam Michael Bauer*[1], Cristian Proistosescu[2,3], and Gernot Wagner[4]

[1]*Department of Physics, University of Illinois Urbana-Champaign, 1110 W Green St Loomis Laboratory, Urbana, IL 61801, USA*
[2]*Department of Climate, Meteorology, and Atmospheric Sciences, University of Illinois Urbana-Champaign, 1301 W Green St, Urbana, IL 61801, USA*
[3]*Department of Earth Science and Environmental Change, University of Illinois Urbana-Champaign, 1301 W Green St, Urbana, IL 61801, USA*
[4]*Columbia Business School, 665 W 130th St, New York, NY 10027, USA*

Forthcoming in *Climatic Change*
March 13, 2024

Abstract

We develop a financial-economic model for carbon pricing with an explicit representation of decision making under risk and uncertainty that is consistent with the Intergovernmental Panel on Climate Change's sixth assessment report. We show that risk associated with high damages in the long term leads to stringent mitigation of carbon dioxide (CO_2) emissions in the near term, and find that this approach provides economic support for stringent warming targets across a variety of specifications. Our results provide insight into how a systematic incorporation of climate-related risk influences optimal emissions abatement pathways.

JEL: G0, G12, Q51, Q54
Keywords: Climate risk, climate policy, asset pricing, cost of carbon

*Corresponding author email: adammb4@illinois.edu

1 Introduction

Climate change's impact on the economy first gained prominence in the economics literature some 30 years ago, when the first climate-economic Integrated Assessment Model (IAM) calculated the cost of a marginal ton of carbon dioxide (CO_2) emissions to society, coined the 'Social Cost of Carbon' (SCC) (Nordhaus, 1992). IAMs have since taken center stage in climate policy discussions, with the resulting SCC estimates being utilized as benchmarks by companies and governments worldwide (World Bank, 2021). To date the most prominent IAM by far – the dynamic integrated climate-economy, or DICE – evaluates climate change impacts within the context of a standard Ramsey growth economy (Nordhaus, 2017; Barrage and Nordhaus, 2023). In this approach, a global social planner considers tradeoffs between emitting CO_2 and incurring damages both now and, largely, in the future, versus abating CO_2 emissions now at some cost. Performing a benefit-cost analysis results in a presently-low and rising optimal SCC over time, with significant global average warming by 2100. Recent efforts have yielded comparatively higher SCC estimates; Rennert et al. (2022), for example, calculates a central SCC of $185 but did not explore the optimal control problem of weighing the benefits and costs of abating CO_2 emissions.[1] It is notable that DICE's optimal warming projections are significantly larger than each warming target – 1.5 °C and 2 °C by 2100 – established in the 2015 Paris Agreement. While this inconsistency has called into question the authority of such models in the climate policy discussion to some (Pindyck, 2013; Stern, 2013), DICE has been made consistent with a warming target of 1.5 °C with alternative, updated damages and discount rate modules (Hänsel et al., 2020).

A limitation of DICE is that it lacks a comprehensive representation of decision-making under risk and uncertainty, a core feature of many 'alternative' climate-economic models (Cai et al., 2016; Cai and Lontzek, 2019; Daniel et al., 2019; Barnett et al., 2020). This is important, as climate change projections are inherently probabilistic, with low probability, extreme impact outcomes presenting the most significant risk to the climate-economic system (Weitzman, 2009). The inherently unpredictable nature of the impacts of climate change has led some to think of climate policy as a form of "insurance" to be taken out against high climate damages (Weitzman, 2012). Conventional IAMs do not allow for such considerations in determining their policy projections. Put in financial-economic terms: conventional IAMs do not allow individuals to 'hedge' against climate impacts.

To address this, there have been considerable advances in climate-economic modeling that include the effects of risk and uncertainty on the SCC and on optimal policy responses to climate change; see Lemoine and Rudik (2017) for a comprehensive review, while Cai and Lontzek (2019) and Lemoine (2021) represent seminal works for including climate-related risk in IAMs.[2] We contribute to this extensive literature by introducing the carbon asset pricing model AR6 (CAP6), a climate-economy IAM that builds on previous financial asset pricing climate-economy models (Daniel et al., 2016, 2019). Our paper makes three primary contributions. The first is along methodological lines: we distill each working group report in the sixth assessment report (AR6) issued by the Intergovernmental Panel on Climate Change (IPCC) (Intergovernmental Panel on Climate Change, 2021, 2022a,b) into workable

[1] This SCC estimate represents a significant increase from the U.S. Interagency Working Group's central estimate of ∼$50 (Committee on Assessing Approaches to Updating the Social Cost of Carbon et al., 2017) and is in line with the U.S. Environmental Protection Agency's recent draft estimates that report a central value of $190 (National Center for Energy Economics, 2022).

[2] We provide a more thorough literature review in Online Appendix A.

IAM components.[3] This allows our model to be up-to-date with the state-of-the-art calibrations for critical model components. Notably, we formulate a new marginal abatement cost curve (MACC) based on AR6 data, providing an update to the well-known McKinsey & Company (2013) MACC.

The second contribution is a computation of optimal carbon prices and associated mitigation policy. Following Daniel et al. (2016, 2019), we embed a representative agent in a binomial, path-dependent tree that allows for risk assessment to endogenously evolve over time. The agent maximizes the Epstein-Zin-Weil utility (Epstein and Zin, 1989; Weil, 1990; Epstein and Zin, 1991) at every node in the tree such that the present-day utility is maximized. Agent discount rates are calibrated to be in-line with a recent expert elicitation (Drupp et al., 2018) and the U.S. Environmental Protection Agency (EPA) latest estimates for the SCC (National Center for Energy Economics, 2022). Notably, we find that the optimal expected warming in our EPA-consistent calibrations is in line with the 2100 warming targets established in the Paris agreement. We find that even if we were pessimistic about the cost of mitigation estimates provided by the IPCC, the EPA-consistent calibration of CAP6 would still support limiting warming to less than 2 °C by 2100, with a discount rate of 2% or lower.[4]

In computing optimal mitigation strategies, we capture uncertainty associated with both climate damages and global temperature rise. For damages, we capture both parametric uncertainty inherent to a given damage function, as well as structural uncertainty associated with different damage function shapes; in other words, in addition to Monte Carlo sampling damage levels for a given damage function, we also account for the fact that it is difficult to determine which damage function is correct in the first place (Pindyck, 2013; Intergovernmental Panel on Climate Change, 2022a). To our knowledge, we are the first to capture this dimension of climate-economic uncertainty. We also account for the marginal damages associated with a probabilistic assessment of climate tipping points (Lenton et al., 2008; Dietz et al., 2021).

Our final contribution is a sensitivity analysis that allows us to identify how each exogenous assumption drives model output. We show that while the expected carbon price depends on the emissions baseline, the expected temperature rise, level of CO_2 concentrations, and incurred economic damages does not. This suggests that our model robustly calculates an economically optimal temperature level for a given calibration; the price of actualizing this temperature level varies across baselines, owing to assumptions about how much emissions are decreasing independently of the policy implemented in CAP6. We find that price uncertainty is dominated by discounting in the near-term and the technological growth rate in the far-term. On the other hand, temperature rise, CO_2 concentration level, and economic damage uncertainty is dominated by discounting for much longer than CO_2 prices, as early inaction leads to warming that cannot be undone later by spending more on abatement (in the absence of significant net-negative emissions or solar geoengineering).

We proceed by presenting the socio-economic setup of CAP6 in section 2, the climate emulator in section 3, and our calibration in section 4. We discuss our results in section 5; section 6 concludes. (For section 2, we provide a brief summary paragraph with key equations and figures for readers who

[3] Nielsen-Gammon and Behl (2021) highlight the need and urgency for standardized, state-of-the-art climate and economic components based on the most up-to-date research for climate-economic modeling.

[4] This rate is significantly below Barrage and Nordhaus (2023)'s "preferred" rate of 4.5% in 2020, but well within the range that has emerged as a broad consensus among economists (Council of Economic Advisors, 2017; Drupp et al., 2018; Newell et al., 2022).

wish to skip the full technical description of our model components.)

2 Socio-economic framework

We consider a representative agent with Epstein-Weil-Zin utility given by (2.1), and embed this individual in a binomial tree structure where their utility is maximized. CO_2 emissions (without any agent mitigation action) follow the shared socio-economic projections used by the IPCC (Figure 2). Climate damage functions are calibrated to IPCC working group (WG) II data (see Figure 3) and our uncertainty parameterization captures both epistemic and parametric uncertainty in the damage functions. Finally, we employ (2.12) as our marginal abatement cost curve (Figure 4) and provide two calibrations: our 'main specification' based solely on the data in AR6, and the 'no free lunches' calibration, which excludes negative costs in the AR6 data.

2.1 Economic utility

CAP6 considers a representative agent with recursive preferences who maximizes their utility throughout time. We choose Epstein-Zin-Weil preferences (Epstein and Zin, 1989; Weil, 1990; Epstein and Zin, 1991), henceforth abbreviated as 'EZ', because of their unique feature of separating risk across states of time and states of nature. This distinction has been shown to be especially relevant for climate economic studies, where risk considerations across different dimensions are key to the outcome (e.g., Cai and Lontzek, 2019, among many others). The discrete time utility, U_t, of a representative agent with EZ preferences is given by

$$U_t = \left([1-\beta]c_t^\rho + \beta\left[\mathbb{E}_t\left(U_{t+1}^\alpha\right)\right]^{\rho/\alpha}\right)^{1/\rho}, \tag{2.1}$$

where $\beta := (1+\delta)^{-1} > 0$ and $\delta > 0$ is the pure rate of time preference (PRTP), $c_t > 0$ is the consumption at time t, $\rho := 1 - 1/\sigma$ and $\sigma > 0$ is the elasticity of intertemporal substitution (EIS), $\alpha := 1 - \psi$ and $\psi > 0$ is agent risk aversion (RA), and $\mathbb{E}_t$ is the expectation operator at time t. When $\alpha = \rho$ (that is, when $\psi = 1/\sigma$), (2.1) collapses into the von Neumann and Morgenstern (1947) expected utility index. Assuming an exogenous growth rate of consumption $g > 0$, in the final period occurring at time T, the utility is given by

$$U_T = \left[\frac{1-\beta}{1-\beta(1+g)^\rho}\right]^{1/\rho} c_T. \tag{2.2}$$

Note that, in the EZ framework, risk aversion across time is parameterized by σ, whereas risk aversion across states of nature is parameterized by ψ.

2.1.1 Tree structure

Following Daniel et al. (2016, 2019), agent utility in CAP6 is optimized within the structure of a binomial tree, therefore embedding the representative agent in a *finite horizon probability landscape.*

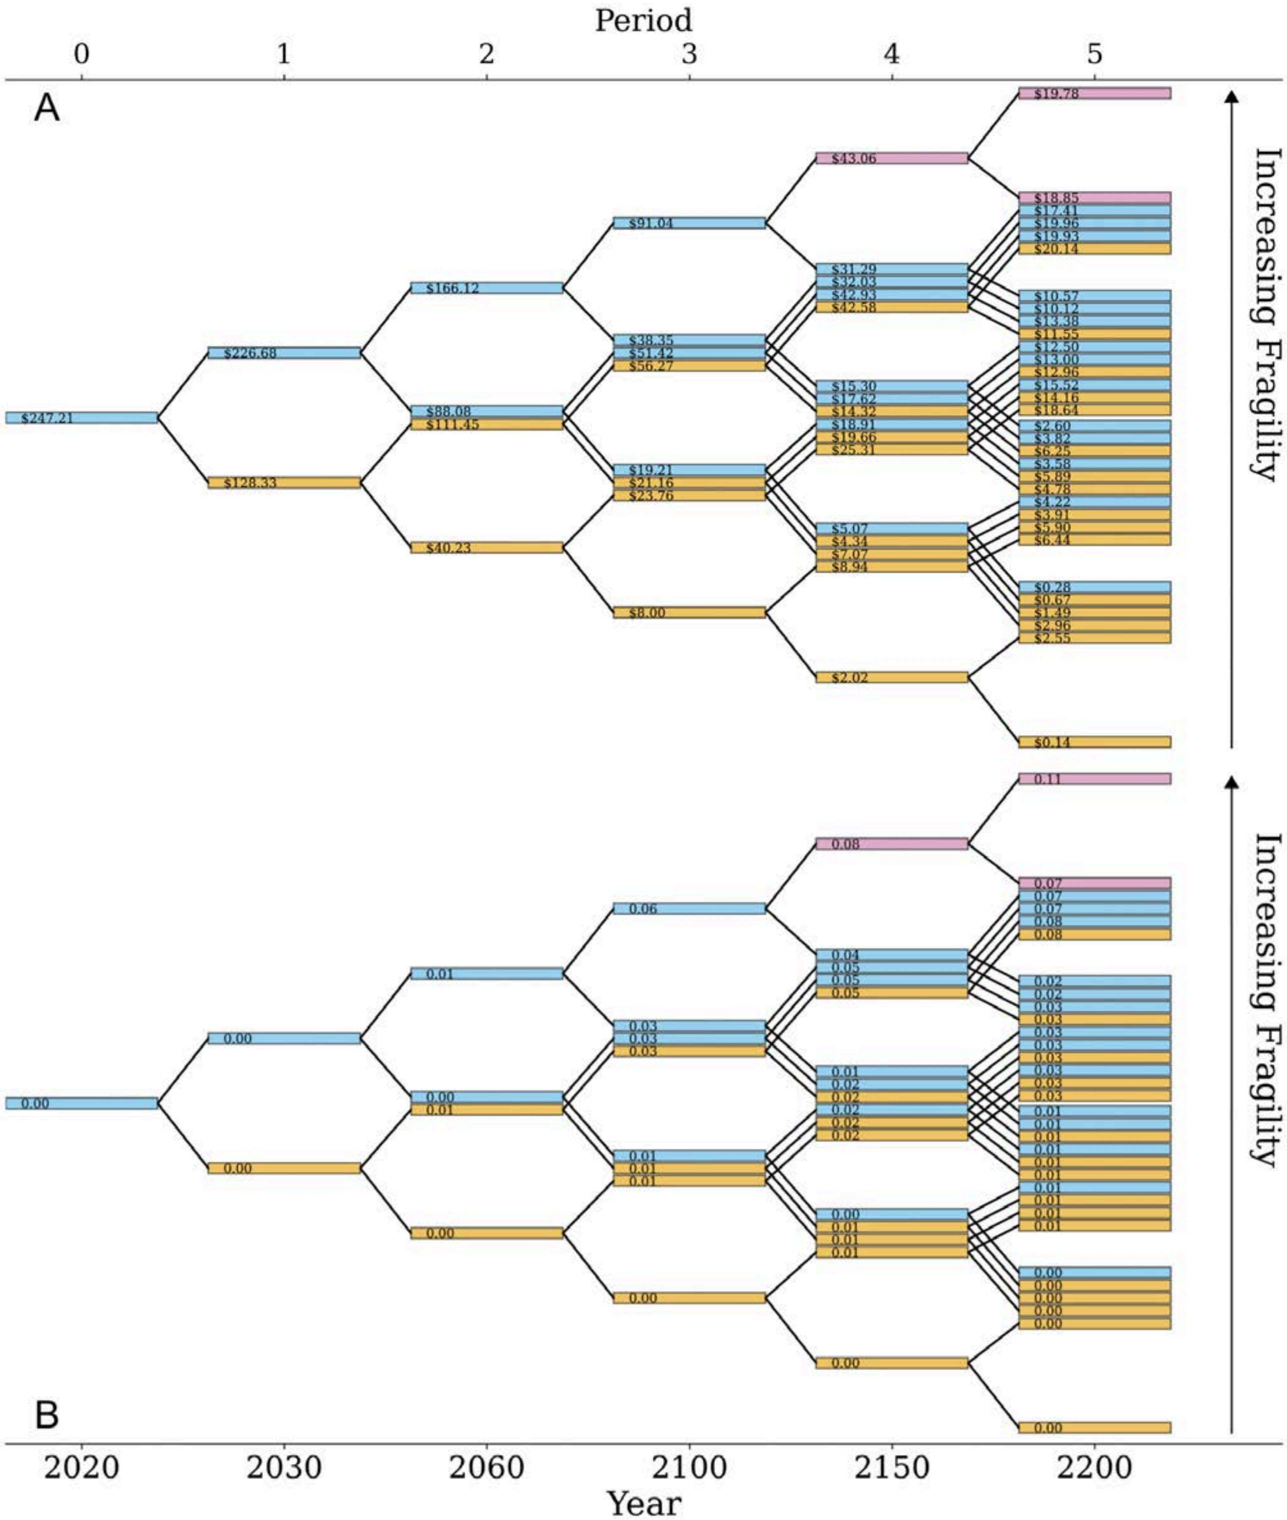

Figure 1: Cost of CO_2 (panel A) and agent experienced climate damages (panel B) at each node. In both panels, we highlight the accessible future states of two agents: one in 2150 (pink boxes) and one in 2030 (gold boxes).

Note: Values are taken from our 2% discount rate featured model run, main specification.

This follows a standard approach employed in financial economics (Cox et al., 1979), and one useful to solve EZ-style models numerically (Epstein and Zin, 1991).

The binomial tree structure of CAP6 is a representation of a time-evolving two dimensional probability distribution of climate damages (see Figure 1 for a schematic). The first dimension is time, while the second is "fragility", the latter of which encodes the potential for high or low climate damages at a moment in time. Throughout, we will refer to the fragility coordinate at a time t as $\theta_t > 0$. Framing the tree structure as a representation of a two dimensional probability distribution allows for the roles of σ and ψ to be clarified: σ parameterizes risk aversion along the *time* dimension, while ψ parameterizes risk aversion along the "*fragility*" dimension. We choose to orient the fragility coordinate such that high (low, resp.) fragility is associated with high (low, resp.) climate damages. By allowing for many agent decisions, and thus the generation of numerous nodes, we are able to coarsely represent the space of possible fragilities, therefore spanning many possible states of the climate and climate impacts. Note that in the limit of infinitely many decisions, fragility is normally distributed owing to every future state being equally likely, so as to not bias any outcome (be it sanguine or catastrophic) within the model structure.

This structure allows for agent risk assessment to evolve endogenously; as an example, consider two agents, one in 2150 and one in 2030 (see Figure 1). The agent in 2150 has only two future states accessible to them from their position in the tree; this represents an individual who knows well the impact of the climate on the economy. The agent in 2030 has a significantly higher number of future states accessible to them; they know less about how climate change impacts the economy, which influences their decision making, as they have to weigh several possible futures with high and low climate damages (or "fragility") all at once.

This approach has the advantage of being easily computationally tractable, while maintaining a structurally endogenous representation of risk and uncertainty resolution. Moreover, it allows for a transparent interpretation of model results and ample sensitivity analyses, which enables our variance decomposition results in § 5.3.1. However, it does suffer from drawbacks: more modern (and computationally expensive and technically challenging) models are able to solve similar optimization problems in continuous-time, on infinite horizons, or both (Bretschger and Vinogradova, 2014; Cai and Lontzek, 2019; Van Den Bremer and Van Der Ploeg, 2021). These considerations can matter for model results: for example, the time horizon used for climate policy models matters owing to the long residency time of CO_2 in the atmosphere. If one sets the time horizon of the model to 2200, then the net-benefits of a unit of CO_2 abatement in 2190 would matter less than one in 2020 because the benefits would not be given time to materialize. Nevertheless, a number of prominent IAMs used in climate policy consider finite horizons (perhaps most notably, the DICE model is solved on a finite horizon, see Nordhaus, 2017; Barrage and Nordhaus, 2023) and our model falls into this class. Moreover, our choice to truncate the time horizon at 2250 aligns with the time where we assume the world reaches net zero emissions without any additional policy in CAP6, which would make, from the policy perspective taken in our model, a carbon tax obsolete (see Figure 2).

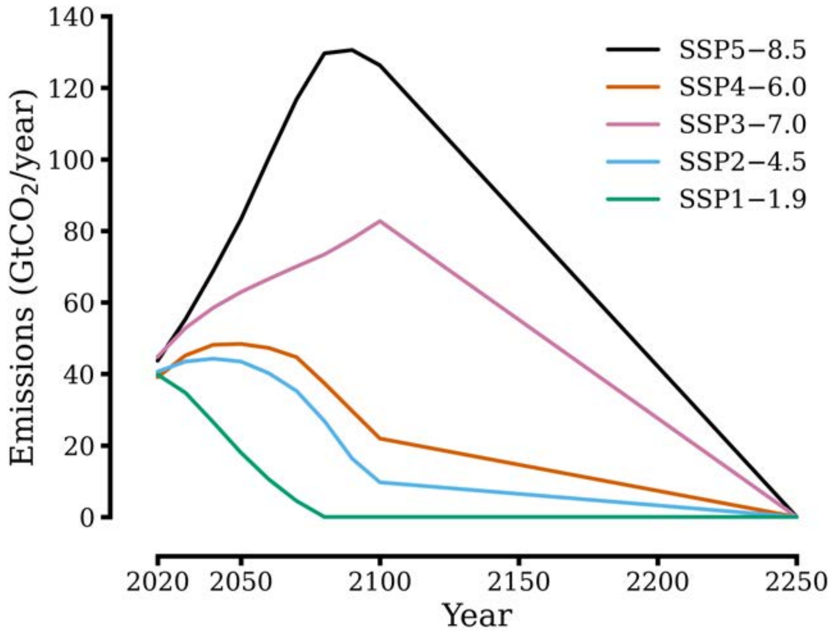

Figure 2: Emissions baselines with their extensions to 2250.

2.1.2 Statement of utility optimization problem

Consider a representative agent embedded within a path-dependent binomial tree with T decision periods, leading to $2^T - 1$ total tree nodes. The individual resides within a standard endowment economy (Summers and Zeckhauser, 2008), where at every period time t they are given an amount $\bar{c}_t > 0$ such that $\bar{c}_t = \bar{c}_0(1+g)^t$. Without loss of generality, set $\bar{c}_0$ to unity. They cannot consume all of $\bar{c}_t$, however, owing to both climate change and climate policy. Climate change can cause the agent to lose some amount of $\bar{c}_t$ due to climate damages, $\mathcal{D}_t \geq 0$. Climate policy allows them to spend some amount of $\bar{c}_t$ to reduce their impact on future climate by mitigating some fraction of emissions x_t with total cost κ_t. The consumption of the agent at each time $t \in \{0, 1, 2, ..., T\}$ is determined by

$$c_0 = \bar{c}_0 \left(1 - \kappa_0(x_0)\right), \tag{2.3}$$

$$c_t = \bar{c}_t \left(1 - \kappa_t(x_t)\right)\left(1 - \mathcal{D}_t(\Psi_t, \theta_t)\right), \qquad \text{for } t \in \{1, 2, ..., T-1\}, \tag{2.4}$$

$$c_T = \bar{c}_T \left(1 - \mathcal{D}_T(\Psi_T, \theta_T)\right), \tag{2.5}$$

where Ψ_t is the cumulative CO_2 emissions. We choose $T = 6$ decision periods in all the calculations in that follow, with our initial and final year being 2020 and 2250, respectively.[5] The net discounted EZ-utility is then maximized to obtain the optimal carbon prices and mitigation policies in § 5; see Online Appendix B for more details on our optimization.

2.2 Emission baselines

There is considerable uncertainty when choosing a 'business-as-usual' emissions scenario for climate-economy IAMs (Hausfather and Peters, 2020). One approach is for the emissions to be a result of economic output (e.g., Golosov et al., 2014). This approach has the advantage of making the emissions baseline endogenous; however, it also tends to exclude important processes relevant to the level and rate of fossil fuel emissions, such as friction in the diffusion of clean energy technologies, which can be captured by more sophisticated energy systems IAMs.

This concern motivates the second approach commonly used by the IPCC, which is to supply a given IAM with a stream of CO_2 emissions exogenously based on plausible future emissions scenarios. The shared socio-economic pathways (SSPs) shown in Figure 2 are an example of this approach, where each baseline represents a "storyline" for future global and regional economic development based on the level of challenges faced by policymakers in mitigation and adaptation. For example, SSP5 is a fossil fuel-based development storyline, with high levels of challenge to mitigation (because of significant fossil fuel development) and low challenges to adaptation (because of expanded wealth). SSP1, on the other hand, is a more sustainable route, with low challenges to both mitigation (because of renewable energy expansion) and adaptation (because of equitable growth and investment in education and health). Combining these socio-economic settings with an energy system model produces the emissions projections seen in Figure 2; see Riahi et al. (2017) for a complete review of the SSP storylines and specifics on the underlying assumptions. This approach has been employed by the US Government in their computations of the SCC (National Center for Energy Economics, 2022),[6], and is our approach here. This implies that our optimal carbon taxes are always with reference to the emissions baseline we assume; we explore the influence of which emissions baseline we choose on our results in § 5.3.

We take emissions data for each SSP at times 2020 – 2100 directly from the SSP database,[7] and select scenarios which span a range of end-of-century radiative forcing amounts. We make one alteration to the projections provided in the database: negative emissions are set to zero.[8] As our model extends out to 2250, we require extensions of the SSPs in the database; we follow the prescription of Meinshausen et al. (2020) for each baseline, which assumes that (a) positive fossil fuel emissions and any net-negative fossil fuel emissions are ramped down to zero by 2250, (b) land use CO2 emissions are zero by 2150, (c) non-fossil fuel greenhouse gas emissions are ramped down by 2250, and (d) land use-related non-CO2 emissions are held constant after 2100. In reality, it is possible that in the absence of a well-designed policy suite that one or more of these assumptions could not hold, which would imply that we are underestimating potential emissions levels in the far-future, and thus long-term climate-economic risk; we explore the relative influence of which emissions baseline we choose in § 5.3. See Figure 2 for the results of our extension procedure.

[5]While one may question the coarseness of our time discretization, it has been shown that including more decision periods in similar models does not significantly affect their output (Coleman et al., 2021).

[6]Note the US EPA uses the so-called RFF-SPs (Rennert et al., 2022) rather than the SSPs used here.

[7]See https://tntcat.iiasa.ac.at/SspDb/dsd?Action=htmlpage&page=10

[8]This assumption only impacts SSP1–1.9, as SSP1–1.9 makes more optimistic assumptions around backstop technology than we do in our cost formulation.

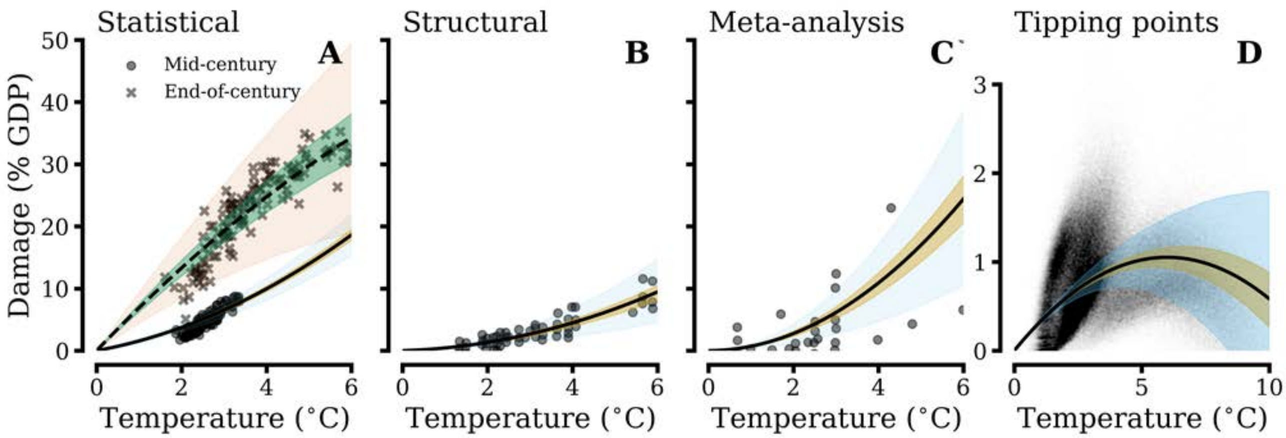

Figure 3: Each of our damage functions by methodology (statistical, structural, and meta-analytic) as well as the marginal damages owing to tipping points.

Note: In each panel, yellow shows ±1 standard deviation in the damage function, while the blue shaded region shows ±2 standard deviations. For the end-of-century estimates in panel A, the green region shown ±1 standard deviation and the salmon shows ±2 standard deviations. The statistical damage function shown assumes SSP2–4.5.

2.3 Damage functions

Our climate damage calculation can be broken down into two components: an *aggregate* climate damage, owing to the total damages incurred by climate change, and a marginal *tipping point* climate damage, which accounts for damages which are incurred by, for example, permafrost melt.

2.3.1 Aggregate climate change damages

Aggregate damages are defined as global damages owing to climate change, and their magnitude is estimated in AR6 by WGII (Intergovernmental Panel on Climate Change, 2022a) (see their Figure Cross-Working Group Box ECONOMIC.1, panels (a)-(c), p. 16-114). We specify three aggregate damage functions: one that is modeled after statistical climate damage modeling efforts (Burke et al., 2018), one estimated using structural estimation techniques (Rose et al., 2017), and a meta-analysis of climate damage estimates (Howard and Sterner, 2017), such that for each we have

$$\mathcal{D}(T') = T'(\varpi_1 + \varpi_2 T') \tag{2.6}$$

where $\varpi_1, \varpi_2 \in \mathbb{R}^+$ are fitted coefficients. We refer to each of these damage functions by their estimation methodology in what follows, i.e., "the statistical damage function" and so on. We supply the fitted coefficients and their uncertainty, as well as a discussion of the qualifications and the limitations of each individual damage function we use, in Online Appendix D. We present the data and fitted curves in Figure 3 (ft. [9]).

[9] We present CAP6 output using only one of each damage function, and compare it to when each damage function is sampled in Online Appendix H.

2.3.2 Tipping point damages

In addition to the aggregate damages accrued owing to climate change, an additional damage potential exists for climate-related tipping points, such as permafrost melt or Amazon dieback. Previous studies parameterize climate tipping points as instantaneous shocks that immediately result in damages (e.g., Lemoine and Traeger, 2016b); however, this is unrealistic, as the consequences of "hitting a tipping point" will take time to be fully realized (Kopp et al., 2016; Armstrong McKay et al., 2022). This effect was captured by Cai and Lontzek (2019); they found that the presence of climate tipping points significantly increases the social cost of carbon.

A recent analysis allows the effect of a given tipping element to be dynamic over time in an IAM, and estimates the marginal damage associated with ten climate tipping points as a function of global average temperature (Dietz et al., 2021). This approach has the advantage of aggregating over the complex dynamic aspects of tipping points and provides a simple "damage function" for marginal damages owing to tipping points. Moreover, this "damage function" implicitly captures the "domino" effect of hitting a tipping point (Lemoine and Traeger, 2016b; Cai et al., 2016) in its damage estimates. However, our use of this approach has the drawback of not capturing aversion to ambiguity surrounding the location of tipping points (Lemoine and Traeger, 2016a), which has been shown to slightly increase the stringency of climate policy. This provides some context to our results, as including the effects of ambiguity aversion to tipping points would increase the resulting carbon price and optimal mitigation level.

We take this additional "damage function" owing to tipping points, $\mathcal{D}_{tp}(T')$, from Dietz et al. (2021) (see their Figure 5c), such that the total damages are given by

$$\mathcal{D}_{tot}(T') = \mathcal{D}(T') + \mathcal{D}_{tp}(T'). \tag{2.7}$$

Note that $\mathcal{D}_{tp}(T')$ has the same functional form as the aggregate damage function, i.e., Eqn. (2.6). See Figure 3D for a visualization and Table 1 in Online Appendix D for the coefficients of this damage function and corresponding uncertainties.

2.3.3 Sampling damage function uncertainty

We sample uncertainty in the damage function in two ways. The first is by sampling the parametric uncertainty in each damage function; that is, the uncertainty in the values of ϖ_1, ϖ_2 in (2.6). The distributions of ϖ_1, ϖ_2 are assumed Gaussian with mean and variance provided in Online Appendix D, Table 1. The second source of uncertainty in the damage function pertains to which damage function (i.e., statistical, structural, or meta-analytic) we specify in the first place. As the IPCC WGII makes no recommendations in this regard, we assign a hyper-parameter in our simulated climate damages that randomly chooses a damage function, thus sampling epistemic uncertainty in the damage function. This methodology allows us to remain agnostic with respect to which damage function we choose.

2.3.4 Calculating damages at a particular decision node

A representative agent in our model at a given decision node only knows the possible end states which can be accessed from their state. They do not know the exact fragility at their own node, or any θ_t for $t < T$, owing to the inherent uncertainty surrounding both the climate system (such as the precise value of climate sensitivity) and economic impacts (such as damage functions). Owing to the agent not knowing the current fragility, the damages assessed at their decision time are dependent on proxies for the relevant damage variables. The two proxies used in our model is the set of possible end states, Θ (which tells us which end states are accessible) and the cumulative CO_2 emissions, Ψ_t (which tells us approximately how warm the world should be, but does not immediately map to the temperature at time t owing to uncertainty in the climate sensitivity). These two variables in concert give us a basis from which we can interpolate end state climate damages backwards in time to any decision node. Moreover, a continually-updating fragility parameter allows the expectation of future damages to co-evolve with agent decisions about mitigation, therefore making risk assessment endogenous within our modeling structure. We calculate the damage at a given node as a probability-weighted average of the current-period damages accessible to each end node across states of fragility, such that

$$\mathcal{D}_{node}(\Psi_t, \theta_t) = \sum_{\theta_T \in \Theta} P\left(\theta_T | \theta_t\right) \mathcal{D}_{tot}(\Psi_t, \theta_t). \tag{2.8}$$

2.4 Cost of mitigation

Calculating the cost of mitigation requires specifying a marginal abatement cost curve (MACC), which relates the price of abatement to the fraction of emissions abated. Such a curve will vary depending on three factors: (1) the current state of emissions mitigation technologies, which in aggregate represent the abatement potential as a function of cost, (2) the availability of a backstop technology, which allows for net-negative emissions, and (3) technological advancement, which makes mitigation costs cheaper over time (Gillingham and Stock, 2018). We discuss the limitations to our approach in Online Appendix E.

2.4.1 Marginal abatement cost curve estimation

Estimating MACCs requires a functional relationship between the fraction of emissions abated, x, the per-ton tax rate, τ, and the emission pathway, E. We use the most recent estimates for the cost of CO_2 emission abatement presented in AR6 WGIII (Intergovernmental Panel on Climate Change, 2022b) (see their Figure SPM.7, p. SPM-50). We make four important assumptions in interpreting the data from AR6 WGIII. First, we assume cost estimates are additive, which is not necessarily the case; however, we expect changes in costs and abatement potential to be small enough to consider them as negligible in this study. Second, we neglect negative costs; that is, whenever WGIII data dictates that costs are $< \$0$, we set the cost to zero. Third, for abatement potentials outside the range provided by the IPCC, we assume the functional relationship between τ and x established for lower abatement potentials holds. Lastly, we assume that the cost of each option is equal to its maximum cost in its respective range, i.e., the cost of an option in the IPCC \$0–\$20 range is assumed to be \$20. Taken

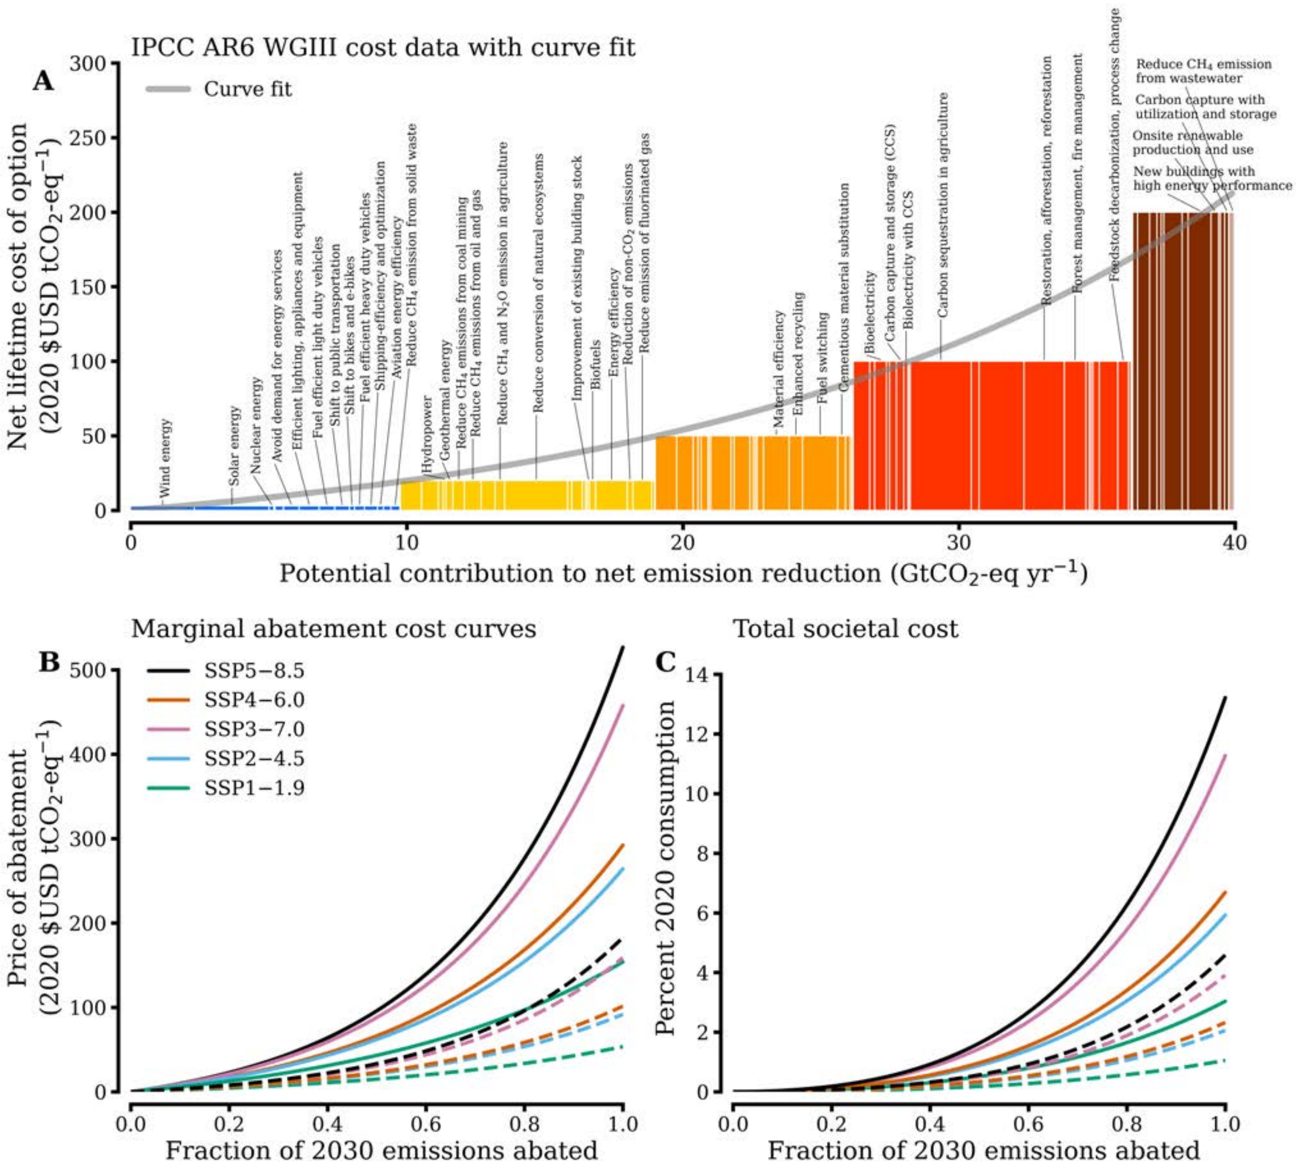

Figure 4: Panel A shows the mitigation potential and cost for each methodology given by the IPCC using their WGIII data. Blue represents zero costs (listed as negative in AR6), yellow is \$0-\$20 range, orange is \$20-\$50, red is \$50-\$100, and maroon is \$100-\$200. Panel B shows the fitted marginal abatement cost curves given by (2.9) and panel C shows the total cost to society given by (2.10) in our 'main specification'. In panels B–C, solid lines correspond to 2030 MACCs, while dashed lines are 2100 MACCs, assuming an exogenous technological growth rate of 1.5% and no endogenous technological growth.

Note: In panel A, the abatement methodology label is only on the bar with the most mitigation potential for a given methodology.

together, these assumptions make our MACC estimation conservative. We then fit an exponential curve to the cost data (see Figure 4A), such that

$$\tau(x) = \tau_0 \left(e^{\xi x} - 1\right), \tag{2.9}$$

where $\tau_0, \xi > 0$ are constants. To evaluate (2.3)–(2.5), we are interested in the total cost to society, $\kappa(\tau)$ for each particular tax rate τ, in units of the fraction of 2020 consumption lost. We use the envelope theorem to calculate $\kappa(\tau)$, such that (see Online Appendix E for the full derivation),

$$\kappa_{MACC}(x) = \frac{E_0 \tau_0}{c_{2020}} \left(\frac{e^{\xi x} - 1}{\xi} - x\right), \tag{2.10}$$

where c_{2020} is the 2020 global consumption in billions of 2020 USD, set to \$61880 (taken from the World Bank[10]) and E_0 is the emissions rate in 2030 in $GtCO_2$ yr^{-1}. A table of fitted values for τ_0 and ξ for each SSP are provided in Table 3 in Online Appendix E, as well as a calculation for the percent of consumption required to abate all emissions. Fits for (2.9) and (2.10) are shown in Figure 4B and Figure 4C, respectively.

2.4.2 Direct air capture technology

Our model represents direct air capture (DAC) via permitting CO_2 removal (National Research Council, 2015). Net CO_2 removal occurs whenever the mitigation exceeds unity; this leads to negative emissions and thus net carbon removal from the atmosphere. The price of net carbon removal is a major source of uncertainty in assessing future climate policy (Johnson et al., 2017), with estimates ranging from \$50 – \$1000 2020 USD per ton of CO_2 removed. Regardless of the specific dollar estimates provided in the literature, DAC faces a common hurdle: scalability (Intergovernmental Panel on Climate Change, 2022b). The parameter x in our MACC is the fraction of 2030 emissions abated; therefore, removing even a small percentage of these emissions from the atmosphere is equivalent to abating billions of tons of CO_2 from the atmosphere in short order. The technology to carry out this task is simply unavailable at present, and it is unclear when it will become fully mature and available at scale.

Note that, before mitigation reaches unity, there is some carbon capture and storage that is assumed to be occurring concurrent with emissions reductions; indeed, by considering the technology-by-technology breakdown of the IPCC's WGIII cost data in Figure 4A, carbon capture and storage is placed in the \$200 2020 USD $tCO_2\text{-eq}^{-1}$ cost bracket. Hence our inclusion of DAC in our MACC formulation represents an abrupt shift from purchasing various abatement technologies (such as solar power or equipment to retrofit buildings) to installing exclusively, and at scale, DAC facilities. The costs of this process are currently assumed to be rather large (International Energy Agency, 2022). However, a breakthrough could certainly occur sometime in the future where DAC becomes deployable at scale for a more economically viable cost (for example, as a result of the uncapped subsidies in the Inflation Reduction Act of 2022 (Yarmuth, 2022)), which would lower the price of DAC considerably and would require a reassessment of our quantitative analysis in § 5.

In light of these considerations, we take a simple approach to adjusting our cost curve to account

[10] https://data.worldbank.org/indicator/NE.CON.TOTL.CD

for DAC technologies by imposing a DAC premium, $\tau_{DAC} > 0$, which is an extra price for carbon removal which shifts τ_0 to $\tau_0 \to \tau_0 + \tau_{DAC}$. Throughout, we essentially price out to-scale DAC leading to net-negative emissions before 2100. This alters our MACC cost curve (2.10) when $x > 1$, such that

$$\kappa_{MACC}(x) = \begin{cases} \frac{E_0 \tau_0}{c_{2020}} \left(\frac{e^{\xi x} - 1}{\xi} - x \right), & 0 \le x \le 1, \\ \\ \frac{E_0 (\tau_0 + \tau_{DAC})}{c_{2020}} \left(\frac{e^{\xi x} - 1}{\xi} - x \right), & x > 1. \end{cases} \tag{2.11}$$

2.4.3 Technological progress

Technological progress in CAP6 is captured by allowing the cost of mitigation to society $\kappa_{MACC}(x)$ to decrease in time as technological proficiency makes mitigation cheaper. Technological progress can occur in two ways: (1) exogenously, where general technological improvement independent of agent choices make mitigation cheaper, and (2) endogenously, where if a given individual invests in mitigation early, the cost of mitigation goes down more over time (Acemoglu et al., 2012). The exogenous (endogenous, resp.) technology advancement rate is given by $\varphi_0 \ge 0$ ($\varphi_1 \ge 0$, resp.). Incorporating these factors into our cost curve results in our final expression for the cost of mitigation to society,

$$\kappa_t(x_t) = \kappa_{MACC}(x_t) \left(1 - \varphi_0 - \varphi_1 X_t\right)^{t-10}, \tag{2.12}$$

where

$$X_t := \frac{\int_0^t x(\zeta) E(\zeta) d\zeta}{\Psi(t)}, \tag{2.13}$$

is the weighted average mitigation up to time t (ft. [11]).

We note that our formulation of endogenous technological change – or "learning by doing" – follows a formulation akin to Wright's law (Wright, 1936), where the reduction in costs of mitigation technologies is proportional to the total deployed mitigation, as opposed to directed technical change in the spirit of Acemoglu et al. (2012) or Lans Bovenberg and Smulders (1995). This is because in our formulation, the social planner chooses levels of abatement, which (via proxy) corresponds to the deployment of clean technologies. As more and more renewable technologies are "deployed" by the planner, Wright's law would suggest that their costs will decrease. Hence the Wright's law-based formulation is the most natural way to incorporate endogenous technological change into our model. This framework has the additional advantage of allowing us to only focus on carbon tax levels rather than including additional policy instruments, such as renewable energy subsidies.

2.4.4 "No free lunches" calibration

Estimating the cost of CO_2 abatement is notoriously challenging. The cost estimates presented above are static, in the sense that they represent the costs of the lifetime of the project and, for example, ignore

[11] Note the technological growth factor is offset by ten years as the cost data from AR6 is for 2030 technologies and our first model period is in 2020.

spillover effects (Intergovernmental Panel on Climate Change, 2022b). However, static estimates fail to capture the impact of the costs (or savings) associated with a given project that outlive the project lifetime itself (Gillingham and Stock, 2018). Such considerations lead some to argue that costs should not be estimated from the "bottom up" as done here, but rather from the "top down." "Top down" estimates generally paint a more pessimistic picture than the "bottom up" methods, positing that the cost of abating CO_2 emissions is actually larger than adding up the cost of each individual option, owing to inertia and friction in the economic system, a set of barriers typically summarized as the "energy paradox" (Jaffe and Stavins, 1994).

To address this concern, we provide an alternative calibration of CAP6 that is more closely aligned to "top-down" MACCs (see, e.g., Barrage and Nordhaus, 2023) by adjusting the MACC to exclude zero-cost abatement technologies; indeed, it has been shown that the degree to which one believes in zero-cost mitigation explains much of the difference between "top-down" and "bottom-up" MACC estimates (Kotchen et al., 2023). We do so by shifting all of the mitigation potential in the IPCC dataset up by one cost bracket; for example, the zero cost methodologies (the blue bars in Figure 4) now have \$20 2020 USD $tCO_2\text{-eq}^{-1}$ lifetime cost, and so on. The highest cost abatement technologies are set to cost \$400 2020 USD $tCO_2\text{-eq}^{-1}$. We coin this MACC calibration as the "no free lunches" MACC, and provide its parameter values in Online Appendix E, Table 3. (ft. [12]).

3 Climate model

Here we present the climate component of our model. We map CO_2 emissions to the temperature anomaly above preindustrial levels, denoted as T', using the transient climate response to emissions (TCRE) (Damon Matthews et al., 2021). The TCRE is defined as a linear scale factor $\lambda > 0$ that maps the cumulative CO_2 emissions, $\Psi(t) := \int_0^t E(\zeta)d\zeta$, to temperature, where $E(t)$ is the emissions baseline. The physical basis for TCRE is a compensation between the diminishing sensitivity of radiative forcing to CO_2 at higher atmospheric concentration and the diminishing ability of the ocean to take up heat and carbon at higher cumulative emissions (Intergovernmental Panel on Climate Change, 2021). We follow the framework laid out in Damon Matthews et al. (2021) to use a TCRE that accounts for non-CO_2 forcing via the parameter $f_{nc} > 0$ which increases the average value and variance of the TCRE. We write our "effective" TCRE – the TCRE including non-CO_2 forcing factors – as

$$\lambda_{eff} := \frac{\lambda}{1 - f_{nc}}. \tag{3.1}$$

The mean value of λ, f_{nc}, and λ_{eff} and their uncertainties are provided in Online Appendix F, Table 4. Using this approach, we are able to reproduce central estimates of warming levels this century reported by WGI in AR6 for each SSP reasonably well, see Online Appendix F, Table 5. Therefore in our

[12] The analogous figure to Figure 4 for the "no free lunches" MACC is provided in Online Appendix E. We also performed a second recalibration that sets the costs of the the $<$ \$0 mitigation options to infinity, coined the "infinite cost" calibration. The figure associated with this calibration is also in Online Appendix E. We do not show the results of CAP6 with this calibration as the final costs of abatement are lower than in the "no free lunches" case, but higher than the 'main specification.' Hence, the results will simply be an interpolation between the main specification and the "no free lunches" results.

calculations of temperature, we use

$$T'(t) = \lambda_{eff}\Psi(t). \tag{3.2}$$

The TCRE approach has a number of advantages: (i) it captures short- and long-term uncertainty in climate warming, (ii) it is relatively simple and transparent, and (iii) emulates state-of-the-art climate models well (Allen et al., 2009; Dvorak et al., 2022).[13] Moreover, the TCRE framework has been used in a number of other climate-economic models (e.g., Dietz and Venmans, 2019; Campiglio et al., 2022).

4 Model calibration

4.1 Featured runs

To calibrate CAP6, we use discount rates in line with recommendations from the US government. Previous analyses use a discount rate of 3% (Committee on Assessing Approaches to Updating the Social Cost of Carbon et al., 2017), but recent studies use 2% in light of recent economic trends (such as falling interest rates) and expert elicitation (Council of Economic Advisors, 2017; Drupp et al., 2018). Indeed, New York State adopted a 2% discount rate in their social cost of carbon calculations (New York State Energy Research and Development Authority and Resources for the Future, 2020). We calibrate our featured runs using 1.5%, 2% and 2.5% discount rates to be consistent with the recent report issued by the EPA (National Center for Energy Economics, 2022) and use the term structures from Bauer and Rudebusch (2020). We also show results using a 3% discount rate for consistency with prior US government estimates (Committee on Assessing Approaches to Updating the Social Cost of Carbon et al., 2017).[14] See Online Appendix G, Table 6 for specifics. We assume $g = 1.5\%$ for all runs. For each discount rate, we assume that $\psi = 10$, in line with trends observed in the U.S. financial market (Schroyen and Aarbu, 2017). For our emissions baseline, we choose SSP2–4.5, as it aligns with recent projections of emissions used by the US EPA (Rennert et al., 2022). Lastly, we assume a modest exogenous technological growth rate of 1.5% and no endogenous technological growth, owing to an inability to reliability calibrate the endogenous technological growth rate parameter φ_1. The choice of no endogenous technological growth makes our technological growth assumptions conservative, given the known link between agent investment in mitigation and rates of growth in clean sectors (Acemoglu et al., 2012).[15]

4.2 Ensemble runs

While risk associated with temperature rise and damage function uncertainty are holistically evaluated in a given run of CAP6, other sources of uncertainty exist and are excluded, such as uncertainty in

[13] Dvorak et al. (2022) showed that the TCRE adequately emulates the response of the more comprehensive FaIR model (Smith et al., 2018), itself a combination of carbon cycle models (Joos et al., 2013) and physical response models (Geoffroy et al., 2013b,a). The TCRE can deviate from more sophisticated models slightly depending on the forcing scenario (Intergovernmental Panel on Climate Change, 2021), but the differences are minor and are therefore ignored in this study.

[14] We do not here take a stand on which discount rate is correct, but do consider the 2% rate as our benchmark, as it is the central rate used by the EPA.

[15] We demonstrate how including endogenous technological growth influences model output in Online Appendix K.

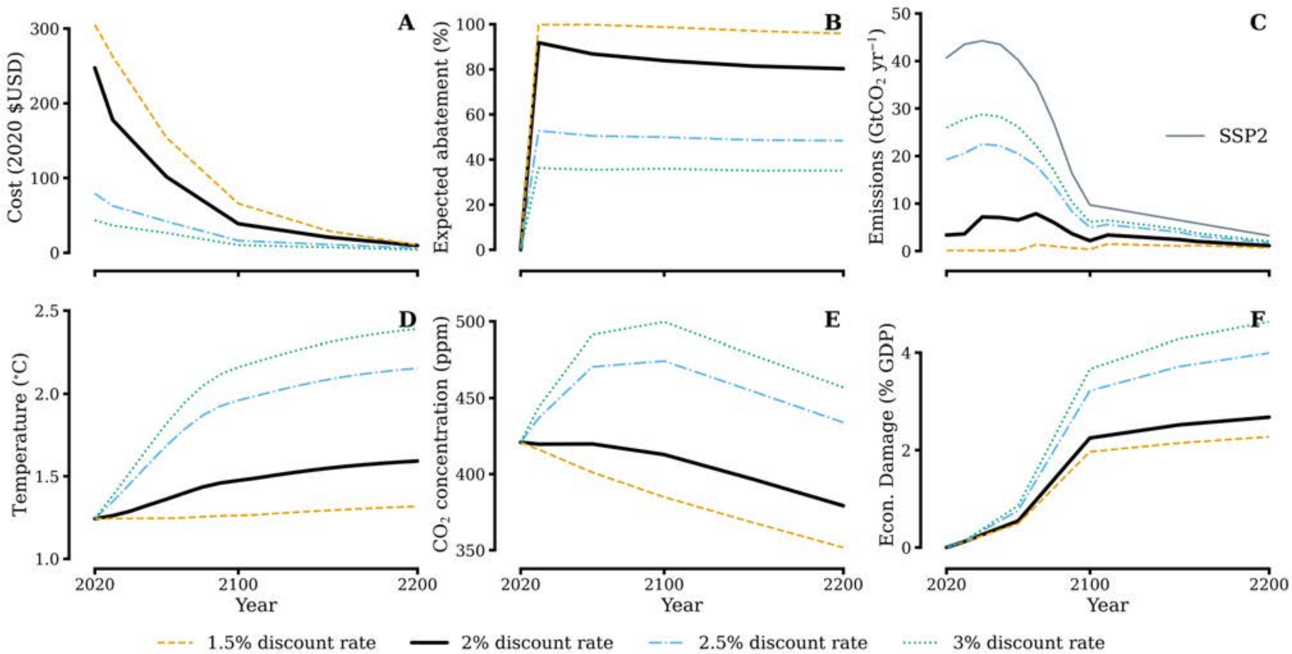

Figure 5: CAP6 output for four discount rates in our main specification.

the rate of technological growth, or which exogenous emissions baseline or discount rate is assumed. Each of these represent a source of epistemic uncertainty in the climate-economic system; indeed, not knowing how much CO_2 will be emitted over the next century, for example, strongly influences the range of possible climate realizations, and thus, climate-related risk (Hawkins and Sutton, 2009; Lehner et al., 2020). To probe the impact of assumptions associated with each of these parameters on model output, we carry out a Monte Carlo analysis. We sample discount rates between the range of 1.5% and 4.25%; we chose the lower bound based on the lower bound considered by the EPA and the upper bound is the preferred rate used in DICE–2016R (Nordhaus, 2017). The value of agent RA has been measured to as high as 15 in wealthy countries and as low as 3 in some European nations (Schroyen and Aarbu, 2017), which defines our range. We choose the modest ranges of 0%–3% for both the exogenous and endogenous rate of technological growth. Note that we use our 'main specification' MACC for the ensemble run analysis. See Online Appendix G, Table 7 for our numerical values.

5 Results

5.1 Main specification

We show the featured model runs of CAP6 in Figure 5. We find that the 2% discount rate policy implies a high cost of carbon and stringent abatement policies, see panels 5A–B. The cost of carbon declines over time; this, however, should not be confused with reduced abatement action over time. Rather, the declining dynamics of carbon prices can be entirely attributed to the improved ability to abate CO_2 emissions owing to technological improvements (see Eqn. (2.12)). This set of mitigation actions leads emissions peaking in 2070, with CO_2 concentrations stabilizing before starting to decrease by mid-century. The expected global temperature change resulting from this emissions policy is less than 1.5 °C by 2100 (∼ 1.47 °C) and less than 2 °C in 2200 (∼ 1.6 °C).

Decreasing the discount rate to 1.5% leads to complete and immediate cessation of emissions (see panel 5B), thus maximizing costs and decreasing 2100 (2200, resp.) warming by 0.2 °C (0.3 °C, resp.) in comparison to the 2% run. Larger discount rates relax the stringent abatement policies seen in the 2% and 1.5% discount rate cases. This results in lower costs and less mitigation action, and consequentially, larger warming and damages. We find that both the 2.5% and 3% discount rates warm beyond the warming target of 1.5 °C by 2100 established in the Paris Agreement. Moreover, the 3% discount rate policy exceeds 2 °C warming by 2100, and the 2.5% discount rate policy barely holds temperatures below 2 °C by 2100 (∼ 1.96 °C by 2100). In the case of the 2.5% and 3% discount rates, CO_2 concentrations rise before falling as emissions cease.[16] The 2.5% (3%, resp.) discount rate individual also tends to lose ∼1% (∼1.4%, resp.) more GDP in 2100 and ∼1.3% (∼2%, resp.) more in 2200 than in the 2% discount rate case, showing the expensive consequences of delayed action in combating climate change.

The intuition behind our declining carbon prices can be found in our structural representation of risk. In the early periods of the model, the social planner faces the risk of catastrophic long-term damages if they choose not to abate any CO_2 emissions (∼50% GDP or higher, if the worst-case climate sensitivity and damage function concurrently materialize); this causes the social planner to mitigate aggressively early on to effectively rule out such catastrophic futures from ever materializing. Technological progress then brings down abatement costs over time (especially if learning-by-doing effects are considered, see Online Appendix K), and drives down the carbon price over time. These two factors combine to cause carbon prices to start high and decline over time.

From this analysis, we find that modeling the cost of climate risk with CAP6 supports stringent mitigation action. We find that the carbon price and corresponding mitigation policy associated with the 2% discount rate saves at least $22 trillion 2020 USD globally in 2100 (assuming global GDP grows annually by 4%) in comparison to the higher discount rate policies. In addition, employing policies with discount rates considered by the EPA result in an expected warming level in line with the targets set forth in the Paris agreement (United Nations Framework Convention on Climate Change, 2015), providing the targets with explicit economic support. When faced with potentially severe damages, the representative agent makes a clear choice: they sacrifice consumption today to abate CO_2 emissions, consistent with our understanding of how risk influences climate mitigation policy.

5.2 Alternative calibration: "no free lunches"

We recalculate our featured runs using the "no free lunches" MACC and show the results in Figure 6. The "no free lunches" cost curve leads to an increase in the optimal price of carbon; the 2020 CO_2 price increases by 20% in the 2% discount rate case. However, the "no free lunches" MACC significantly influences the efficacy of the optimal price in abating CO_2 emissions. For example, the 2% discount rate policy now abates only 70% of emissions (as opposed to ∼ 85% in the main specification). This emissions pathway reaches ∼ 1.7 °C of warming by 2100 and ∼ 1.9 °C warming by 2200, notably

[16] We use the carbon cycle model of Joos et al. (2013) to compute carbon concentrations for our optimal mitigation pathways, see Online Appendix F.

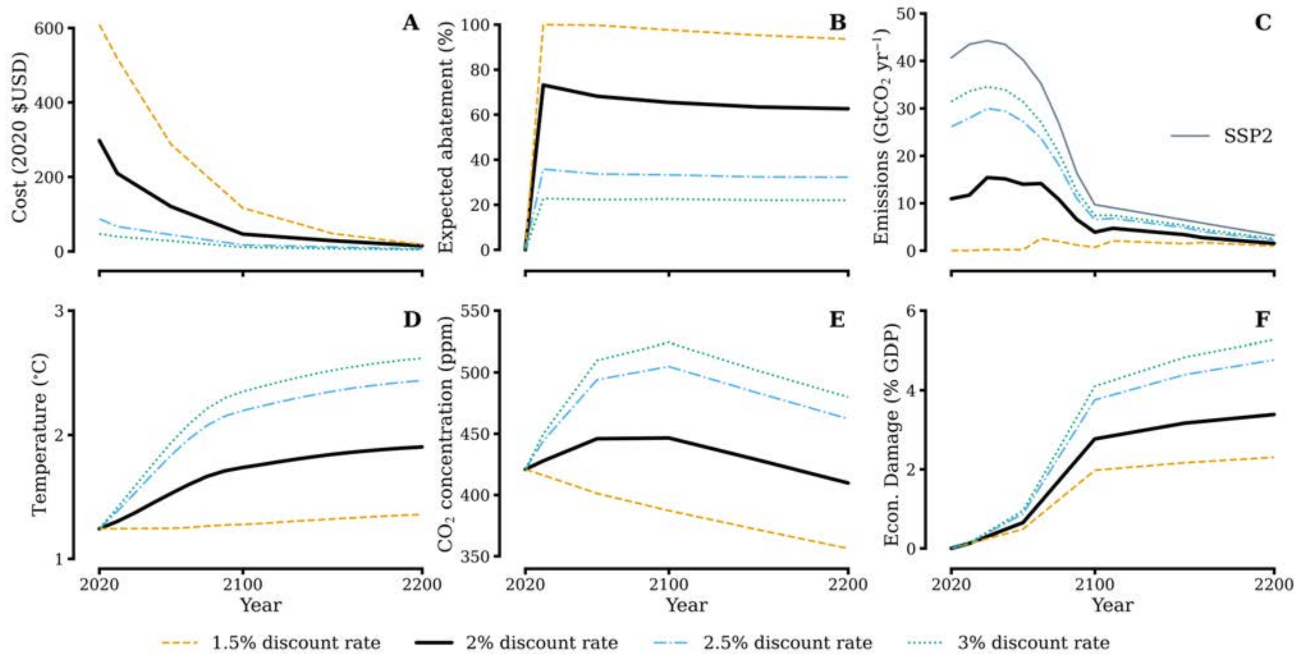

Figure 6: CAP6 output for four discount rates using the "no free lunches" cost curve calibration.

maintaining less than 2 °C warming. This shows that even if the cost of abatement is considerably higher than the IPCC foretells, keeping total warming below 2 °C is still optimal within CAP6 when a 2% discount rate is used.

Running CAP6 with the "no free lunches" calibration and a 2.5% or 3% discount rates show similar results as the 2% rate, with higher optimal prices, more near-term warming, and higher CO_2 concentrations. In this case, however, we find that using a 2.5% or 3% discount rate exceeds 2 °C warming in 2100, thus exceeding the upper bound of targeted warming in the Paris agreement. This shows that if abatement turns out to be more costly than we expect, using a higher discount rate in climate policy makes the world's ability of achieving the warming targets in the Paris agreement far more tenuous.

The only exception to the pattern above – the "no free lunches" MACC leading to less abatement and more warming – is the 1.5% discount rate policy, which still abates nearly 100% of emissions in the near term. This can be explained by this agent having both a low discount rate and low risk tolerance, and therefore sacrifices considerable consumption to minimize both experienced and potential future damages owing to climate change.

5.3 Ensemble model analysis

We probe the influence of uncertainty in exogenous model parameters on CO_2 price paths, temperature change, CO_2 concentrations, and economic damages incurred in our ensemble runs, shown in Figure 7. We find that CO_2 price paths decline over time, regardless of socio-economic specification, owing to agent risk response and technological progress. The level of CO_2 price varies between baselines because the MACC is baseline dependent (see Eqn. (2.12)); for the same fraction of emissions abated, agents pay different prices depending on the baseline. Finally, cost variance is highly stratified across baselines,

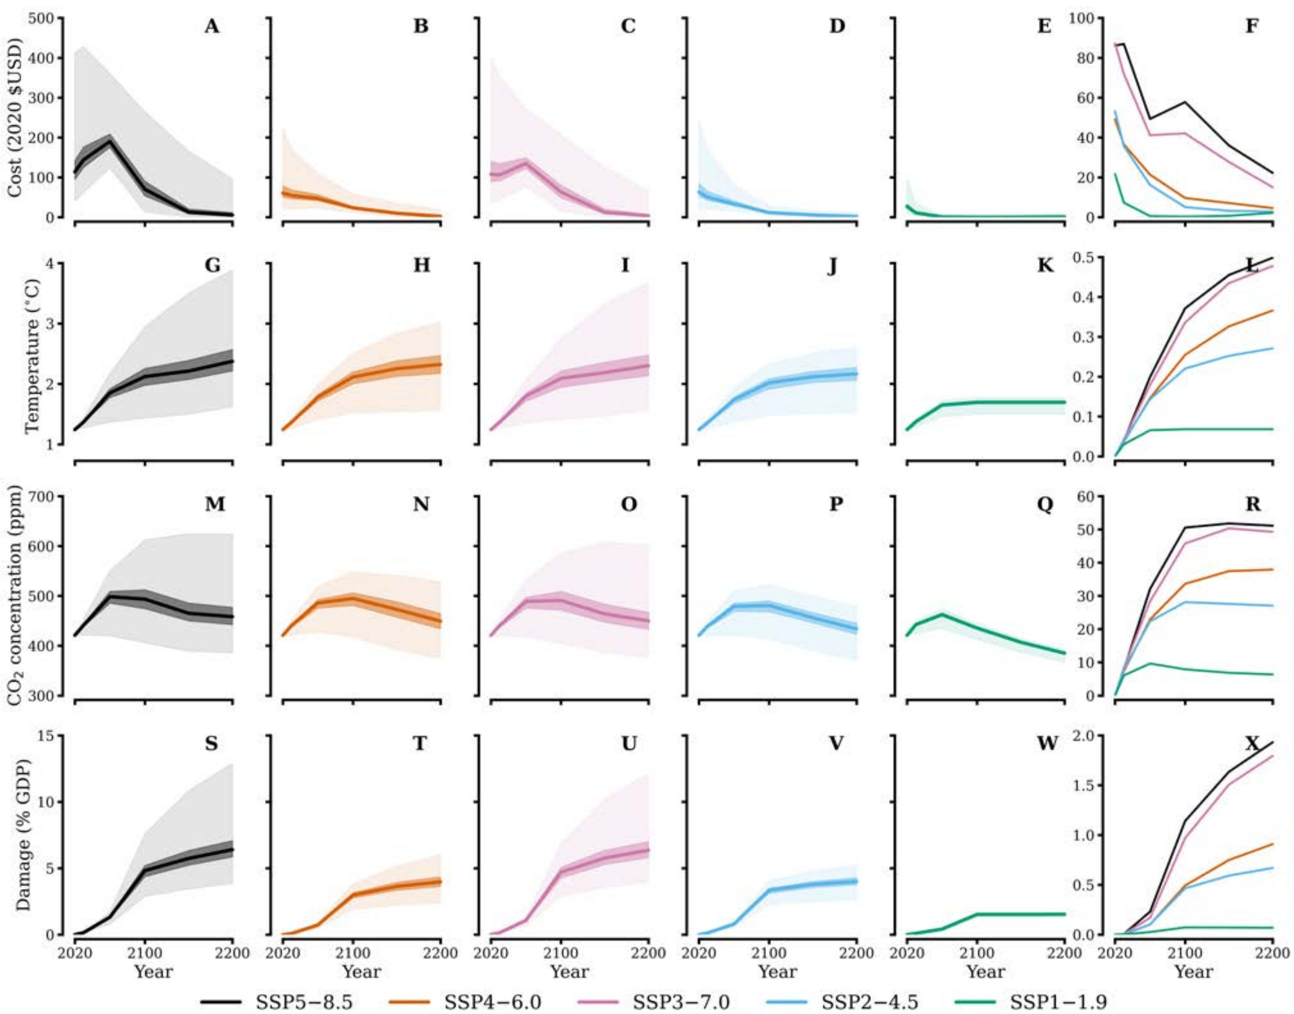

Figure 7: Cost (top row, panels A–E), temperature (second from top, panels G–K), CO_2 concentrations (third from top, panels M–Q), and economic damages (bottom row, panels S–W) from our ensemble model runs. Dark (light, resp.) shaded region represents the 36^{th}–64^{th} (1^{st}–99^{th}, resp.) percentile range, solid lines represent the median time series. In the final column (panels F, L, R, and X) we plot the standard deviation of each parameter distribution in time.

see panel 7F.

Central estimates of temperature, CO_2 concentrations, and economic damages,[17] however, do not see significant differences in central estimates across baselines as was observed in CO_2 prices. This owes to suggested policy in CAP6 being consistent across baselines; the only difference is the price of implementing said policy. Hence, the impact variables are relatively insensitive to baseline choice. This is a notable result, as it implies CAP6 finds an optimal outcome across emissions baselines for a given calibration. The variance in each impact variable however, displayed in panels 7L,R,X, is sensitive to the choice in baseline, with high (low, resp.) emissions scenarios having the highest (lowest, resp.) amount of variance. This can be explained by considering the consequences of inaction (i.e., high discount rate policies). In a high emissions scenario such as SSP5–8.5, inaction leads to more emissions, and thus higher impacts than in a low emissions scenario such as SSP2–4.5. Hence, the variance in each impact variable are all higher for high emissions scenarios than in low emissions scenarios.

5.3.1 Variance decomposition of ensemble results

The significant stratification of uncertainty in our output variables shown in Figure 7 motivates further study; is it high discount rates that control prices, for example, or rates of technological change? To this end, we perform a regression analysis of CO_2 price and the impact variables studied above at every point in time against parameter values, and plot the fraction of total r^2 attributable to each parameter in Figure 8 (see Online Appendix I for details and supporting figures).

For prices, we find that the discount rate (i.e., EIS and PRTP) dominate uncertainty in the near term (i.e., prior to 2100). This owes to these parameters dictating individual attitudes towards time-related risk and discounting. In early periods of the model, climate damages are highly uncertain. Therefore, any abatement action that is taken is with the intent to rule out the most catastrophic outcomes and secure future welfare; the extent to which individuals respond to this threat of catastrophe is governed by the discount rate, thus determining the level of early mitigation action and driving costs. On longer timescales (past 2100), climate damages have been more distinctly realized, and the number of possible futures have narrowed. Individuals must come to grips with their damaged future, and generally begin investing more stringently in emissions abatement. This comes at a cost, a cost that is determined by how much cheaper abatement technologies have become in the time it took to reach this decision. In particular, high prices in late periods are almost entirely attributable to low rates of technological change across SSPs.

For the impact variables, however, a different story emerges: the influence of the discount rate is pronounced for much longer than in the case of CO_2 prices. This owes to inactivity early on leading to long-term consequences in the form of climate-economic impacts that cannot simply be fixed by more spending on abatement.[18] Indeed, while technological change can certainly halt any further increase in global mean surface temperature, for example, it cannot undo past malfeasance.[19] Hence, the discount

[17] We refer to this set of variables as "impact variables" for the remainder of this discussion.

[18] This conclusion relies on a high cost of net-negative emissions; if a breakthrough in direct air capture (DAC) technologies occurs, then we would expect the variance explained in impact variables owing to technological growth to be higher, as net-negative emissions would enable long-run temperatures, CO_2 concentrations, and economic losses to be changed, perhaps significantly so, depending on how expensive DAC turns out to be.

[19] An important qualification to this conclusion is that we do not consider solar geoengineering, which could lead to

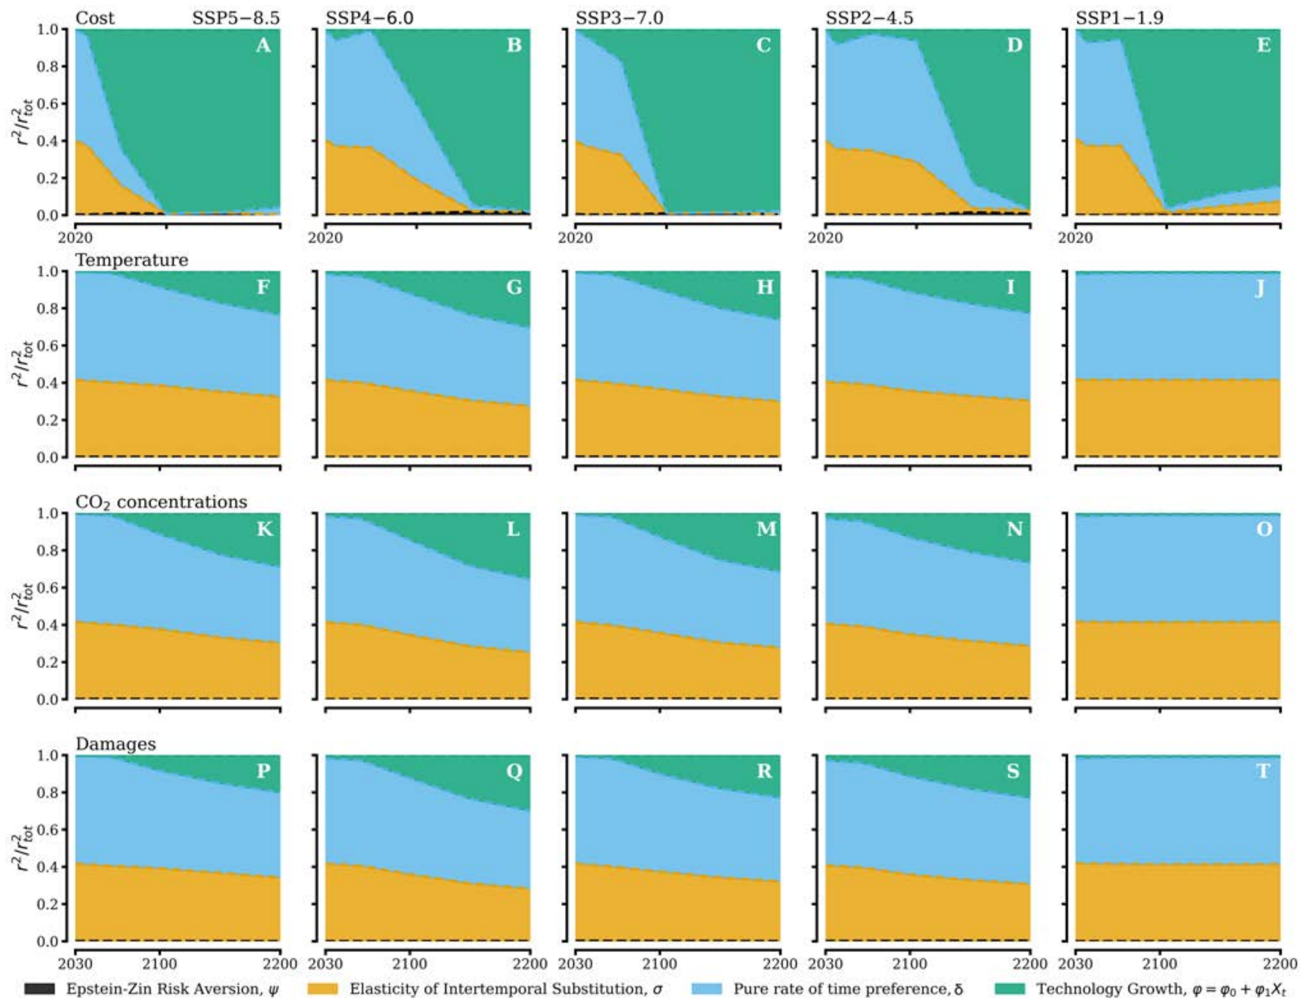

Figure 8: Fraction of total variance (calculated as total r^2) attributable to each model parameter for carbon prices (top row, panels A–E), temperature (second row, panels F–J), CO_2 concentrations (third row, panels K–O), and economic damages (bottom row, panels P–T). Each column represents a different SSP.

Note: Cost variance (top row) begins in 2020 whereas temperature, CO_2 concentrations, and economic damages (bottom three rows) begin in 2030, as the model is initialized with the same climate conditions and no damages incurred, leading to zero variance in 2020 for the latter three variables.

rate has a much more pronounced influence on far-distant temperature rise, atmospheric CO_2 levels, and economic damages than in the case of CO_2 prices.

Interestingly, Figure 8 shows that the influence of RA (i.e., the value of ψ) is suppressed for CAP6 output uncertainty[20] relative to other model inputs. We postulate that this owes to the risk aversion captured by ψ (i.e., the Epstein-Weil-Zin sense of risk aversion across states of nature) is relatively less important to risk across states of time. Given the large residence time of CO_2 in the atmosphere, it stands to reason that the impact of risk aversion with respect to time would dwarf the impact of risk aversion across states of nature. Indeed, the results of Figure 8 provide resounding support for this theory: risk aversion across states of time (captured by EIS) drowns out the influence of risk across states of nature (as captured by RA).

6 Conclusion

Over a decade ago, Lord Nicholas Stern wrote that "Presenting the [climate] problem as risk-management is likely to point strongly towards a policy for a rapid transition to a low-carbon economy" (Stern, 2013).[21] Our framework takes this view seriously, and, in the final analysis, shows the wisdom in Stern's words. By treating CO_2 as a risky asset and calculating the optimal CO_2 price and associated abatement policy using U.S. EPA-consistent discount rates, we find that optimal policy limits warming below 2 °C in 2100 for each discount rate we considered. Practically speaking, this corresponds to cutting > 70% of CO_2 emissions in relatively short order; a "rapid transition to a low-carbon economy" indeed. Our results flip the conventional view of climate policy on its head; rather than abating progressively more CO_2 emissions as time goes on (and damages are felt more acutely), our model suggests stringent early abatement as a 'hedge' against potentially severe damages associated with climate change.

Evidently our framework for computing optimal climate policies is idealized, and in practice, a number of additional considerations are necessary for formulating robust climate policy. For example, we compute a globally "optimal carbon tax" as a proxy for the overall strength of climate policy, not as an actual policy guide.[22] Prospects for such a common global carbon tax are bleak, to put it mildly. Therefore, useful extensions of this work would analyze the transition risk towards zero emissions policies, i.e., by considering asset stranding and adjustment costs (Campiglio et al., 2022), the potential for a 'run on fossil fuels' induced by an expected transition away from fossil fuel use (Barnett, 2023), or considering the distributional effects of heterogeneous climate policy mixes in different nations (as explored in Clausing and Wolfram, 2023). More work in this direction could prove both scientifically and economically insightful as well as immediately applicable in a wide variety of policy settings.

increased spending influencing temperature, CO_2 concentration, and economic damages levels in both the short- and long-term.

[20]This is not to say that RA has no impact on price levels, as increasing (decreasing, resp.) RA does slightly raise (lower, resp.) near term prices, see Online Appendix J.

[21]Others, like Nordhaus (2007), criticized Stern at the time, while Weitzman (2007) argued that Stern was "right for the wrong reasons", reasons subsequently developed in Weitzman (2009, 2012).

[22]Another limitation is that we compute the optimal carbon tax with a single exogenous discount rate. In reality, the discount rate will respond to the level of risk (Lucas, 1976) and is uncertain on long time horizons (Weitzman, 1998). Allowing for a dynamic discount rate in our framework is a potentially fruitful avenue of future work.

Author ORCIDs

- Adam Michael Bauer: 0000-0002-7471-8934
- Cristian Proistosescu: 0000-0002-1717-124X
- Gernot Wagner: 0000-0001-6059-0688

Acknowledgements

The authors thank Simon Dietz, Romain Fillon, Bob Kopp, Geoffrey Heal, Glenn Hubbard, Bob Litterman, Bruce McCarl, James Rising, Chris Smith, Thomas Stoerk, Andrew Wilson, three anonymous reviewers, and seminar audiences at Columbia University, the Center for Social and Environmental Futures at the University of Colorado Boulder, the 2022 fall meeting of the American Geophysical Union, and the 2023 summer conferences of the Association of Environmental and Resource Economists and the European Association of Environmental and Resource Economists for providing useful feedback on this work. The authors thank Alaa Al Khourdajie for providing the data from AR6 WGIII's cost of mitigation figure, W. Matthew Alampay Davis and Steve Rose for helpful discussions of climate damage functions, David C. Lafferty for his contributions to Figure 3, and Jaydeep Pillai for testing the public release version of the code. AMB thanks Columbia Business School for their hospitality while this work was being completed, and acknowledges support from the Gies College of Business Office of Risk Management and Insurance at the University of Illinois Urbana-Champaign, Columbia Business School, and a National Science Foundation Graduate Research Fellowship grant No. DGE 21-46756. CP was supported by the Gies College of Business Office of Risk Management and Insurance Research at the University of Illinois Urbana-Champaign. Computations were performed on the Keeling computing cluster, a computing resource operated by the School of Earth, Society and the Environment (SESE) at the University of Illinois Urbana-Champaign.

Disclosure statements

Adam Michael Bauer

At the time of submission, I hold a short-term consultancy position at the World Bank's Climate Change Group; the work being completed under this consultancy is unrelated to this publication (hence why I do not declare an affiliation to the World Bank Group in the above). I have no other conflicts of interest to disclose.

Cristian Proistosescu

I have no conflicts of interest to disclose.

Gernot Wagner

I am on the corporate advisory board of CarbonPlan. I have no other conflicts of interest to disclose.

Author contributions

Cristian Proistosescu and Gernot Wagner conceived of the study. Adam Michael Bauer wrote the code, designed numerical experiments, performed literature review, and made the figures. The first draft of the paper was written by Adam Michael Bauer, and all authors assisted in editing this draft to shape the final submitted manuscript. All authors have approved the submitted verison.

Data Availability

The code for the Carbon Asset Pricing model – AR6 (CAP6) can be found at the following Github repository: github.com/adam-bauer-34/cap6.

References

D. Acemoglu, P. Aghion, L. Bursztyn, and D. Hemous. The Environment and Directed Technical Change. *American Economic Review*, 102(1):131–166, Feb. 2012. ISSN 0002-8282. doi: 10.1257/aer.102.1.131. URL https://pubs.aeaweb.org/doi/10.1257/aer.102.1.131.

M. R. Allen, D. J. Frame, C. Huntingford, C. D. Jones, J. A. Lowe, M. Meinshausen, and N. Meinshausen. Warming caused by cumulative carbon emissions towards the trillionth tonne. *Nature*, 458(7242):1163–1166, Apr. 2009. ISSN 0028-0836, 1476-4687. doi: 10.1038/nature08019. URL http://www.nature.com/articles/nature08019.

D. I. Armstrong McKay, A. Staal, J. F. Abrams, R. Winkelmann, B. Sakschewski, S. Loriani, I. Fetzer, S. E. Cornell, J. Rockström, and T. M. Lenton. Exceeding 1.5°C global warming could trigger multiple climate tipping points. *Science*, 377(6611):eabn7950, Sept. 2022. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.abn7950. URL https://www.science.org/doi/10.1126/science.abn7950.

M. Barnett. A Run on Oil? Climate Policy and Stranded Assets Risk. Working Paper, 2023. URL https://www.ssrn.com/abstract=4346525.

M. Barnett, W. Brock, and L. P. Hansen. Pricing Uncertainty Induced by Climate Change. *The Review of Financial Studies*, 33(3):1024–1066, Mar. 2020. ISSN 0893-9454, 1465-7368. doi: 10.1093/rfs/hhz144. URL https://academic.oup.com/rfs/article/33/3/1024/5735312.

L. Barrage and W. Nordhaus. Policies, Projections, and the Social Cost of Carbon: Results from the DICE-2023 Model. Technical Report w31112, National Bureau of Economic Research, Cambridge, MA, Apr. 2023. URL http://www.nber.org/papers/w31112.pdf.

M. D. Bauer and G. D. Rudebusch. The Rising Cost of Climate Change: Evidence from the Bond Market. *Federal Reserve Bank of San Francisco, Working Paper Series*, pages 1.000–39.000, July 2020. doi: 10.24148/wp2020-25. URL https://www.frbsf.org/economic-research/publications/working-papers/2020/25/.

L. Bretschger and A. Vinogradova. Growth and Mitigation Policies with Uncertain Climate Damage. Working Paper 5085, CESifo, Munich, 2014. URL https://www.cesifo.org/en/publications/2014/working-paper/growth-and-mitigation-policies-uncertain-climate-damage.

M. Burke, W. M. Davis, and N. S. Diffenbaugh. Large potential reduction in economic damages under UN mitigation targets. *Nature*, 557(7706):549–553, May 2018. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-018-0071-9. URL http://www.nature.com/articles/s41586-018-0071-9.

Y. Cai and T. S. Lontzek. The Social Cost of Carbon with Economic and Climate Risks. *Journal of Political Economy*, 127(6):2684–2734, Dec. 2019. ISSN 0022-3808, 1537-534X. doi: 10.1086/701890. URL https://www.journals.uchicago.edu/doi/10.1086/701890.

Y. Cai, T. M. Lenton, and T. S. Lontzek. Risk of multiple interacting tipping points should encourage rapid CO2 emission reduction. *Nature Climate Change*, 6(5):520–525, May 2016. ISSN 1758-678X, 1758-6798. doi: 10.1038/nclimate2964. URL http://www.nature.com/articles/nclimate2964.

E. Campiglio, S. Dietz, and F. Venmans. Optimal Climate Policy as If the Transition Matters. Working Paper 10139, CESifo, Munich, 2022. URL https://www.cesifo.org/en/publications/2022/working-paper/optimal-climate-policy-if-transition-matters.

K. A. Clausing and C. Wolfram. Carbon Border Adjustments, Climate Clubs, and Subsidy Races When Climate Policies Vary. *Journal of Economic Perspectives*, 37(3):137–162, Sept. 2023. ISSN 0895-3309. doi: 10.1257/jep.37.3.137. URL https://www.aeaweb.org/articles?id=10.1257/jep.37.3.137.

T. F. Coleman, N. S. Dumont, W. Li, W. Liu, and A. Rubtsov. Optimal Pricing of Climate Risk. *Computational Economics*, Aug. 2021. ISSN 0927-7099, 1572-9974. doi: 10.1007/s10614-021-10179-6. URL https://link.springer.com/10.1007/s10614-021-10179-6.

Committee on Assessing Approaches to Updating the Social Cost of Carbon, Board on Environmental Change and Society, Division of Behavioral and Social Sciences and Education, and National Academies of Sciences, Engineering, and Medicine. *Valuing Climate Changes: Updating Estimation of the Social Cost of Carbon Dioxide*. National Academies Press, Washington, D.C., 2017. ISBN 9780309454209. doi: 10.17226/24651. URL https://www.nap.edu/catalog/24651.

Council of Economic Advisors. Discounting for public policy: Theory and recent evidence on the merits of updating the discount rate. Technical report, Washington, DC, 2017. URL https://obamawhitehouse.archives.gov/sites/default/files/page/files/201701_cea_discounting_issue_brief.pdf.

J. C. Cox, S. A. Ross, and M. Rubinstein. Option pricing: A simplified approach. *Journal of Financial Economics*, 7(3):229–263, Sept. 1979. ISSN 0304405X. doi: 10.1016/0304-405X(79)90015-1. URL https://linkinghub.elsevier.com/retrieve/pii/0304405X79900151.

H. Damon Matthews, K. B. Tokarska, J. Rogelj, C. J. Smith, A. H. MacDougall, K. Haustein, N. Mengis, S. Sippel, P. M. Forster, and R. Knutti. An integrated approach to quantifying uncertainties in the remaining carbon budget. *Communications Earth & Environment*, 2(1):7, Jan. 2021.

ISSN 2662-4435. doi: 10.1038/s43247-020-00064-9. URL https://www.nature.com/articles/s43247-020-00064-9.

K. Daniel, R. Litterman, and G. Wagner. Applying Asset Pricing Theory to Calibrate the Price of Climate Risk. Technical Report w22795, National Bureau of Economic Research, Cambridge, MA, Nov. 2016. URL http://www.nber.org/papers/w22795.pdf.

K. D. Daniel, R. B. Litterman, and G. Wagner. Declining CO_2 price paths. *Proceedings of the National Academy of Sciences*, 116(42):20886–20891, Oct. 2019. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1817444116. URL http://www.pnas.org/lookup/doi/10.1073/pnas.1817444116.

S. Dietz and F. Venmans. Cumulative carbon emissions and economic policy: In search of general principles. *Journal of Environmental Economics and Management*, 96:108–129, July 2019. ISSN 00950696. doi: 10.1016/j.jeem.2019.04.003. URL https://linkinghub.elsevier.com/retrieve/pii/S0095069618302122.

S. Dietz, J. Rising, T. Stoerk, and G. Wagner. Economic impacts of tipping points in the climate system. *Proceedings of the National Academy of Sciences*, 118(34):e2103081118, Aug. 2021. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.2103081118. URL https://pnas.org/doi/full/10.1073/pnas.2103081118.

M. A. Drupp, M. C. Freeman, B. Groom, and F. Nesje. Discounting Disentangled. *American Economic Journal: Economic Policy*, 10(4):109–134, Nov. 2018. ISSN 1945-7731, 1945-774X. doi: 10.1257/pol.20160240. URL https://pubs.aeaweb.org/doi/10.1257/pol.20160240.

M. T. Dvorak, K. C. Armour, D. M. W. Frierson, C. Proistosescu, M. B. Baker, and C. J. Smith. Estimating the timing of geophysical commitment to 1.5 and 2.0 °C of global warming. *Nature Climate Change*, 12(6):547–552, June 2022. ISSN 1758-678X, 1758-6798. doi: 10.1038/s41558-022-01372-y. URL https://www.nature.com/articles/s41558-022-01372-y.

L. G. Epstein and S. E. Zin. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica*, 57(4):937, July 1989. ISSN 00129682. doi: 10.2307/1913778. URL https://www.jstor.org/stable/1913778?origin=crossref.

L. G. Epstein and S. E. Zin. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis. *Journal of Political Economy*, 99(2):263–286, Apr. 1991. ISSN 0022-3808, 1537-534X. doi: 10.1086/261750. URL https://www.journals.uchicago.edu/doi/10.1086/261750.

O. Geoffroy, D. Saint-Martin, G. Bellon, A. Voldoire, D. J. L. Olivié, and S. Tytéca. Transient climate response in a two-layer energy-balance model. part ii: Representation of the efficacy of deep-ocean heat uptake and validation for cmip5 aogcms. *Journal of Climate*, 26(6):1859–1876, 2013a. ISSN 08948755, 15200442. URL http://www.jstor.org/stable/26192253.

O. Geoffroy, D. Saint-Martin, D. Olivié, A. Voldoire, G. Bellon, and S. Tytéca. Transient climate response in a two-layer energy-balance model. part i: Analytical solution and parameter calibration using cmip5 aogcm experiments. *Journal of Climate*, 26:1841–1857, 03 2013b. doi: 10.1175/JCLI-D-12-00195.1.

K. Gillingham and J. H. Stock. The Cost of Reducing Greenhouse Gas Emissions. *Journal of Economic Perspectives*, 32(4):53–72, Nov. 2018. ISSN 0895-3309. doi: 10.1257/jep.32.4.53. URL https://pubs.aeaweb.org/doi/10.1257/jep.32.4.53.

M. Golosov, J. Hassler, P. Krusell, and A. Tsyvinski. Optimal Taxes on Fossil Fuel in General Equilibrium. *Econometrica*, 82(1):41–88, 2014. ISSN 0012-9682. doi: 10.3982/ECTA10217. URL http://doi.wiley.com/10.3982/ECTA10217.

M. C. Hänsel, M. A. Drupp, D. J. Johansson, F. Nesje, C. Azar, M. C. Freeman, B. Groom, and T. Sterner. Climate economics support for the UN climate targets. *Nature Climate Change*, 10(8): 781–789, 2020.

Z. Hausfather and G. P. Peters. Emissions – the 'business as usual' story is misleading. *Nature*, 577(7792):618–620, Jan. 2020. doi: 10.1038/d41586-020-00177-3. URL https://www.nature.com/articles/d41586-020-00177-3.

E. Hawkins and R. Sutton. The Potential to Narrow Uncertainty in Regional Climate Predictions. *Bulletin of the American Meteorological Society*, 90(8):1095–1108, Aug. 2009. ISSN 0003-0007, 1520-0477. doi: 10.1175/2009BAMS2607.1. URL https://journals.ametsoc.org/doi/10.1175/2009BAMS2607.1.

P. H. Howard and T. Sterner. Few and Not So Far Between: A Meta-analysis of Climate Damage Estimates. *Environmental and Resource Economics*, 68(1):197–225, Sept. 2017. ISSN 0924-6460, 1573-1502. doi: 10.1007/s10640-017-0166-z. URL http://link.springer.com/10.1007/s10640-017-0166-z.

Intergovernmental Panel on Climate Change. Climate change 2021: The physical science basis. 2021. URL https://report.ipcc.ch/ar6wg1/pdf/IPCC_AR6_WGI_FinalDraft_FullReport.pdf.

Intergovernmental Panel on Climate Change. Climate change 2022: Impacts, adaptation, and vulnerability. 2022a. URL https://report.ipcc.ch/ar6wg2/pdf/IPCC_AR6_WGII_FinalDraft_FullReport.pdf.

Intergovernmental Panel on Climate Change. Climate change 2022: Mitigation of climate change. 2022b. URL https://report.ipcc.ch/ar6wg3/pdf/IPCC_AR6_WGIII_FinalDraft_FullReport.pdf.

International Energy Agency. Direct Air Capture 2022 – Analysis. Technical report, International Energy Agency, Apr. 2022. URL https://www.iea.org/reports/direct-air-capture-2022.

A. B. Jaffe and R. N. Stavins. The energy paradox and the diffusion of conservation technology. *Resource and Energy Economics*, 16(2):91–122, May 1994. ISSN 09287655. doi: 10.1016/0928-7655(94)90001-9. URL https://linkinghub.elsevier.com/retrieve/pii/0928765594900019.

K. Johnson, D. Martin, X. Zhang, C. DeYoung, and A. Stolberg. Carbon Dioxide Removal Options: A Literature Review Identifying Carbon Removal Potentials and Costs. 2017. URL http://deepblue.lib.umich.edu/handle/2027.42/136610.

F. Joos, R. Roth, J. S. Fuglestvedt, G. P. Peters, I. G. Enting, W. von Bloh, V. Brovkin, E. J. Burke, M. Eby, N. R. Edwards, T. Friedrich, T. L. Frölicher, P. R. Halloran, P. B. Holden, C. Jones, T. Kleinen, F. T. Mackenzie, K. Matsumoto, M. Meinshausen, G.-K. Plattner, A. Reisinger, J. Segschneider, G. Shaffer, M. Steinacher, K. Strassmann, K. Tanaka, A. Timmermann, and A. J. Weaver. Carbon dioxide and climate impulse response functions for the computation of greenhouse gas metrics: a multi-model analysis. *Atmospheric Chemistry and Physics*, 13(5):2793–2825, Mar. 2013. ISSN 1680-7324. doi: 10.5194/acp-13-2793-2013. URL https://acp.copernicus.org/articles/13/2793/2013/.

R. E. Kopp, R. L. Shwom, G. Wagner, and J. Yuan. Tipping elements and climate–economic shocks: Pathways toward integrated assessment. *Earth's Future*, 4(8):346–372, 2016.

M. J. Kotchen, J. A. Rising, and G. Wagner. The costs of "costless" climate mitigation. *Science*, 382(6674):1001–1003, Dec. 2023. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.adj2453. URL https://www.science.org/doi/10.1126/science.adj2453.

A. Lans Bovenberg and S. Smulders. Environmental quality and pollution-augmenting technological change in a two-sector endogenous growth model. *Journal of Public Economics*, 57(3):369–391, July 1995. ISSN 00472727. doi: 10.1016/0047-2727(95)80002-Q. URL https://linkinghub.elsevier.com/retrieve/pii/004727279580002Q.

F. Lehner, C. Deser, N. Maher, J. Marotzke, E. M. Fischer, L. Brunner, R. Knutti, and E. Hawkins. Partitioning climate projection uncertainty with multiple large ensembles and CMIP5/6. *Earth System Dynamics*, 11(2):491–508, May 2020. ISSN 2190-4987. doi: 10.5194/esd-11-491-2020. URL https://esd.copernicus.org/articles/11/491/2020/.

D. Lemoine. The Climate Risk Premium: How Uncertainty Affects the Social Cost of Carbon. *Journal of the Association of Environmental and Resource Economists*, 8(1):27–57, Jan. 2021. ISSN 2333-5955, 2333-5963. doi: 10.1086/710667. URL https://www.journals.uchicago.edu/doi/10.1086/710667.

D. Lemoine and I. Rudik. Managing Climate Change Under Uncertainty: Recursive Integrated Assessment at an Inflection Point. *Annual Review of Resource Economics*, 9(1):117–142, Oct. 2017. ISSN 1941-1340, 1941-1359. doi: 10.1146/annurev-resource-100516-053516. URL https://www.annualreviews.org/doi/10.1146/annurev-resource-100516-053516.

D. Lemoine and C. P. Traeger. Ambiguous tipping points. *Journal of Economic Behavior & Organization*, 132:5–18, Dec. 2016a. ISSN 01672681. doi: 10.1016/j.jebo.2016.03.009. URL https://linkinghub.elsevier.com/retrieve/pii/S0167268116300221.

D. Lemoine and C. P. Traeger. Economics of tipping the climate dominoes. *Nature Climate Change*, 6(5):514–519, May 2016b. ISSN 1758-678X, 1758-6798. doi: 10.1038/nclimate2902. URL https://www.nature.com/articles/nclimate2902.

T. M. Lenton, H. Held, E. Kriegler, J. W. Hall, W. Lucht, S. Rahmstorf, and H. J. Schellnhuber. Tipping elements in the Earth's climate system. *Proceedings of the National Academy of Sciences*, 105(6):1786–1793, Feb. 2008. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.0705414105. URL https://pnas.org/doi/full/10.1073/pnas.0705414105.

R. E. Lucas. Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy*, 1:19–46, Jan. 1976. ISSN 01672231. doi: 10.1016/S0167-2231(76)80003-6. URL https://linkinghub.elsevier.com/retrieve/pii/S0167223176800036.

McKinsey & Company. *Pathways to a low-carbon economy: Version 2 of the global greenhouse gas abatement cost curve*. McKinsey & Company, Stockholm, Sept. 2013. URL https://www.mckinsey.com/capabilities/sustainability/our-insights/pathways-to-a-low-carbon-economy.

M. Meinshausen, Z. R. J. Nicholls, J. Lewis, M. J. Gidden, E. Vogel, M. Freund, U. Beyerle, C. Gessner, A. Nauels, N. Bauer, J. G. Canadell, J. S. Daniel, A. John, P. B. Krummel, G. Luderer, N. Meinshausen, S. A. Montzka, P. J. Rayner, S. Reimann, S. J. Smith, M. van den Berg, G. J. M. Velders, M. K. Vollmer, and R. H. J. Wang. The shared socio-economic pathway (SSP) greenhouse gas concentrations and their extensions to 2500. *Geoscientific Model Development*, 13(8):3571–3605, Aug. 2020. ISSN 1991-9603. doi: 10.5194/gmd-13-3571-2020. URL https://gmd.copernicus.org/articles/13/3571/2020/.

National Center for Energy Economics. *Supplementary Material for the Regulatory Impact Analysis for the Supplemental Proposed Rulemaking, "Standards of Performance for New, Reconstructed, and Modified Sources and Emissions Guidelines for Existing Sources: Oil and Natural Gas Sector Climate Review"*. U.S. Environmental Protection Agency, Washington D.C., Sept. 2022.

National Research Council. *Climate Intervention: Carbon Dioxide Removal and Reliable Sequestration*. National Academies Press, Washington, D.C., June 2015. ISBN 9780309305297. doi: 10.17226/18805. URL http://www.nap.edu/catalog/18805.

New York State Energy Research and Development Authority and Resources for the Future. Estimating the cost of carbon: two approaches. Technical report, Oct. 2020. URL https://media.rff.org/documents/RFF_NYSERDA_Valuing_Carbon_Synthesis_Memo.pdf.

R. G. Newell, W. A. Pizer, and B. C. Prest. A Discounting Rule for the Social Cost of Carbon. *Journal of the Association of Environmental and Resource Economists*, 9(5):1017–1046, Sept. 2022. ISSN 2333-5955, 2333-5963. doi: 10.1086/718145. URL https://www.journals.uchicago.edu/doi/10.1086/718145.

J. Nielsen-Gammon and M. Behl. Improving the value of climate data and models for assessing climate impacts and policies. Technical report, University Corporation for Atmospheric Research (UCAR), 2021. URL https://opensky.ucar.edu/islandora/object/usclivar:131.

W. D. Nordhaus. An Optimal Transition Path for Controlling Greenhouse Gases. *Science*, 258(5086): 1315–1319, Nov. 1992. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.258.5086.1315. URL https://www.sciencemag.org/lookup/doi/10.1126/science.258.5086.1315.

W. D. Nordhaus. A review of the stern review on the economics of climate change. *Journal of economic literature*, 45(3):686–702, 2007.

W. D. Nordhaus. Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, 114(7):1518–1523, Feb. 2017. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1609244114. URL https://pnas.org/doi/full/10.1073/pnas.1609244114.

R. S. Pindyck. Climate Change Policy: What Do the Models Tell Us? *Journal of Economic Literature*, 51(3):860–872, Sept. 2013. ISSN 0022-0515. doi: 10.1257/jel.51.3.860. URL https://pubs.aeaweb.org/doi/10.1257/jel.51.3.860.

K. Rennert, F. Errickson, B. C. Prest, L. Rennels, R. G. Newell, W. Pizer, C. Kingdon, J. Wingenroth, R. Cooke, B. Parthum, D. Smith, K. Cromar, D. Diaz, F. C. Moore, U. K. Müller, R. J. Plevin, A. E. Raftery, H. Ševčíková, H. Sheets, J. H. Stock, T. Tan, M. Watson, T. E. Wong, and D. Anthoff. Comprehensive Evidence Implies a Higher Social Cost of CO_2. *Nature*, Sept. 2022. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-022-05224-9. URL https://www.nature.com/articles/s41586-022-05224-9.

K. Riahi, D. P. van Vuuren, E. Kriegler, J. Edmonds, B. C. O'Neill, S. Fujimori, N. Bauer, K. Calvin, R. Dellink, O. Fricko, W. Lutz, A. Popp, J. C. Cuaresma, S. Kc, M. Leimbach, L. Jiang, T. Kram, S. Rao, J. Emmerling, K. Ebi, T. Hasegawa, P. Havlik, F. Humpenöder, L. A. Da Silva, S. Smith, E. Stehfest, V. Bosetti, J. Eom, D. Gernaat, T. Masui, J. Rogelj, J. Strefler, L. Drouet, V. Krey, G. Luderer, M. Harmsen, K. Takahashi, L. Baumstark, J. C. Doelman, M. Kainuma, Z. Klimont, G. Marangoni, H. Lotze-Campen, M. Obersteiner, A. Tabeau, and M. Tavoni. The Shared Socioeconomic Pathways and their energy, land use, and greenhouse gas emissions implications: An overview. *Global Environmental Change*, 42:153–168, Jan. 2017. ISSN 09593780. doi: 10.1016/j.gloenvcha.2016.05.009. URL https://linkinghub.elsevier.com/retrieve/pii/S0959378016300681.

S. K. Rose, D. B. Diaz, and G. J. Blanford. Understanding the Social Cost of Carbon: A Model Diagnostic and Inter-comparison Study. *Climate Change Economics*, 08(02):1750009, May 2017. ISSN 2010-0078, 2010-0086. doi: 10.1142/S2010007817500099. URL https://www.worldscientific.com/doi/abs/10.1142/S2010007817500099.

F. Schroyen and K. O. Aarbu. Attitudes Towards Large Income Risk in Welfare States: An International Comparison, Dec. 2017. URL https://papers.ssrn.com/abstract=3089690.

C. J. Smith, P. M. Forster, M. Allen, N. Leach, R. J. Millar, G. A. Passerello, and L. A. Regayre. FAIR v1.3: a simple emissions-based impulse response and carbon cycle model. *Geoscientific Model Development*, 11(6):2273–2297, June 2018. ISSN 1991-9603. doi: 10.5194/gmd-11-2273-2018. URL https://gmd.copernicus.org/articles/11/2273/2018/.

N. Stern. The Structure of Economic Modeling of the Potential Impacts of Climate Change: Grafting Gross Underestimation of Risk onto Already Narrow Science Models. *Journal of Economic Literature*, 51(3):838–859, Sept. 2013. ISSN 0022-0515. doi: 10.1257/jel.51.3.838. URL https://pubs.aeaweb.org/doi/10.1257/jel.51.3.838.

L. Summers and R. Zeckhauser. Policymaking for posterity. *Journal of Risk and Uncertainty*, 37 (2-3):115–140, Dec. 2008. ISSN 0895-5646, 1573-0476. doi: 10.1007/s11166-008-9052-y. URL http://link.springer.com/10.1007/s11166-008-9052-y.

United Nations Framework Convention on Climate Change. *Adoption of the Paris Agreement. I: Proposal by the President*. United Nations Office, Geneva, 2015. URL https://unfccc.int/documents/184656.

T. S. Van Den Bremer and F. Van Der Ploeg. The Risk-Adjusted Carbon Price. *American Economic Review*, 111(9):2782–2810, Sept. 2021. ISSN 0002-8282. doi: 10.1257/aer.20180517. URL https://pubs.aeaweb.org/doi/10.1257/aer.20180517.

J. von Neumann and O. Morgenstern. *Theory of games and economic behavior*. Princeton University Press, 1947.

P. Weil. Nonexpected Utility in Macroeconomics. *The Quarterly Journal of Economics*, 105(1): 29, Feb. 1990. ISSN 00335533. doi: 10.2307/2937817. URL https://academic.oup.com/qje/article-lookup/doi/10.2307/2937817.

M. L. Weitzman. Why the Far-Distant Future Should Be Discounted at Its Lowest Possible Rate. *Journal of Environmental Economics and Management*, 36(3):201–208, Nov. 1998. ISSN 00950696. doi: 10.1006/jeem.1998.1052. URL https://linkinghub.elsevier.com/retrieve/pii/S009506969891052X.

M. L. Weitzman. A review of the stern review on the economics of climate change. *Journal of economic literature*, 45(3):703–724, 2007.

M. L. Weitzman. On Modeling and Interpreting the Economics of Catastrophic Climate Change. *Review of Economics and Statistics*, 91(1):1–19, Feb. 2009. ISSN 0034-6535, 1530-9142. doi: 10.1162/rest.91.1.1. URL https://direct.mit.edu/rest/article/91/1/1-19/57734.

M. L. Weitzman. GHG Targets as Insurance Against Catastrophic Climate Damages. *Journal of Public Economic Theory*, 14(2):221–244, Mar. 2012. ISSN 10973923. doi: 10.1111/j.1467-9779.2011.01539.x. URL http://doi.wiley.com/10.1111/j.1467-9779.2011.01539.x.

World Bank. State and Trends of Carbon Pricing 2021. Serial, World Bank, Washington, DC, May 2021. URL https://openknowledge.worldbank.org/handle/10986/35620.

T. P. Wright. Factors Affecting the Cost of Airplanes. *Journal of the Aeronautical Sciences*, 3(4): 122–128, Feb. 1936. ISSN 1936-9956. doi: 10.2514/8.155. URL https://arc.aiaa.org/doi/10.2514/8.155.

J. A. Yarmuth. Inflation Reduction Act of 2022, Aug. 2022. URL https://www.congress.gov/bill/117th-congress/house-bill/5376.

Online Appendix: Carbon Dioxide as a Risky Asset

Adam Michael Bauer*[1], Cristian Proistosescu[2,3], and Gernot Wagner[4]

[1] *Department of Physics, University of Illinois Urbana-Champaign, 1110 W Green St Loomis Laboratory, Urbana, IL 61801, USA*
[2] *Department of Climate, Meteorology, and Atmospheric Sciences, University of Illinois Urbana-Champaign, 1301 W Green St, Urbana, IL 61801, USA*
[3] *Department of Earth Science and Environmental Change, University of Illinois Urbana-Champaign, 1301 W Green St, Urbana, IL 61801, USA*
[4] *Columbia Business School, 665 W 130th St, New York, NY 10027, USA*

Forthcoming in *Climatic Change*
March 13, 2024

Abstract

We develop a financial-economic model for carbon pricing with an explicit representation of decision making under risk and uncertainty that is consistent with the Intergovernmental Panel on Climate Change's sixth assessment report. We show that risk associated with high damages in the long term leads to stringent mitigation of carbon dioxide emissions in the near term, and find that this approach provides economic support for stringent warming targets across a variety of specifications. Our results provide insight into how a systematic incorporation of climate-related risk influences optimal emissions abatement pathways.

JEL: G0, G12, Q51, Q54
Keywords: Climate risk, climate policy, asset pricing, cost of carbon

*Corresponding author email: adammb4@illinois.edu

A Brief literature review

There are three primary ways that risk and uncertainty are incorporated into climate-economic integrated assessment models. The first such approach is to augment DICE with stochastic components and reframe the model into a dynamic stochastic optimal control problem. This approach has yielded a number of fruitful insights (see, e.g., Lemoine and Traeger (2016a,b)). For example, the seminal work of Cai and Lontzek (2019) show that the possibility of hitting a climate tipping point substantially increases the SCC, and thus the stringency of optimal mitigation policy, using a continuous-time version of DICE with many stochastic components.

Another approach is to formulate climate policy from the perspective of dynamic stochastic general equilibrium (DSGE) models. This approach was pioneered by Golosov et al. (2014), who derive a simple expression for the marginal externality damage from carbon emissions (analogously, the optimal carbon price or SCC). This expression shows that the optimal carbon price can be – in a stylized setting – decomposed into three contributing factors: (i) the discount rate, (ii) the elasticity of damage associated with a marginal ton of emissions, and (iii) the rate of depreciation of carbon stocks in the atmosphere. However, this study does not include temperature uncertainty, and utilizes a logarithmic utility, which causes the role of uncertainty to be substantially suppressed. Van Den Bremer and Van Der Ploeg (2021) extend the DGSE framework to include recursive preferences, finding that the influence of temperature and climate damage uncertainty increase the SCC. Similar conclusions were drawn by Hambel et al. (2021), whose formulation allows for multiple, additive climate shocks, as well as for considering the influence of climate change on both GDP levels and growth rates.

A third approach is to employ methods from financial economics to explore the influence of uncertainty on carbon prices. Dietz et al. (2018) utilize a simple analytic model derive the consumption-based capital asset pricing model "beta" (Lucas, 1978) for climate mitigation projects. They find that the sign of the "climate beta" is positive, and that the discounted expected net benefits of carbon emissions abatement are increasing in the "climate beta". However, they do not utilize recursive preferences in their approach. Lemoine (2021) formulates a simple analytic expression that highlights the various channels of uncertainty associated with the SCC and signs each; the collective effect is positive. Barnett et al. (2020) build a dynamic structural model which includes decision making under uncertainty, nonlinear impulse response functions, and dynamic valuation, and find that the influence of uncertainty is multiplicative across economic and climate channels. Each of these contributions provide relatively simple – yet powerful – explanations, in financial economic terms, of how uncertainty influences optimal carbon pricing.

B Statement of optimization problem

Put together, solving CAP6 is equivalent to solving the following optimization problem:

$$\max_{\{x_t\}_{t\in\{0,1,...,T-1\}}} U_0(x_t), \tag{B.1}$$

$$\text{Such that}: x_t \in \mathbb{R}^+, \tag{B.2}$$

$$U_t = \left[(1-\beta)c_t^\rho + \beta\left(\mathbb{E}_t\left[U_{t+1}^\alpha\right]^{\rho/\alpha}\right)\right]^{1/\rho} \tag{B.3}$$

$$U_T = \left(\frac{1-\beta}{1-\beta(1+g)^\rho}\right)^{1/\rho} c_T, \tag{B.4}$$

$$c_t = \bar{c}_t(1-\kappa_t(x_t))(1-\mathcal{D}_t(\Psi_t,\theta_t)), \qquad \forall t \in \{0,1,...,T-1\} \tag{B.5}$$

$$c_T = \bar{c}_T(1-\mathcal{D}_t(\Psi_T,\theta_T)), \tag{B.6}$$

$$\bar{c}_t = \bar{c}_0(1+g)^t, \tag{B.7}$$

$$\Psi_t = \int_0^t E_\zeta d\zeta, \tag{B.8}$$

$$\mathcal{D}_t(\Psi_t,\theta_t) = \sum_{\theta_t\in\Theta_t} P(\theta_T|\theta_t)\mathcal{D}_{tot}(\Psi_t,\theta_t) \tag{B.9}$$

$$\kappa_t(x_t) = \kappa_{MACC}(x_t)(1-\varphi_0-\varphi_1 X_t)^{t-10}, \tag{B.10}$$

$$X_t = \frac{\int_0^t x_\zeta E_\zeta d\zeta}{\Psi_t}, \tag{B.11}$$

$$\beta = \frac{1}{1+\delta} \tag{B.12}$$

$$0 \le P(\theta_T|\theta_t) \le 1, \tag{B.13}$$

$$\delta,\rho,\alpha,g,\varphi_0,\varphi_1,\tau_{DAC},E_t \text{ given and positive}, \tag{B.14}$$

$$\mathcal{D}_0(\Psi_t,\theta_t) = 0. \tag{B.15}$$

See the main text for details regarding functional forms, calibrations, and results after numerically solving the model.

C Prototypical model run

Given that our climate-economy model is unlike most other such models in the literature, an example of how each of the components laid out above interact in one model "run" is warranted. First, let us establish some important concepts and recurring values that will be essential for our understanding. We have chosen $T = 6$ decision periods. This implies that we have a total of $n := 2^T - 1 = 63$ decision nodes in the tree. Decisions are made at times t such that $t \in \{2020, 2030, 2060, 2100, 2150, 2200\}$, and an additional period (with no decisions being made) occurs at $t = 2250$ to establish the terminal period conditions. As the binomial tree is path dependent, it immediately follows that the number of unique paths through the tree is equal to the number of nodes in the *final* period, given by $n_f := 2^{T-1} = 32$. Any vector of length n (which represents the value of a given variable, say mitigation, at each *node* in the tree) can be readily translated into a set of *paths* of shape $n_f \times T$ through the tree (which represents the values of a given variable at the nodes in each *path* through the tree). Note that Figure 1 in the main text is a helpful visual guide for our entire discussion.

Step 1: Simulate climate damages

The first step is to simulate potential climate damages. This comes *before* agent utility is optimized, as decisions about utility are made *within the context* of the landscape of potential damages. Once the landscape of damages are calculated (and we will be more precise about what is meant by "landscape" in our discussion below), then damages are interpolated in our utility calculations. Note that in the following discussion $N_{MC} = 3 \times 10^6$ refers to the number of draws taken in our Monte Carlo samples of TCRE and damage function parametric uncertainty.

Climate damages are simulated using the following prescription. First, we specify an emissions baseline by choosing an SSP. Once specified, there is a range of possible cumulatively emitted CO_2 at each point in time, depending on hypothetical agent mitigation policy. Let the maximum cumulative emissions (associated with no mitigation) at a time t be represented by Ψ_t^*. Cumulative emissions Ψ_t therefore always lie in the range $0 \leq \Psi_t \leq \Psi_t^*$. We discretize the range of potential cumulative emissions at each point in time by applying a constant scaling $0 \leq m \leq 1$ to the SSP and computing damages for each value of m. In our runs, we choose $M = 101$ values of m. To recapitulate: we choose a value of m such that $0 \leq m \leq 1$, resulting in a time series of cumulative emissions $\Psi_t = m\Psi_t^*$ that is manifestly less than or equal to the maximum permissible amount Ψ_t^* for all t.

For a given time series of cumulative emissions, the corresponding temperature change is uncertain owing to the uncertainty in the TCRE. We draw N_{MC} samples of the TCRE from a rectified normal distribution with best estimate and variance taken from Table 4 and evaluate (3.2), which results in N_{MC} time series of global temperature change. For each temperature time series, we at random choose a damage function (statistical, structural, or meta-analytic) and evaluate (2.6) for the chosen damage function and the additional tipping points piece. The total damage is given by (2.7). This procedure results in N_{MC} time series of climate damages. The climate damage time series are then ordered by severity of the final period damages (thus establishing an orientation of the "fragility" dimension), and grouped in N_{MC}/n_f sized bundles. An average is then taken over each bundle, resulting in n_f time series of climate damages. The averaging procedure is necessary to make the simulated climate damages congruent with the dimensionality of the binomial tree.

The procedure described above has resulted in a $n_f \times T$ matrix of climate damages, ordered from high to low. Continuing for every value of m results in a $M \times n_f \times T$ *landscape* of climate damages. This is coined as a landscape owing to its encapsulation of the potential extent of climate damages. The M-dimension contains information about the extent of emissions; the T-dimension contains information about the timing of damages; and the n_f-dimension contains the extent of climate damages based on the uncertainty in TCRE and the damage function. With this landscape now calculated, we can turn our attention to how the economic utility is maximized within it.

Step 2: Utility maximization

We optimize the economic utility given by (2.1) using a genetic algorithm (Goldberg, 1989). The genetic algorithm is a stochastic optimization routine, where a set of random solution vectors are generated and their "fitness" is determined. The vectors with high fitness are stored for the next round (they "survive"), and vectors with low fitness are discarded (they "die"). The low fitness vectors are replaced with another set of random vectors (the "offspring" of the more fit vectors) whose fitness is compared to the incumbents'. This process continues until minimal changes in the highest fitness value are recorded for a number of rounds; the vector corresponding to the highest fitness is then said to be the "optimal" solution vector. The genetic algorithm is best suited for objective functions with unknown or difficult to evaluate gradients, making it ideal for CAP6. In our use case, the randomly selected solution vectors are mitigation vectors, and a given vector's fitness is its 2020 economic utility. In what follows, allow $\vec{x}$ be a vector of mitigation values with length n.

EZ utility captures future risk by allowing the utility at time t be dependent on the utility at time $t+1$ (see Eqn. (2.1)). Evaluating the utility must therefore begin at the final period, and is then evaluated *backwards* to $t = 2020$. Thus, the first step is to evaluate the final period utility (2.2) where the final period consumption is given by (2.5) for each final state node. (Recall there are $n_f = 32$ nodes in the final period.) The assumed SSP and the mitigation vector $\vec{x}$ are used to calculate the emissions time series for every path through the tree, and thus the cumulative emissions at each end node. The cumulative emissions are used in (2.8) to calculate the damages at each node.

For each node before the final period, the mitigation action *up to but not including* a given node is used to calculate the cumulative emissions at that node. The cost of mitigation is found using (2.12), and the damages are found using (2.8). These in tandem determine the consumption by (2.4). The consumption and the following period utility are used in (2.1) to determine the utility. This continues for each node, and each randomly generated vector, until the genetic algorithm finds the mitigation vector with the highest utility.

Step 3: Visualize model output

The most fit mitigation vector $\vec{x}^*$ translates into the output shown in Figures 5, 3, 6, 9, and 7 in the following way. To calculate the cost, we apply (2.9) at each node, including the technological growth prefactor found in (2.12). We calculate the expected mitigation using (2.13). We use $\vec{x}^*$ to calculate the emissions at each node, which readily translates into the concentrations at each node using (F.4) and the *expected* warming at each node using (3.2) assuming the mean value of TCRE. Economic damages for each node are calculated using $\vec{x}^*$ in (2.8). Averaging over the cost, expected mitigation, emissions, temperature, CO_2 concentrations, and damage amount in each period gives the time series shown in Figures in the main text and the Online Appendix.

D Supplementary discussion: damage functions

In Table 1 we show the calibrated values and uncertainties for the free parameters in (2.6). Below, we provide a technical description of how the values in Table 1 are computed.

D.1 Discussion of IPCC aggregate damage functions

D.1.1 Statistically estimated damage function

The statistically estimated damage function (Burke et al., 2018) builds on previous work involving the nonlinear response of economic productivity to temperature (Burke et al., 2015), following methodologies laid out more generally in Carleton and Hsiang (2016). This damage function relies on the specification of a certain horizon where damages set in, and choose the natural markers of 2049 and 2099 (mid-century and end of century, respectively). The mid-century and end of century estimates are starkly different, as in this framework climate change slows economic growth, therefore requiring sufficient time for damages to compound. Damages are also different depending on which SSP one chooses; this owes to the fact that each SSP contains different assumptions around adaptation, technological growth, and so on. Finally, the warming levels represented in Burke et al. (2018) are relative to a 1986–2005 baseline, *not* relative preindustrial temperature levels. The IPCC's representation of this damage function differs from the original publication in three ways: they only report end-of-century estimates; they aggregate damage estimates across SSPs without indicating the differences between each; and they report the temperature change as relative to preindustrial rather than to a 1986–2005 baseline.

Table 1: Fitted parameters for the damage function (2.6) based on Burke et al. (2018), Dietz et al. (2021) and Intergovernmental Panel on Climate Change (2022).

Damage function	$\bar{\varpi}_2$ $[K^{-2}]$	σ_{ϖ_2} $[K^{-2}]$	$\bar{\varpi}_1$ $[K^{-1}]$	σ_{ϖ_1} $[K^{-1}]$
Statistically estimated				
SSP1, mid-century	5.36×10^{-3}	7.13×10^{-4}	8.93×10^{-3}	1.12×10^{-3}
SSP2, mid-century	3.09×10^{-3}	4.76×10^{-4}	1.24×10^{-2}	1.90×10^{-3}
SSP3, mid-century	2.95×10^{-3}	4.74×10^{-4}	1.18×10^{-2}	1.89×10^{-3}
SSP4, mid-century	3.50×10^{-3}	7.14×10^{-4}	5.83×10^{-3}	1.19×10^{-3}
SSP5, mid-century	3.40×10^{-3}	5.20×10^{-4}	1.14×10^{-2}	1.75×10^{-3}
SSP1, end-of-century	-1.24×10^{-3}	2.49×10^{-4}	7.07×10^{-2}	1.42×10^{-2}
SSP2, end-of-century	-2.33×10^{-3}	4.75×10^{-4}	7.21×10^{-2}	1.47×10^{-2}
SSP3, end-of-century	-2.81×10^{-3}	5.93×10^{-4}	7.20×10^{-2}	1.52×10^{-2}
SSP4, end-of-century	-1.11×10^{-3}	3.42×10^{-4}	4.67×10^{-2}	1.43×10^{-2}
SSP5, end-of-century	-1.33×10^{-3}	3.45×10^{-4}	5.56×10^{-2}	1.45×10^{-2}
Structurally estimated	2.30×10^{-3}	8.53×10^{-4}	2.05×10^{-3}	7.59×10^{-4}
Meta analysis	6.85×10^{-3}	2.43×10^{-3}	2.98×10^{-4}	1.06×10^{-4}
Climate tipping points	4.8×10^{-1}	4×10^{-2}	-4×10^{-2}	1×10^{-2}

We correct these inconsistencies in our formulation to be consistent with the original publication. We include explicitly the time dependence of this damage function in our simulated climate damages, allowing for the decision periods of 2030 and 2060 to use the mid-century estimates and each decision period from 2100 onward to use end of century estimates. This of course is not perfect, as damages are expected to continue growing past 2100 in their framework, but we lack projection data to extend their framework to longer time horizons. Therefore, our estimates of climate damages in the long run are to be considered as conservative. We also change the fit to damage function data based on which SSP we consider. Finally, we correct the temperature baseline by shifting the abscissa by ~ 0.8 °C to correctly represent temperature anomalies relative to preindustrial levels.

A final qualifier to our use of this damage function is our parameterization of uncertainty. The uncertainty range for these estimates is large, and net-benefits of climate change are not ruled out even in the long term (though they are exceptionally rare). The extent of this uncertainty is largely driven by the assumed economic response to climate change and the discount rate chosen in their model, and no range is given for the estimates of economic damages for a given climate model's projection; only the median estimate is reported for each climate model. We also suppress climate model uncertainty in their presented results so as to not double count climate uncertainty, resulting in a more narrow uncertainty envelope for damages estimates. We present our formulation in Figure 3A, taking care to allow uncertainty to broaden between 2049 and 2099, consistent with the original publication.

D.1.2 Structurally estimated damage function

In the case of the structurally estimated damage function (Rose et al., 2017), three IAMs' (DICE (Nordhaus, 1992), PAGE (Hope et al., 1993), and FUND (Tol, 1999)) output are aggregated to form a range of climate damages estimates as a function of temperature. The central value of climate damages is close to that of DICE–2023 (Barrage and Nordhaus, 2023). The uncertainty associated with this damage function results from sampling the input parameter distribution of each IAM (Rose et al., 2017). We present our formulation of this damage function based on IPCC data in Figure 3B.

D.1.3 Meta-analytic damage function

The meta-analytic damage function (Howard and Sterner, 2017) is derived from a synthesis of studies found in the literature, where care was taken to account for duplicates of studies and methodology. We use the preferred damage function from Howard and Sterner (2017), and assign an uncertainty envelope which encompasss much of the spread in the data reported by the IPCC, see Figure 3C. One limitation of this approach is that it is unclear if a set of damage estimates using different models and estimation types can be joined together in this way to form one unified "damage function"; moreover, it is also unclear if the uncertainty found in the data can truly be labeled as "parametric" or simply a by-product of disagreements in the literature.

D.1.4 Synthesis

The inability to properly compare damage estimates across studies and methodologies led WGII to conclude that a reliable range of damage estimates could not be determined; there is no single 'correct' damage function that we can specify in this work (Intergovernmental Panel on Climate Change, 2022). We resolve this issue by taking a conservative approach and sampling all of the damage functions mentioned above with equal probabilities; in this way, we remain agnostic about which damage function is the 'correct' one, and sample the space of possible damage functions in addition to uncertainty inherent to a specific climate damage estimation methodology.

Despite the issues with individual damage functions described above, our approach to sampling all available damage functions has the benefit that, at minimum, we sample a variety of damage function shapes and scales. The statistically estimated damage function has a concave shape at end-of-century. Furthermore, this damage function is time dependent, capturing the impact of climate change impacting economic growth; this has been shown to be an important factor in climate policy (Moore and Diaz, 2015). The structurally estimated damage function, in contrast, is convex, with low damages in the short run which slowly rise in temperature. Finally, the meta-analytic damage function is also convex, but rises much faster than the structurally estimated damage function.

D.2 Damage function calibration

We fit the damage function data in the following way. For each damage function, we require that the concavity of the damage function is preserved, i.e., $\partial^2\mathcal{D}/\partial T'^2 \gtrless 0$, depending on the damage function being considered. To solve for the damage function coefficients as presented in (2.6), we require knowing the damages for two data points, generically labeled as $(T_1, \mathcal{D}_1)$ and $(T_2, \mathcal{D}_2)$. Then we can write

$$\mathcal{D}_1 = T_1(\varpi_2 T_1 + \varpi_1), \tag{D.1}$$

$$\mathcal{D}_2 = T_2(\varpi_2 T_2 + \varpi_1), \tag{D.2}$$

and, solving the above for ϖ_1 and ϖ_2, results in

$$\varpi_1 = \frac{\mathcal{D}_1 T_2^2 - \mathcal{D}_2 T_1^2}{T_2 T_1 (T_2 - T_1)}, \tag{D.3}$$

$$\varpi_2 = \frac{\mathcal{D}_2 T_1 - \mathcal{D}_1 T_2}{T_2 T_1 (T_2 - T_1)}. \tag{D.4}$$

Having established the mean state, we can now introduce uncertainty into (D.3) and (D.4). We do so by allowing $\mathcal{D}_1$ to be uncertain, assigning it a Gaussian distribution $\tilde{\mathcal{D}}_1$ with mean $\bar{\mathcal{D}}_1$ and standard deviation $\sigma_{\mathcal{D}_1}$. We link this to a distribution of $\mathcal{D}_2$ by invoking the condition $\partial^2\mathcal{D}/\partial T'^2 \gtrless 0$, immediately resulting in the condition $\varpi_2 \gtrless 0$. Using (D.4), we arrive at

$$\tilde{\mathcal{D}}_2 \gtrless \tilde{\mathcal{D}}_1 \left(\frac{T_2}{T_1}\right). \tag{D.5}$$

Eqn. (D.5) is generic for any damage function, but our model, we want either a concave up or concave down damage function. To accomplish this, we include an additional factor $\Lambda > 0$ to (D.5) such that the inequality is ensured, i.e.,

$$\tilde{\mathcal{D}}_2 = \Lambda \tilde{\mathcal{D}}_1 \left(\frac{T_2}{T_1}\right), \qquad \text{such that } \Lambda \gtrless 1. \tag{D.6}$$

Therefore, if $\Lambda > 1$, we have a concave up damage function, and if $\Lambda < 1$, we have a concave down damage function. Setting $T_1 = 3$ °C and $T_2 = 10$ °C, we fit values for $\bar{\mathcal{D}}_1$, $\sigma_{\mathcal{D}_1}$, and Λ to each set of damage function data resulting in the values presented in Table 1. See Table 2 for the values of our calibration coefficients.

E Supplementary discussion: cost of of mitigation

E.1 Marginal abatement cost curve alternative calibrations

As a sensitivity test of our marginal abatement cost curve (MACC), we increased the cost of each mitigation option by one cost bracket, eliminating the zero-cost mitigation options (i.e., "free lunch" options) that the IPCC reports in their WGIII report. The resulting cost figure is in Figure 1.

Table 2: Fitted parameters for the damage function calibration equation (D.6) based on Burke et al. (2018), Dietz et al. (2021) and Intergovernmental Panel on Climate Change (2022).

Damage function	$\bar{\mathcal{D}}_1$ [−]	$\sigma_{\mathcal{D}_1}$ [−]	Λ
Statistically estimated			
SSP1, mid-century	0.075	0.01	2.5
SSP2, mid-century	0.065	0.01	2.0
SSP3, mid-century	0.062	0.01	2.0
SSP4, mid-century	0.049	0.01	2.5
SSP5, mid-century	0.065	0.01	2.1
SSP1, end-of-century	0.2	0.04	0.87
SSP2, end-of-century	0.195	0.04	0.75
SSP3, end-of-century	0.19	0.04	0.69
SSP4, end-of-century	0.13	0.04	0.82
SSP5, end-of-century	0.155	0.04	0.82
Structurally estimated	0.027	0.01	2.8
Meta analysis of climate damages	0.063	0.022	3.3

Table 3: Fitted coefficients for (2.9), the cost of abating all emissions, $\tau_a := \tau(x = 1)$, and the percent of consumption required to abate all emissions, $\kappa_a := \kappa_{MACC}(x = 1)$, based on AR6 WGIII data for each SSP in our 'main specification' and our "no free lunches" alternative calibration. All dollar values are in 2020 $USD.

SSP	ξ	τ_0 [\$ tCO_2-eq^{-1}]	τ_a [\$ tCO_2-eq^{-1}]	κ_a [%]
Main specification				
1	1.9	27.5	153.88	3.0
2	2.4	27.5	264.09	5.9
3	2.9	27.5	457.57	11.3
4	2.5	27.5	292.15	6.69
5	3.0	27.5	526.58	13.2
"No free lunches"				
1	1.8	58.9	297.42	6.1
2	2.3	58.9	528.58	11.9
3	2.8	58.9	909.69	22.5
4	2.3	58.9	528.58	13.4
5	2.9	58.9	1011.56	26.3

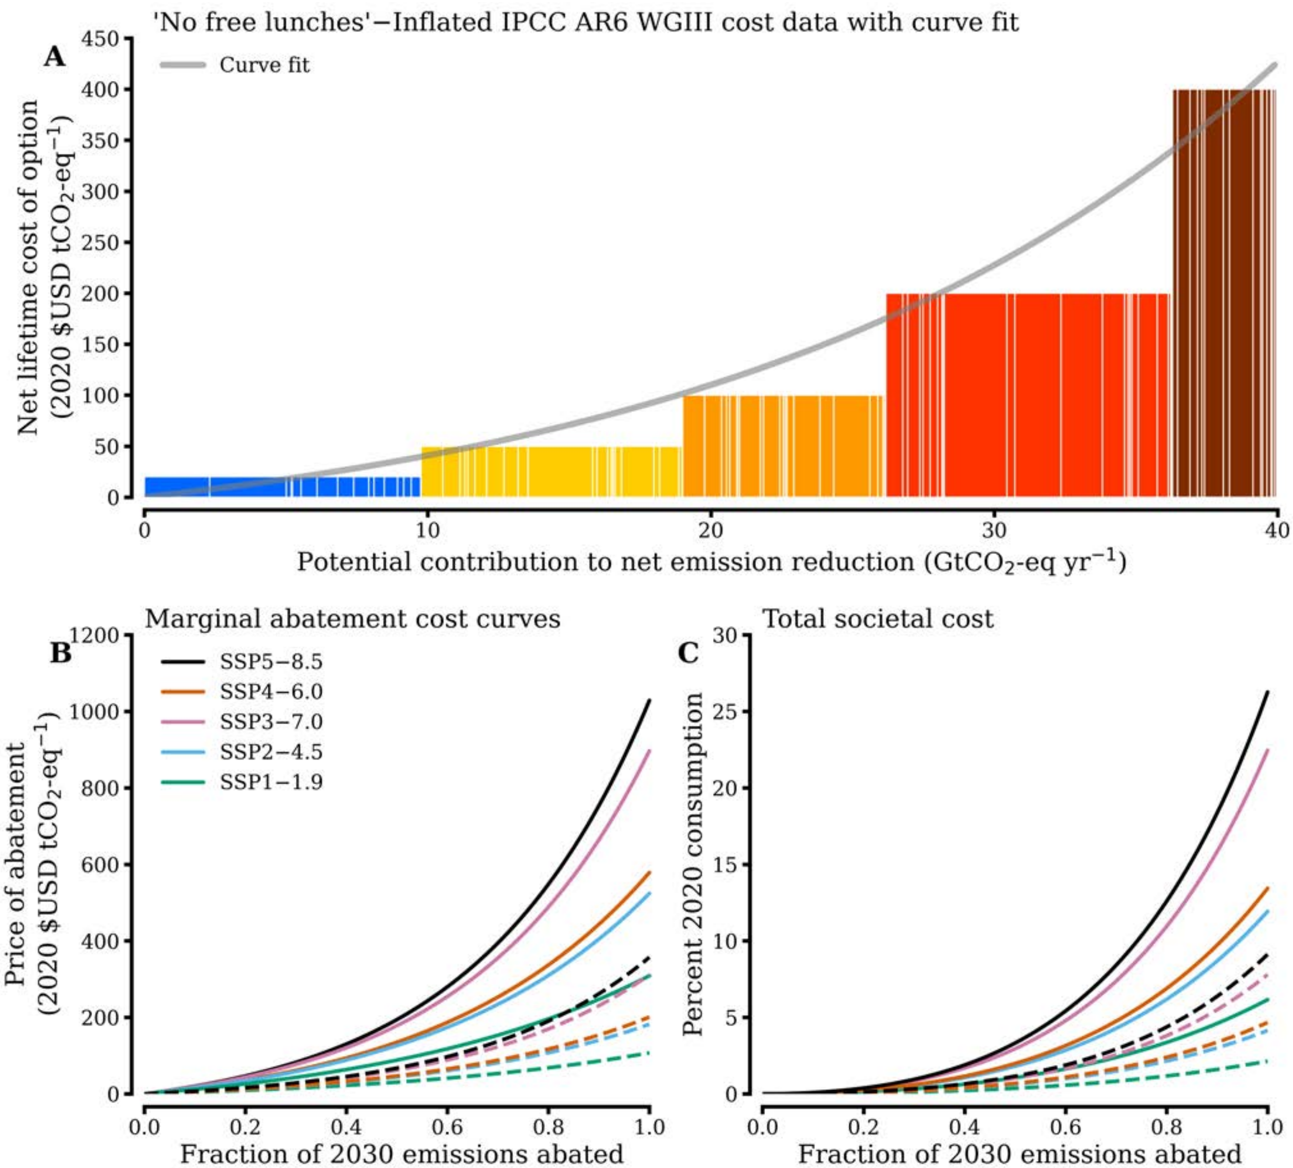

Figure 1: Panel A shows the mitigation potential and cost for each methodology given by the IPCC using their WGIII data after adjusting for the "no free lunches" calibration. Blue bars represent the \$0-\$20 range, yellow is \$20-\$50, orange is \$50-\$100, red is \$100-\$200, and maroon is our new cost bracket \$400. Our curve fit is in grey. Panel B shows the fitted marginal abatement cost curves and panel C shows the total cost to society. In panels B–C, solid lines correspond to 2030, while dashed lines are cost curves in 2100, assuming an exogenous technological growth rate of 1.5% and no endogenous technological growth.

As another sensitivity test of our marginal abatement cost curve (MACC), we cut out the < \$0 abatement potential reported by the IPCC WGIII data and fit a curve to the nonzero cost options. The resulting cost figure is in Figure 2. Note that this marginal abatement cost curve (MACC) results in costs that are lower than the "no free lunches" calibration. Hence, we do not present Climate Asset Pricing model – AR6 runs with this cost curve specified, as the results will be simple interpolations between the main specification results and the "no free lunches" results.

E.2 Limitations of our cost of abatement approach

A major qualification to our results regards two assumptions in our cost of CO_2 abatement parameterization. The first major assumption is that abatement technologies are essentially instantly able to be deployed; we do not capture real-world inertia, represented in other energy systems IAMs, that cap the rate of decarbonization owing to the delayed availability of abatement technologies, stranded assets, limited construction times, and other factors (Ha-Duong et al., 1997; Richels and Blanford, 2008; Vogt-Schilb et al., 2018). This limitation, however, is common in other IAMs such as DICE (Nordhaus, 2017) which have been widely used to study optimal climate-economic policy (Committee on Assessing Approaches to Updating the Social Cost of Carbon et al., 2017). Secondly, our MACC assumes that the sacrificed consumption to abate CO_2 emissions does not feedback on other aspects of the economy, such as growth or productivity (Hogan and Jorgenson, 1991). Including a more sophisticated abatement cost parameterization (i.e., through representing investments in abatement capital explicitly) or the feedback of mitigation policy on growth would be an interesting direction for future work. These limitations provide important context for our results.

E.3 Full derivation of total cost to society, Eqn. (2.10)

First, assume a representative agent optimizes consumption $c(\tau)$ such that $dc(\tau)/d\tau = -E(x(\tau)) = -E(\tau)$, where we have dropped the dependence of the emissions on mitigation action for clarity. Then by simple integration the consumption is given by

$$c(\tau) = \bar{c} - \underbrace{\int_0^\tau E(\zeta) d\zeta}_{=:K(\tau)}, \tag{E.1}$$

where $\bar{c} > 0$ is the baseline endowed consumption and $K(\tau)$ is the cost to society in monetary units (i.e., dollars). Eqn. (E.1) would be correct if the government was to waste the entirety of the policy proceeds, given by $E(\tau)\tau$. We instead assume that the proceeds are refunded in a lump sum (Mankiw et al., 2009), thus requiring an alteration to $K(\tau)$ such that

$$K(\tau) = \int_0^\tau E(\zeta) d\zeta - E(\tau)\tau. \tag{E.2}$$

The lump sum refund does not allow for CO_2 tax proceeds to be used to decrease distortionary taxes unrelated to CO_2 emissions; this would lower the *net* cost of CO_2 even further (Goulder, 1995; Jorgenson, 2013). Rewriting the emissions as $E(\tau) = E_0(1 - x(\tau))$ where E_0 is the (SSP–dependent) 2030 emissions in $GtCO_2$ yr^{-1}, we have

$$K(\tau) = E_0 \left(\tau x(\tau) - \int_0^\tau x(\zeta) d\zeta \right). \tag{E.3}$$

Note that E_0 is the 2030 emissions for consistency with the cost data presented by WGIII. Now using (2.9) and its inverse in (E.3), carrying out the integral, and dividing by 2020 consumption results

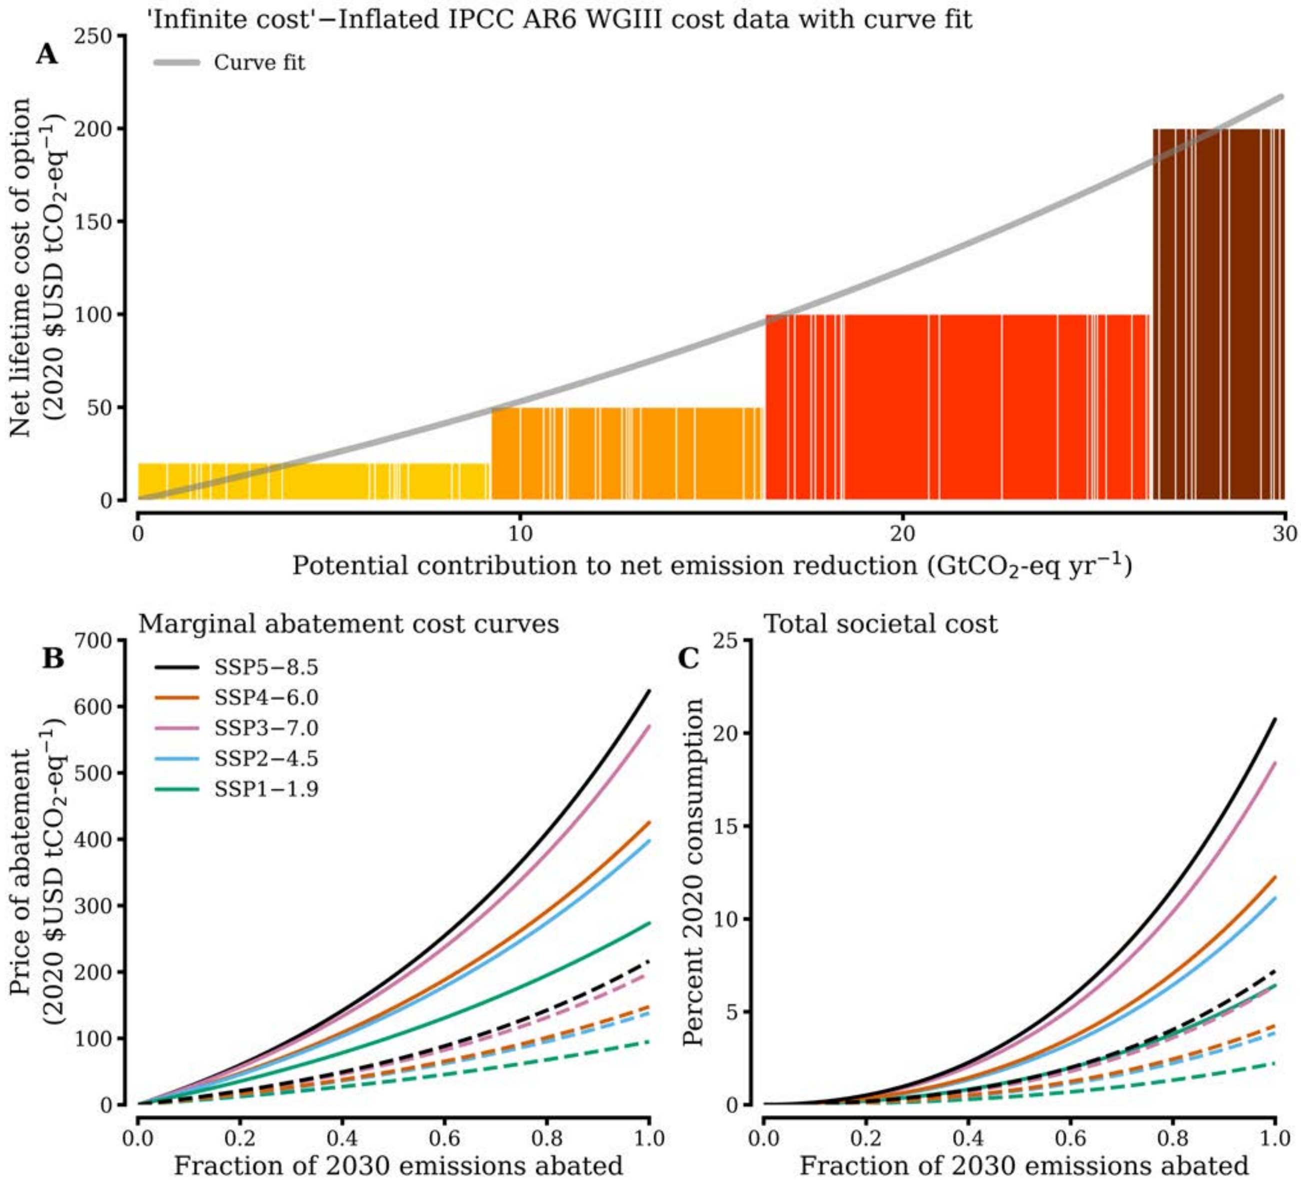

Figure 2: Panel A shows the mitigation potential and cost for each methodology given by the IPCC using their WGIII data after adjusting for the "infinite cost" calibration. Yellow bars are the $0-$20 range, orange is $20-$50, red is $50-$100, and maroon is $100-$200. Our curve fit is in grey. Panel B shows the fitted marginal abatement cost curves and Panel C shows the total cost to society. In panels B–C, solid lines correspond to 2030, while dashed lines are cost curves in 2100, assuming an exogenous technological growth rate of 1.5% and no endogenous technological growth.

Table 4: Values of fitting coefficients a_i and timescales τ_i used in (F.2) (taken from Joos et al. (2013)), as well as the best estimate and standard deviation of TCRE (taken from Intergovernmental Panel on Climate Change (2021) and Damon Matthews et al. (2021)).

Fitting Coefficient		Timescale [years]	
a_0	0.2173	τ_1	394.4
a_1	0.2240	τ_2	36.54
a_2	0.2824	τ_3	4.304
a_3	0.2763		
TCRE Parameters			
$\bar{\lambda} = 0.45$ °C $(1000 \text{ GtCO}_2)^{-1}$		$\sigma_\lambda = 0.18$ °C $(1000 \text{ GtCO}_2)^{-1}$	
$\bar{f}_{nc} = 0.14$		$\sigma_{f_{nc}} = 0.11$	
$\bar{\lambda}_{eff} = 0.52$ °C $(1000 \text{ GtCO}_2)^{-1}$		$\sigma_{\lambda_{eff}} = 0.21$ °C $(1000 \text{ GtCO}_2)^{-1}$	

in the total cost to society in terms of fractional 2020 consumption loss, given by $\kappa_{MACC}(x)$, as

$$\kappa_{MACC}(x) = \frac{E_0\tau_0}{c_{2020}}\left(\frac{e^{\xi x}-1}{\xi} - x\right), \tag{E.4}$$

where c_{2020} is the 2020 global consumption in billions of 2020 \$USD. This completes our derivation.

F Supplementary discussion: climate model

In Table 5 we compare the average warming levels using our effective TCRE approach and the weighted model averages presented by the IPCC in AR6.

F.1 Carbon cycle model

For a given emission time series the corresponding CO_2 concentration time series can be found by convolving emissions with the impulse response function (IRF) of a pulse of CO_2 emissions, denoted as $\mathcal{I}(t)$, such that

$$\mathcal{C}_E(t) = E(t) * \mathcal{I}(t). \tag{F.1}$$

In Joos et al. (2013), it is shown that the IRF for a pulse of CO_2 can be sufficiently represented by a superposition of exponentials, given by

$$\mathcal{I}(t) := a_0 + a_1 e^{-t/\tau_1} + a_2 e^{-t/\tau_2} + a_3 e^{-t/\tau_3}. \tag{F.2}$$

See Table 4 for the numerical values of the fitting coefficients a_i and timescales τ_i in (F.2).

The final component of the concentration time series accounts for pre-2020 CO_2 that is present in the atmosphere when an agent begins emitting. This ensures that our carbon cycle model not only acts to take new CO_2 out of the atmosphere, but continues to remove CO_2 from past emissions. To account for this extra CO_2 in the atmosphere, we make the assumption that the majority of CO_2 before 2020 is old, such that the time it has been in the atmosphere is much greater than τ_2. This implies that

Table 5: Shown are the central estimate and the 5%-95% range of warming levels in three time periods, for three emissions baselines, using our effective TCRE approach and what is reported by the IPCC in their Table 4.5.

Time period	Effective TCRE range (°C)	AR6 range (°C)
SSP2–4.5		
Near-term: 2021–2040	1.5 (1.3 to 1.6)	1.5 (1.2 to 1.8)
Mid-term: 2041–2060	1.9 (1.5 to 2.4)	2.0 (1.6 to 2.5)
Long-term: 2081–2100	2.6 (1.7 to 3.5)	2.7 (2.1 to 3.5)
SSP3–7.0		
Near-term: 2021–2040	1.5 (1.3 to 1.7)	1.5 (1.2 to 1.8)
Mid-term: 2041–2060	2.1 (1.6 to 2.7)	2.1 (1.7 to 2.6)
Long-term: 2081–2100	3.6 (2.1 to 5.1)	3.6 (2.8 to 4.6)
SSP5–8.5		
Near-term: 2021–2040	1.5 (1.3 to 1.7)	1.6 (1.3 to 1.9)
Mid-term: 2041–2060	2.3 (1.6 to 2.9)	2.4 (1.9 to 3.0)
Long-term: 2081–2100	4.6 (2.4 to 6.8)	4.4 (3.3 to 5.7)

there is a constant fraction that remains, and a piece that is still decaying. Hence, the remaining CO_2 in the atmosphere is given by

$$\mathcal{C}_{pre-2020}(t) = \mathcal{C}_{2020}\left(\frac{a_0 + a_1 e^{-t/\tau_1}}{a_0 + a_1}\right), \tag{F.3}$$

where $\mathcal{C}_{2020} = 420.87$ ppm.[1] Therefore, we can write the total carbon concentrations time series for a given individual as

$$\mathcal{C}(t) = \mathcal{C}_{2020}\left(\frac{a_0 + a_1 e^{-t/\tau_1}}{a_0 + a_1}\right) + E(t) * \mathcal{I}(t). \tag{F.4}$$

We note that (F.4) is used only to compute carbon concentrations as a result of optimal policy in Figures 5–8; we do not utilize carbon concentrations in our optimization routine, as our temperature parameterization relies on cumulative emissions only.

G Supplementary discussion: discount rate calibration

In Table 6 we show the term structures of the discount rates used in our featured runs. In Table 7 we show the ranges of parameters that are sampled in our ensemble runs.

Table 6: Term structures for each discount rate in featured CAP6 runs.

Discount rate [%]	δ [%]	η
1.5	0.1	0.93
2	0.2	1.20
2.5	0.5	1.42
3	0.8	1.53

Table 7: Ranges of values for each model parameter sampled in the ensemble runs.

Parameter	Symbol	Range
Risk aversion	ψ	3 – 15
Elasticity of intertemporal substitution	σ	0.55 – 1.08
Pure rate of time preference	δ	0.1% – 1.47%
Exogenous rate of technological growth	φ_0	0% – 3%
Endogenous rate of technological growth	φ_1	0% – 3%

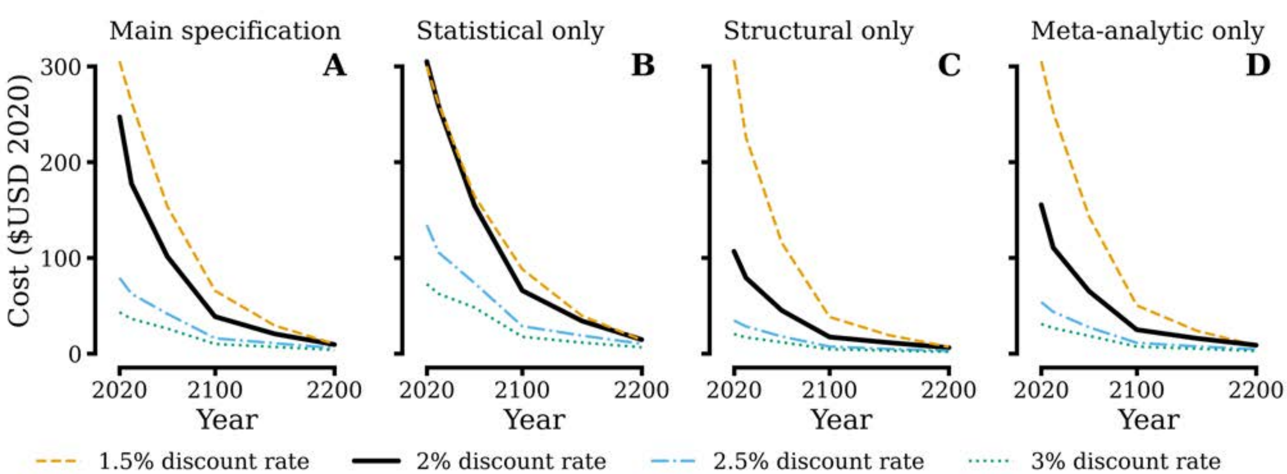

Figure 3: Price paths for each damage function. All damage functions are sampled in panel A.

H Isolating individual damage functions

We isolate the influence of each damage function on carbon price paths in Figure 3 by isolating a single damage function and re-running our featured model runs. For comparison, we also provide our featured runs in panel 3A. Beginning with the statistically estimated damage function, we find that prices are higher in the near term in comparison to the other damage functions, with the exception of the 1.5% discount rate run. By comparison, running CAP6 with a convex damage function (i.e., the structural and meta-analytic damage functions) results in lower prices in the near term, with the exception of the 1.5% discount rate runs. This shows that for sufficiently low discount rates, individual preferences can supercede the specifics of model components in 'optimal' policy considerations.

I Regression analysis

Regression coefficients in Figure 8 are calculated by fitting a linear regression between each parameter value and carbon costs. The one exception is technological growth, which is time dependent and given by

$$\varphi := \varphi_0 + \varphi_1 X_t. \tag{I.1}$$

In 2100 and later, technological change is nonlinearly related to carbon costs. We therefore fit a quadratic to carbon costs as a function of total technological growth from 2100 on. Figures 4, 5, 6, and 7 show the intermediate step in computing the results shown in Figure 8.

J Impact of Epstein-Zin risk aversion on prices

Shown in Figure 8 is the influence of changing the Epstein-Zin risk aversion parameter, ψ, on CO_2 prices. Increasing (decreasing, resp.) ψ causes an increase (decrease, resp.) in the optimal carbon tax, consistent with other studies (e.g., Cai and Lontzek, 2019).

K Including learning by doing

We run CAP6 with learning by doing (LbD) included for both our main specification and "no free lunches" MACC in Figure 9. Note we use a 2% discount rate for each curve in Figure 9, $\varphi_1 = 1.5\%$ when LbD is enabled, and all other calibration parameters are the same as in our 'main specification' runs above. We find that including LbD causes a relatively minor change in the the present-day carbon price for both MACC, and lowers the overall cost burden of the optimal abatement policy (i.e., the integrated cost over time). This is owed to prices declining faster as consumption is spent on mitigation, thus enabling more abatement in the near-term for cheaper costs. Furthermore, enabling LbD lowers the expected optimal warming by ~ 0.05 °C in 2100 for both MACCs. For the 'main specification' MACC, warming in 2200 is lower by ~ 0.1 °C, whereas for the "no free lunches" MACC 2200 warming is lower by ~ 0.12 °C.

A notable result from this exercise is that by including LbD effects, the 2% discount rate policy with our 'main specification' cost curve stays below the 1.5 °C warming target in 2200; recall this threshold was exceeded when LbD was excluded. Hence, we can expect that the feasibility of reaching the warming targets set forth in Paris are highly sensitive to such outcomes; given that the rate of endogenous technological change is difficult to empirically ground, this represents a significant source of uncertainty in policy projections and a target for future research.

[1] Taken from https://keelingcurve.ucsd.edu/

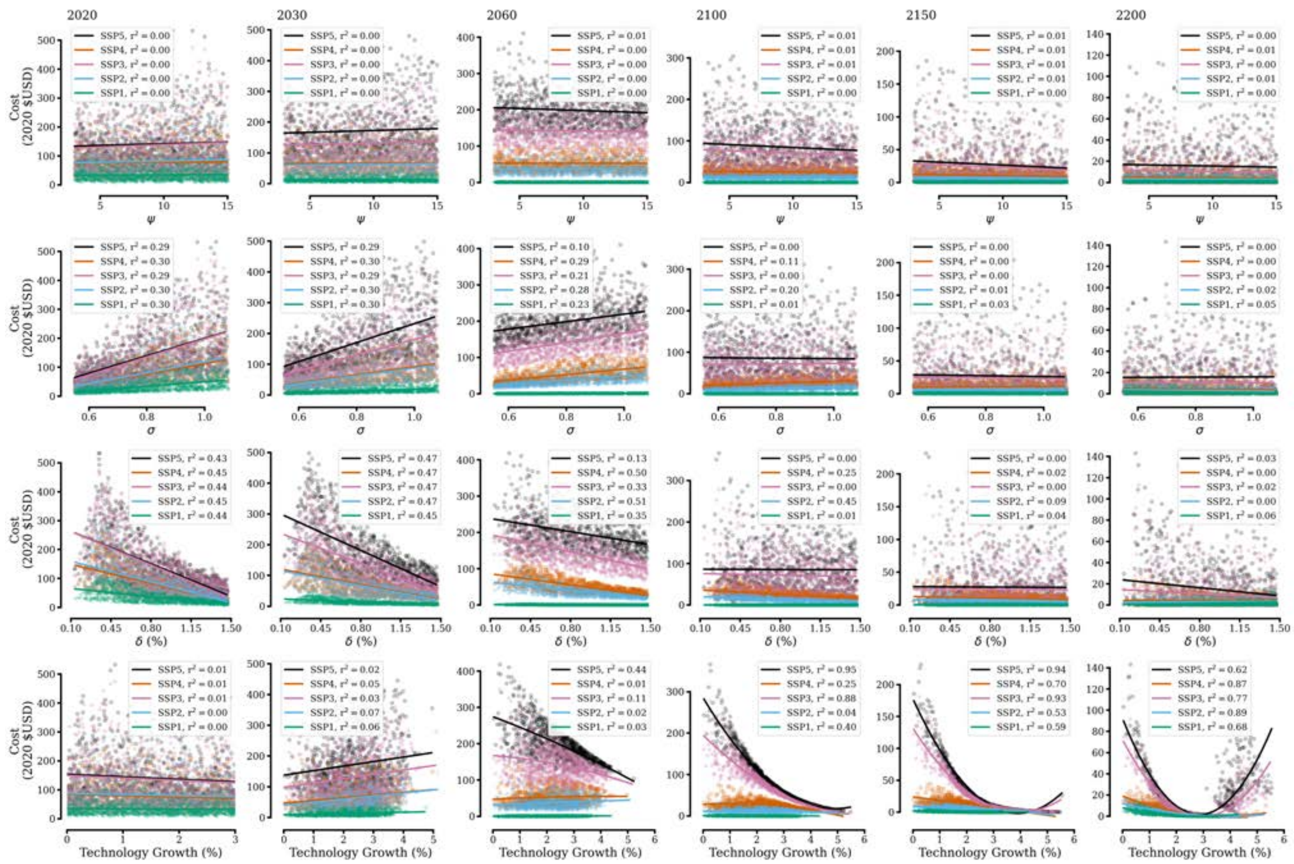

Figure 4: In each row, we plot the regression of each parameter against carbon costs in that period. r^2 values are given for each regression in the legend of each panel.

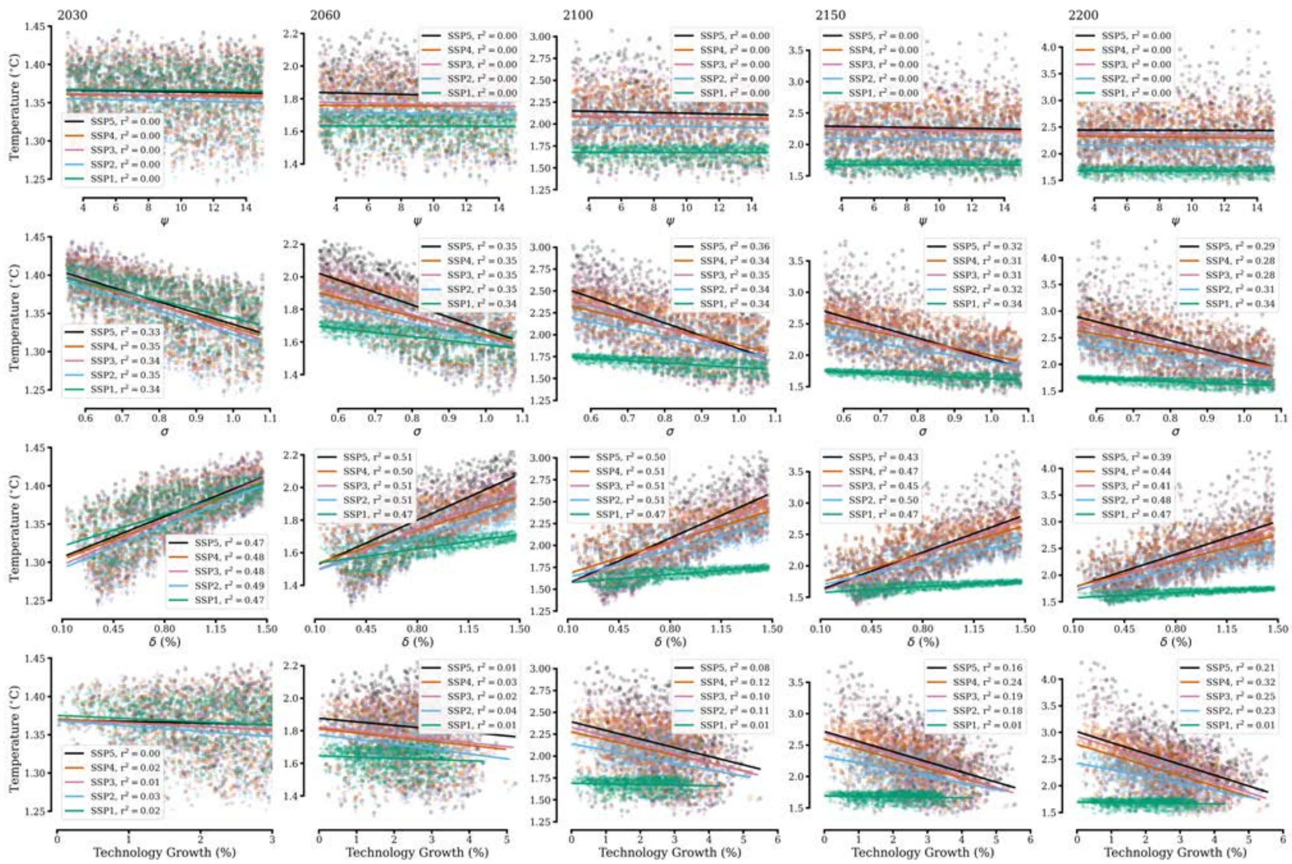

Figure 5: In each row, we plot the regression of each parameter against temperature in that period. r^2 values are given for each regression in the legend of each panel.

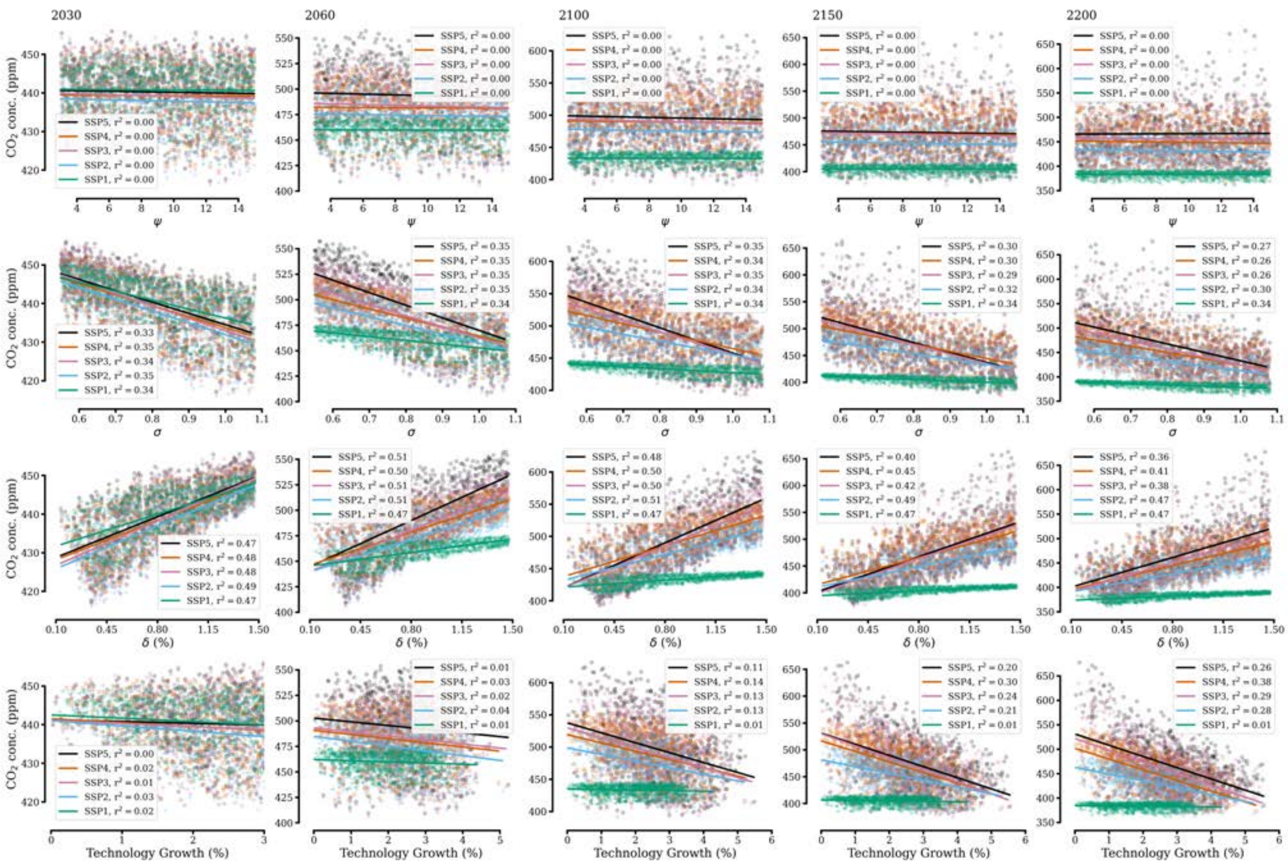

Figure 6: In each row, we plot the regression of each parameter against CO_2 concentrations in that period. r^2 values are given for each regression in the legend of each panel.

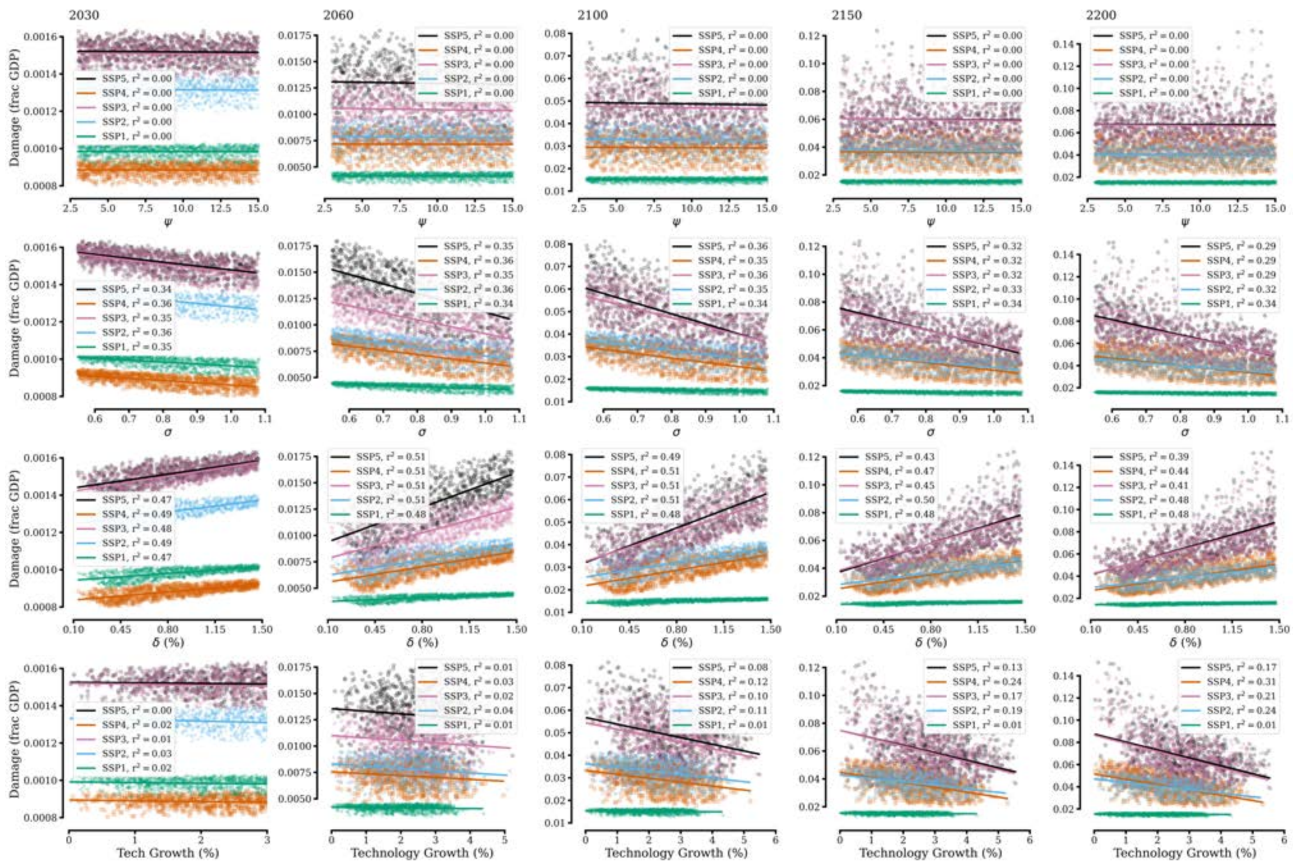

Figure 7: In each row, we plot the regression of each parameter against economic damages in that period. r^2 values are given for each regression in the legend of each panel.

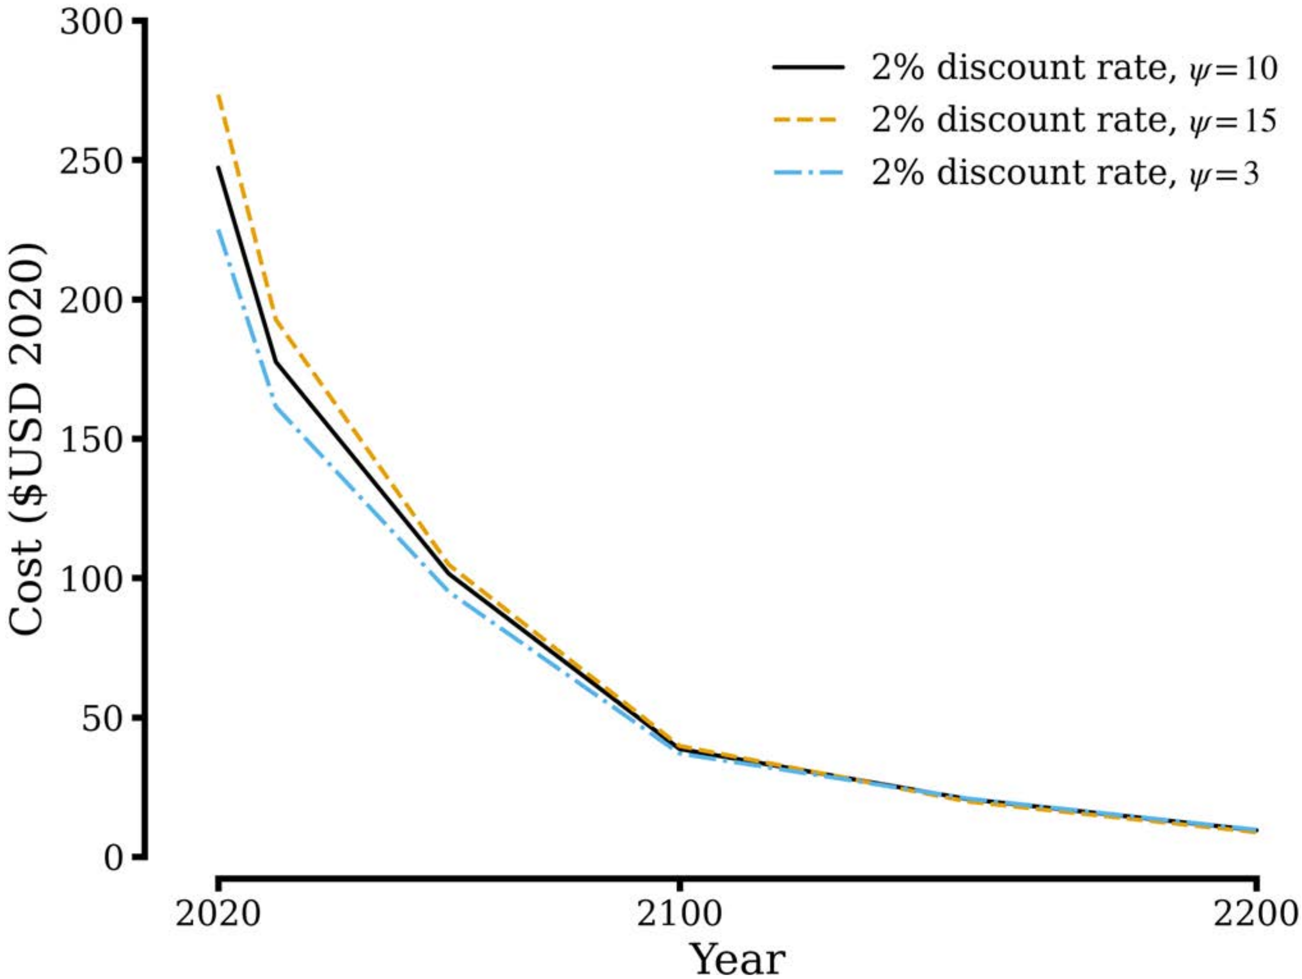

Figure 8: Shown is the resulting price path for different choices of risk aversion, holding all other model inputs constant in our preferred calibration.

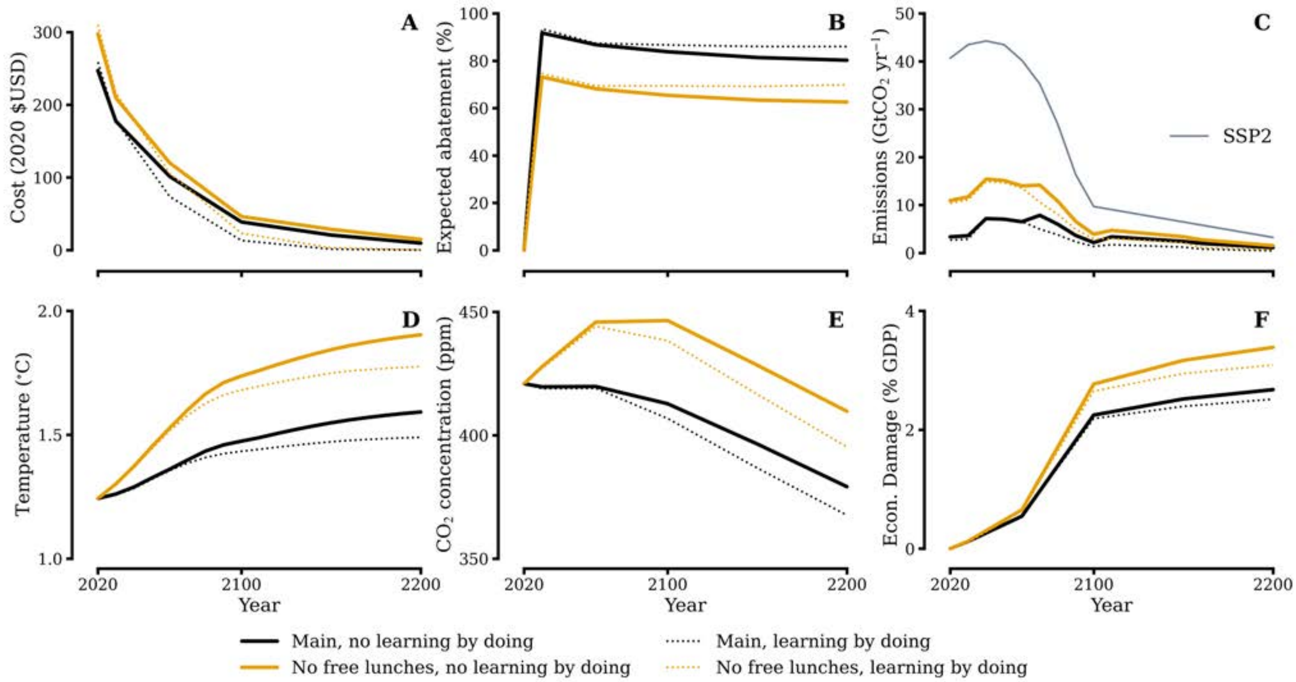

Figure 9: We show model output using our preferred 2% discount rate and toggling which MACC we use ('main' or "no free lunches") with or without learning by doing.

Note: Learning by doing implies that $\varphi_1 = 1.5\%$; no learning by doing corresponds to $\varphi_1 = 0\%$. All other parameters are the same as in our main specification.

References

M. Barnett, W. Brock, and L. P. Hansen. Pricing Uncertainty Induced by Climate Change. *The Review of Financial Studies*, 33(3):1024–1066, Mar. 2020. ISSN 0893-9454, 1465-7368. doi: 10.1093/rfs/hhz144. URL https://academic.oup.com/rfs/article/33/3/1024/5735312.

L. Barrage and W. Nordhaus. Policies, Projections, and the Social Cost of Carbon: Results from the DICE-2023 Model. Technical Report w31112, National Bureau of Economic Research, Cambridge, MA, Apr. 2023. URL http://www.nber.org/papers/w31112.pdf.

M. Burke, S. M. Hsiang, and E. Miguel. Global non-linear effect of temperature on economic production. *Nature*, 527(7577):235–239, Nov. 2015. ISSN 0028-0836, 1476-4687. doi: 10.1038/nature15725. URL http://www.nature.com/articles/nature15725.

M. Burke, W. M. Davis, and N. S. Diffenbaugh. Large potential reduction in economic damages under UN mitigation targets. *Nature*, 557(7706):549–553, May 2018. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-018-0071-9. URL http://www.nature.com/articles/s41586-018-0071-9.

Y. Cai and T. S. Lontzek. The Social Cost of Carbon with Economic and Climate Risks. *Journal of Political Economy*, 127(6):2684–2734, Dec. 2019. ISSN 0022-3808, 1537-534X. doi: 10.1086/701890. URL https://www.journals.uchicago.edu/doi/10.1086/701890.

T. A. Carleton and S. M. Hsiang. Social and economic impacts of climate. *Science*, 353(6304):aad9837, Sept. 2016. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.aad9837. URL https://www.science.org/doi/10.1126/science.aad9837.

Committee on Assessing Approaches to Updating the Social Cost of Carbon, Board on Environmental Change and Society, Division of Behavioral and Social Sciences and Education, and National Academies of Sciences, Engineering, and Medicine. *Valuing Climate Changes: Updating Estimation of the Social Cost of Carbon Dioxide*. National Academies Press, Washington, D.C., 2017. ISBN 9780309454209. doi: 10.17226/24651. URL https://www.nap.edu/catalog/24651.

H. Damon Matthews, K. B. Tokarska, J. Rogelj, C. J. Smith, A. H. MacDougall, K. Haustein, N. Mengis, S. Sippel, P. M. Forster, and R. Knutti. An integrated approach to quantifying uncertainties in the remaining carbon budget. *Communications Earth & Environment*, 2(1):7, Jan. 2021. ISSN 2662-4435. doi: 10.1038/s43247-020-00064-9. URL https://www.nature.com/articles/s43247-020-00064-9.

S. Dietz, C. Gollier, and L. Kessler. The climate beta. *Journal of Environmental Economics and Management*, 87:258–274, Jan. 2018. ISSN 00950696. doi: 10.1016/j.jeem.2017.07.005. URL https://linkinghub.elsevier.com/retrieve/pii/S0095069617304734.

S. Dietz, J. Rising, T. Stoerk, and G. Wagner. Economic impacts of tipping points in the climate system. *Proceedings of the National Academy of Sciences*, 118(34):e2103081118, Aug. 2021. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.2103081118. URL https://pnas.org/doi/full/10.1073/pnas.2103081118.

D. E. Goldberg. *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley Pub. Co, Reading, Mass, 1989. ISBN 9780201157673.

M. Golosov, J. Hassler, P. Krusell, and A. Tsyvinski. Optimal Taxes on Fossil Fuel in General Equilibrium. *Econometrica*, 82(1):41–88, 2014. ISSN 0012-9682. doi: 10.3982/ECTA10217. URL http://doi.wiley.com/10.3982/ECTA10217.

L. H. Goulder. Environmental taxation and the double dividend: A reader's guide. *International Tax and Public Finance*, 2(2):157–183, Aug. 1995. ISSN 0927-5940, 1573-6970. doi: 10.1007/BF00877495. URL http://link.springer.com/10.1007/BF00877495.

M. Ha-Duong, M. J. Grubb, and J.-C. Hourcade. Influence of socioeconomic inertia and uncertainty on optimal CO_2-emission abatement. *Nature*, 390(6657):270–273, Nov. 1997. ISSN 0028-0836, 1476-4687. doi: 10.1038/36825. URL http://www.nature.com/articles/36825.

C. Hambel, H. Kraft, and E. Schwartz. Optimal carbon abatement in a stochastic equilibrium model with climate change. *European Economic Review*, 132:103642, Feb. 2021. ISSN 00142921. doi: 10.1016/j.euroecorev.2020.103642. URL https://linkinghub.elsevier.com/retrieve/pii/S0014292120302725.

W. W. Hogan and D. W. Jorgenson. Productivity trends and the cost of reducing CO_2 emissions. *The Energy Journal*, 12(1):67–85, 1991. URL https://www.jstor.org/stable/41322403.

C. Hope, J. Anderson, and P. Wenman. Policy analysis of the greenhouse effect. *Energy Policy*, 21(3):327–338, Mar. 1993. ISSN 03014215. doi: 10.1016/0301-4215(93)90253-C. URL https://linkinghub.elsevier.com/retrieve/pii/030142159390253C.

P. H. Howard and T. Sterner. Few and Not So Far Between: A Meta-analysis of Climate Damage Estimates. *Environmental and Resource Economics*, 68(1):197–225, Sept. 2017. ISSN 0924-6460, 1573-1502. doi: 10.1007/s10640-017-0166-z. URL http://link.springer.com/10.1007/s10640-017-0166-z.

Intergovernmental Panel on Climate Change. Climate change 2021: The physical science basis. 2021. URL https://report.ipcc.ch/ar6wg1/pdf/IPCC_AR6_WGI_FinalDraft_FullReport.pdf.

Intergovernmental Panel on Climate Change. Climate change 2022: Impacts, adaptation, and vulnerability. 2022. URL https://report.ipcc.ch/ar6wg2/pdf/IPCC_AR6_WGII_FinalDraft_FullReport.pdf.

F. Joos, R. Roth, J. S. Fuglestvedt, G. P. Peters, I. G. Enting, W. von Bloh, V. Brovkin, E. J. Burke, M. Eby, N. R. Edwards, T. Friedrich, T. L. Frölicher, P. R. Halloran, P. B. Holden, C. Jones, T. Kleinen, F. T. Mackenzie, K. Matsumoto, M. Meinshausen, G.-K. Plattner, A. Reisinger, J. Segschneider, G. Shaffer, M. Steinacher, K. Strassmann, K. Tanaka, A. Timmermann, and A. J. Weaver. Carbon dioxide and climate impulse response functions for the computation of greenhouse gas metrics: a multi-model analysis. *Atmospheric Chemistry and Physics*, 13(5):2793–2825, Mar. 2013. ISSN 1680-7324. doi: 10.5194/acp-13-2793-2013. URL https://acp.copernicus.org/articles/13/2793/2013/.

D. W. Jorgenson. *Double dividend: environmental taxes and fiscal reform in the United States*. MIT Press, Cambridge, MA, 2013. ISBN 9780262027090.

D. Lemoine. The Climate Risk Premium: How Uncertainty Affects the Social Cost of Carbon. *Journal of the Association of Environmental and Resource Economists*, 8(1):27–57, Jan. 2021. ISSN 2333-5955, 2333-5963. doi: 10.1086/710667. URL https://www.journals.uchicago.edu/doi/10.1086/710667.

D. Lemoine and C. P. Traeger. Ambiguous tipping points. *Journal of Economic Behavior & Organization*, 132:5–18, Dec. 2016a. ISSN 01672681. doi: 10.1016/j.jebo.2016.03.009. URL https://linkinghub.elsevier.com/retrieve/pii/S0167268116300221.

D. Lemoine and C. P. Traeger. Economics of tipping the climate dominoes. *Nature Climate Change*, 6(5):514–519, May 2016b. ISSN 1758-678X, 1758-6798. doi: 10.1038/nclimate2902. URL https://www.nature.com/articles/nclimate2902.

R. E. Lucas. Asset Prices in an Exchange Economy. *Econometrica*, 46(6):1429–1445, 1978. ISSN 0012-9682. doi: 10.2307/1913837. URL https://www.jstor.org/stable/1913837.

N. G. Mankiw, M. Weinzierl, and D. Yagan. Optimal Taxation in Theory and Practice. *Journal of Economic Perspectives*, 23(4):147–174, Nov. 2009. ISSN 0895-3309. doi: 10.1257/jep.23.4.147. URL https://pubs.aeaweb.org/doi/10.1257/jep.23.4.147.

F. C. Moore and D. B. Diaz. Temperature impacts on economic growth warrant stringent mitigation policy. *Nature Climate Change*, 5(2):127–131, Feb. 2015. ISSN 1758-678X, 1758-6798. doi: 10.1038/nclimate2481. URL http://www.nature.com/articles/nclimate2481.

W. D. Nordhaus. An Optimal Transition Path for Controlling Greenhouse Gases. *Science*, 258(5086):1315–1319, Nov. 1992. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.258.5086.1315. URL https://www.sciencemag.org/lookup/doi/10.1126/science.258.5086.1315.

W. D. Nordhaus. Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, 114(7):1518–1523, Feb. 2017. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1609244114. URL https://pnas.org/doi/full/10.1073/pnas.1609244114.

R. G. Richels and G. J. Blanford. The value of technological advance in decarbonizing the U.S. economy. *Energy Economics*, 30(6):2930–2946, Nov. 2008. ISSN 01409883. doi: 10.1016/j.eneco.2008.06.005. URL https://linkinghub.elsevier.com/retrieve/pii/S014098830800087X.

S. K. Rose, D. B. Diaz, and G. J. Blanford. Understanding the Social Cost of Carbon: A Model Diagnostic and Inter-comparison Study. *Climate Change Economics*, 08(02):1750009, May 2017. ISSN 2010-0078, 2010-0086. doi: 10.1142/S2010007817500099. URL https://www.worldscientific.com/doi/abs/10.1142/S2010007817500099.

R. Tol. Safe policies in an uncertain climate: an application of FUND. *Global Environmental Change*, 9(3):221–232, Oct. 1999. ISSN 09593780. doi: 10.1016/S0959-3780(99)00011-4. URL https://linkinghub.elsevier.com/retrieve/pii/S0959378099000114.

T. S. Van Den Bremer and F. Van Der Ploeg. The Risk-Adjusted Carbon Price. *American Economic Review*, 111(9):2782–2810, Sept. 2021. ISSN 0002-8282. doi: 10.1257/aer.20180517. URL https://pubs.aeaweb.org/doi/10.1257/aer.20180517.

A. Vogt-Schilb, G. Meunier, and S. Hallegatte. When starting with the most expensive option makes sense: Optimal timing, cost and sectoral allocation of abatement investment. *Journal of Environmental Economics and Management*, 88:210–233, Mar. 2018. ISSN 00950696. doi: 10.1016/j.jeem.2017.12.001. URL https://linkinghub.elsevier.com/retrieve/pii/S0095069617308392.