Risk, Inequality, and Climate Change

Joshua Bernstein* Francis Dennig† Maddalena Ferranna‡

Marc Fleurbaey§

August 31, 2022

Abstract

We study risk and inequality in an integrated assessment model, and examine their joint impact on optimal carbon taxation. We provide a simple decomposition of the carbon tax, and we show analytically that the interaction of risk and inequality increases the optimal carbon tax when ex-ante high consumption regions are less exposed to climate change damages. Under a standard calibration of global damage risk, the presence of unequal damage incidence calls for a \$117 per ton increase in the optimal carbon tax. 10% of this increase is driven by the interaction of risk and inequality.

Keywords Risk; Inequality; Climate Change; Carbon Price; Integrated Assessment Model

*Indiana University. jmbernst@iu.edu.

†Yale-NUS College.

‡University of Southern California, ferranna@usc.edu

§Paris School of Economics, marc.fleurbaey@psemail.eu

1 Introduction

Carbon taxation forms the backbone of policy designed to combat the effects of climate change. As such, the design of optimal taxes must overcome two key challenges. First, there is great uncertainty surrounding the pace and magnitude of climate change impacts, so that carbon tax policy is subject to risk concerns. Second, these damage impacts are unlikely to be spread evenly among the world population, so that carbon tax policy is subject to inequality concerns.[1] In light of these joint risk and inequality concerns, it is important to understand their impacts, both separately and together, on the determination of optimal carbon taxation.

In this paper, we model the interaction of risk and inequality, and study their joint impact on optimal carbon taxation. In order to study risk and inequality simultaneously, we extend the integrated assessment RICE model (Nordhaus (2010)) in 2 dimensions.[2] First, we follow Dennig et al. (2015) and allow for consumption inequality within each global sub-region that itself depends on the regional level of climate damages. This dependence captures the idea that climate change can affect inequality directly. Second, we attach uncertainty to various structural model parameters and specify that taxes must be set before the uncertainty is resolved in order to capture the effects of risk on optimal carbon taxation.

Using this model, our first contribution is to provide an analytical decomposition that sheds light on the key forces that determine the optimal path of carbon taxes. Intuitively, we show that risk and inequality each call for higher carbon taxes since the planner is both risk and inequality averse.

Our decomposition emphasizes how the interaction of parameter uncertainty and inequality may lead to even higher carbon taxes than the sum of their separate effects would imply.

[1]For example, recent work by the World Bank (Hallegatte et al. (2015)) finds that poor populations are more likely to be vulnerable to weather and health shocks caused by climate change.

[2]See http://www.econ.yale.edu/ nordhaus/homepage/RICEmodels.htm for documentation. RICE is a disaggregation of the DICE model into twelve regions, but does not account for inequalities between rich and poor within regions.

This "super-additivity" property follows from the fact that parameter uncertainty induces risk not only over aggregate consumption, but also over its distribution across individuals.

Our decomposition shows that this distribution risk manifests in 2 ways. First, if parameter uncertainty increases the variance of the global consumption distribution relative to the case without risk, then the optimal carbon taxes are higher. Intuitively, a higher average variance across states of parameter uncertainty implies that individuals with high consumption gain the most when aggregate consumption is higher, but lose the least when aggregate consumption is lower. Higher carbon taxation therefore provides additional insurance against this unequal exposure to aggregate risk.

Second, carbon taxes are higher if parameter uncertainty lowers the covariance between individual consumption and the benefits of carbon taxation (through lower damages) relative to the case without risk. Intuitively, a lower covariance implies that individuals with high consumption gain the less when aggregate benefits are larger, and lose more when aggregate benefits are smaller. Higher carbon taxation serves to exploit this pattern of redistribution through a reduction in climate damages.

With these analytical insights in mind, our second contribution is to investigate the effects of risk and inequality in a quantitative sense. As expected, we find that risk and inequality each cause optimal taxes to increase. For example, under a standard calibration of damage risk (Nordhaus (2018)), we find that the optimal carbon tax increases from \$35 in the absence of risk and inequality to \$167 in the presence of both. Furthermore, we show that 17% of this increase is driven by the joint interaction of risk and inequality, which provides quantitative evidence for the "super-additivity" property we illuminated analytically. We confirm that this result is driven by the effects described above. For example, when poorer individuals are disproportionally exposed to climate damages, states of nature with higher damages exacerbate inequality, while states of nature with lower damages do not decrease inequality enough to fully offset the risk in higher damage states.

Related Literature. Our work relates to two broad strands of the literature on stochastic integrated assessment models. First, we build on work that studies how structural risk affects optimal mitigation policies. See Jensen and Traeger (2014) for growth risk, Daniel et al. (2016), Crost and Traeger (2014), and Cai et al. (2013) for damage uncertainty, Jensen and Traeger (2021), Kelly and Tan (2015), Hwang et al. (2017), Leach (2007), and Kelly and Kolstad (1999) for climate sensitivity, and Lemoine (2021), Lontzek et al. (2015) for tipping points. Relative to these papers, our key contribution is to consider how risk permeates in a more disaggregated model, with both regional and sub-regional income inequality, and to study the interaction of risk with unequal damage incidence of climate impacts inside a given region.

Second, our theoretical decomposition of the optimal carbon price builds on the literature analyzing the effects of correlation between exogenous risks and consumption risk on optimal mitigation policy (Howarth (2003); Sandsmark and Vennemo (2007); Dietz et al. (2018); Lemoine and Traeger (2016). Our disaggregated approach extends this literature by examining the additional effect of correlation between exogenous risk and unequal damage incidence on the optimal carbon price.

The paper proceeds as follows: Section 2 introduces the model and explains how we model both risk and unequal damage incidence. In section 3 we derive decompositions that highlight the key forces determining optimal carbon taxes and how they differ in the presence of different combinations of risk and unequal damage incidence. These decompositions inform interpretation of the numerical exercises presented in section 4. Section 5 concludes.

2 The Model Economy

Our model extends the RICE model of Nordhaus (2010) to allow for both parameter uncertainty and within-region inequality. We call this variant of integrated assessment models the

NICER model (Nested Inequality Climate Economy model with Risk).[3] Given this, we now provide a concise exposition of the key features of the model, with further details provided in appendix B.[4]

2.1 Parameter Uncertainty

The bulk of our analysis concerns a planner who must choose paths for the carbon tax, and global allocation of resources subject to ex-ante uncertainty over various model parameters that we describe below. To this end, let uncertainty be represented by a set of states of nature $\mathbb{S} = \{1, .., S\}$ with generic element s and corresponding probability π_s, where $\sum_{s \in \mathbb{S}} \pi_s = 1$. These probabilities reflect the information available to the planner when she makes her policy and allocation choices. For example, the planner may not know for sure what the future growth rate of the economy will be. To capture this in a simple manner, we can set $\mathbb{S} = \{1, 2\}$, where $s = 1$ is a low growth state and $s = 2$ is a high growth state corresponding to different realized values for the relevant growth rate parameter in the model.

Given a state of nature s, the economy lasts for $T + 1$ periods, where time is indexed by $t \in \{0, .., T\}$, and the planner makes her choices just before period $t = 0$ begins. In this sense, we do not allow the planner to learn about the model, or to re-optimize her choices given new information she may obtain.

2.2 Regional Heterogeneity

As in the original RICE model, there are R regions, where each region r is characterized by region-specific production and damage functions, which lead to heterogeneity in consumption levels across individuals living in different regions. Let Q_{rst}, Λ_{rst} and D_{rst} denote,

[3]The label comes from its ancestor NICE, which featured consumption disaggregation both between and within regions, but no risk (Dennig et al., 2015).

[4]Further information on the equations of the RICE model can be found in the supplementary materials of Nordhaus (2010)

respectively, gross output, mitigation costs and damages in region r, period t, and state s. Then, net output Y_{rst} is defined as:

$$Y_{rst} = \frac{1 - \Lambda_{rst}}{1 + D_{rst}} Q_{rst} \tag{1}$$

Net output is used either for consumption or for savings (i.e. investment in capital). Let S_{rst} be the (endogenous) regional savings rate. Then, regional aggregate consumption can be written as:

$$C_{rst} = (1 - S_{rst})Y_{rst}$$

where $S_{rst}Y_{rst}$ is the output invested in capital, which will contribute to the next generation production level Q_{rst+1}. Finally, let L_{rt} denote the number of individuals living in region r at time t, where we have used the fact that population is deterministic to suppress the dependence on states of the world.

In equilibrium, the evolution of these aggregate variables is jointly determined with the evolution of climate variables such as atmospheric temperature, as outlined in appendix B. Of particular importance, is the dependence of damages on temperature T_{rst},

$$D_{rst} = \psi_{rs} T_{rst}^2 \tag{2}$$

where the coefficients $\{\psi_{rs}\}$ are allowed to depend on the state s thus introducing uncertainty to the damages inflicted by increases in temperature caused by climate change.

2.3 Within-Region Inequality

In order to model within-region inequality, we build on Dennig et al. (2015). Specifically, we divide the population of region r at time t in state s into quintiles and define c_{irst} as the per-capita consumption of individuals in quintile i, region r, at time t, and in state s. We model the distribution of aggregate regional consumption across quintiles using the following equation:

$$c_{irst} = C_{rst}\Big[q_{ir}(1 + D_{rst}) - d_{ir}D_{rst})\Big]\frac{5}{L_{rt}} \tag{3}$$

where q_{ir} and d_{ir} have the following interpretation.

First, assume that there are no climate damages so that $D_{rst} = 0$. Then $c_{irst} = C_{rst}q_{ir}\frac{5}{L_{rt}}$, which means that q_{ir} is the fraction of aggregate regional consumption allocated to quintile i, thus reflecting exogenous consumption inequality that would occur in the absence of climate change and its associated damages.[5] In the numerical exercises, we use the estimates of q_{ir} from Dennig et al. (2015), which are based on World Bank estimates of global income inequality.

Now, suppose that there are climate damages, and recall that $C_{rst}D_{rst}$ measures the total amount of consumption that is lost due to climate change. Given this, d_{ir} measures the fraction of aggregate consumption loss that falls on quintile i.

The relative sizes of d_{ir} and q_{ir} play an important role in determining post-climate damage consumption inequality: individuals in quintile i get a proportion q_{ir} of total, pre-damage, consumption, and pay a proportion d_{ir} of total damages. The sum of the two components determines the share of post-damage consumption that individuals in quintile i get to consume. For example, if individuals in quintile i consume 10% of aggregate consumption in the absence of climate change ($q_{ir} = 0.1$), but bear 20% of aggregate climate damages ($d_{ir} = 0.2$), their final share of aggregate consumption will be smaller than 10%, with the size of the gap being an increasing function of total damages.

While there is ample evidence that climate damages are likely to be unequally distributed across the global population, there is little evidence to guide our numerical choices for d_{ir}. Therefore, in order to study a range of possible outcomes in a parsimonious manner, we

[5] The consumption allocated to quintile i, $C_{rst}q_{ir}$, is then equally divided among the number of individuals in quintile i, $\frac{L_{rt}}{5}$.

follow Dennig et al. (2015), and posit an explicit relationship between q_{ir} and d_{ir}:

$$d_{ir} = \frac{q_{ir}^{\xi}}{\sum_i q_{ir}^{\xi}} \tag{4}$$

In this formula, ξ is the income elasticity of damage, and determines the relationship between the pre-existing consumption distribution and the distribution of damages. For example, if $\xi = 1$, overall consumption inequality is unaffected by the presence of climate damages, as $d_{ir} = q_{ir}$ and damages are spread proportionally across quintiles. In contrast, if $\xi = 0$, then $d_{ir} = \frac{1}{5}$ regardless of the distribution q_{ir}, so that poorer quintiles lose a greater share of consumption due to climate damage then richer quintiles.

In general, $\xi < 1$ implies that the poorest bear a larger share of climate damages, so that climate change increases consumption inequality. In contrast, when $\xi > 1$ the rich pay the most of damages, and climate change decreases existing inequality.

2.4 Planning Objective

Our model economy is highly dis-aggregated, featuring 60 region-quintile consumption levels in each period and state of nature. In order to compare different time-state paths of region-quintile consumption levels implied by a choice of carbon taxes $\vec{\tau} = (\tau_0, \ldots, \tau_T)$, we use the objective

$$W(\vec{\tau}) = \mathbb{E} \sum_{t,r,i} \beta^t \frac{L_{rt}}{5} \left(\frac{c_{irst}(\vec{\tau})^{1-\eta}}{1-\eta} \right) \tag{5}$$

where β is the preference discount factor, η the coefficient of both risk and inequality aversion of the planner, and $\mathbb{E}$ the expectation operator associated with the probability distribution $\{\pi_s\}_{\forall s}$, such that, for any random variable x with realizations x_s, $\mathbb{E}x = \sum_s \pi_s x_s$. The dependence of both welfare and consumption on the vector of carbon taxes $\vec{\tau} = (\tau_0, \ldots, \tau_T)$

emphasizes that the planner chooses a path of taxes to maximize this measure of social welfare subject to technological and climate constraints.

In order to make our results easily comparable to the the bulk of the climate change policy literature, we restrict attention to the selection of a uniform global carbon tax that applies to all region-quintiles equally. This rules out the application of differentiated carbon taxes (Chichilnisky and Heal, 1994), which we leave for future research. [6]

A crucial feature of our analysis is understanding how the interaction of risk and inequality impacts the path of optimal carbon taxes. To isolate the effect of this interaction, we compare the optimal taxes using the previous objective to the optimal taxes chosen in an economy with the same treatment of parameter uncertainty, but in which there is no consumption inequality. In this economy, all individuals achieve the same level of consumption in state s and period t, c_{st}, which is evaluated by the analogous objective function

$$W^A(\vec{\tau}) = \sum_t \beta^t L_t \mathbb{E}\left[\frac{c_{st}(\vec{\tau})^{1-\eta}}{1-\eta}\right] \tag{6}$$

where $L_t \equiv \sum_r L_{rt}$ is the global population in period t, and the A stands for "aggregated model". We refer to the model with this objective function as DICER.[7] Comparing results from our NICER and DICER models then allows us to isolate the effects of parameter uncertainty (deterministic model DICE vs. stochastic version DICER), within-region inequality (deterministic model DICE vs. deterministic model NICE), and their interaction (stochastic model DICER vs. stochastic model NICER).

[6] In light of the current system of nationally determined contributions instituted by the Paris agreement, the mitigation pathway the world is likely to follow will be somewhere in between the extremes of complete uniformity and total differentiation.

[7] Technically, this is a slight abuse of terminology since the original DICE model did not feature differentiated regions, while our DICER model aggregates over these regions in order to achieve uniform per-capita consumption.

3 The Risk-and-Equity-adjusted Carbon Tax

Our goal in this paper is to understand how the interaction of parameter uncertainty and between/within-region inequality impacts the optimal path of carbon taxes. To this end, we now analytically decompose expressions for the optimal carbon taxes in both the DICER and NICER models. These decompositions shed light on the key forces that cause taxes to change in the presence of both risk and inequality.

In the following analysis, we will use an additional expectation operator, $\hat{\mathbb{E}}$, which denotes the average value of a variable with respect to all dimensions except time: $\hat{\mathbb{E}}x_{irst} = \frac{1}{5}\sum_{r,i}\frac{L_{rt}}{L_t}\mathbb{E}x_{irst}$, where $L_t = \sum_r L_{rt}$ is total population in period t. In addition, whenever we encounter a variable with a missing index, it means that we are considering the average value of the variable along that dimension. For example, $x_{irt} = \mathbb{E}x_{irst}$; $x_{rst} = \frac{1}{5}\sum_i x_{irst}$; $x_{ist} = \sum_r \frac{L_{rt}}{L_t}x_{irst}$; if two or more indexes are missing, then we are considering the average value of the variable along those dimensions, for example $x_{st} = \frac{1}{5}\sum_{r,i}\frac{L_{rt}}{L_t}x_{irst}$, and so on.

3.1 Optimal Carbon Taxation in DICER

We begin with the DICER model where the planning objective is given by (6). Details of all derivations are contained in appendix B.

Introducing a carbon tax creates both economic benefits and costs. The optimal tax is set such that the marginal benefit equals the marginal cost. Define λ_t as the per-capita marginal monetary cost of a marginal increase in the carbon tax in period t, τ_t:

$$\lambda_t \equiv \sum_r \frac{L_{rt}}{L_t}\frac{\Lambda'_{rt}}{1-\Lambda_{rt}}\frac{Y_{rt}}{L_{rt}}$$

where $-\frac{\Lambda'_{rt}}{1-\Lambda_{rt}}Y_{rt}$ is the total output loss for region r using (1), and $\Lambda'_{rt} = \frac{\partial \Lambda_{rt}}{\partial \tau_t}$ is the marginal impact of the tax on mitigation costs. This cost represents the monetary loss in period t, as production has to shift from carbon intensive technologies to less intensive ones due to the imposition of the tax.

The benefit of taxation arises from a reduction in emissions, which causes a reduction in the carbon stock and hence a smaller atmospheric temperature increase for all future periods. Let $T'_{sj} = \frac{\partial T_{sj}}{\partial \tau_t}$ be the marginal impact of a tax at t on the increase in temperature at time j and in the state of the world s. Changing T_{sj} affects output in that period, and, from (1), results in a monetary regional benefit equal to $-\frac{D'_{rsj}}{1+D_{rsj}} T'_{sj} Y_{rsj}$, with $D'_{rsj} = \frac{\partial D_{rsj}}{\partial T_{sj}}$. Define δ_{sj} as the per-capita monetary benefit in period j and state s of a marginal increase in the tax τ_t:

$$\delta_{sj} \equiv -T'_{sj} \sum_r \frac{L_{rj}}{L_j} \frac{D'_{rsj}}{1+D_{rsj}} \frac{Y_{rsj}}{L_{rj}} \qquad \forall j \geq t$$

This monetary benefit at j is valued according to its marginal impact on welfare at j and in state s, $c_{sj}^{-\eta}\delta_{sj}$. Finally, this welfare change has to be evaluated in monetary units at period t, so as to make it comparable to the monetary loss at t induced by the policy. Therefore, we divide it by the marginal consumption value at t, $c_t^{-\eta}$. The total return of the policy, at the margin, is the value at t of the sum of future per-capita monetary benefits δ_{sj}.

The optimal tax is chosen such that, at the margin, the monetary mitigation costs are equal to the total discounted monetary benefits (see Appendix B):[8]

$$\lambda_t \simeq \sum_{j>t} \beta^{j-t} \frac{L_j}{L_t} \left(\frac{\mathbb{E}[c_{sj}^{-\eta}\delta_{sj}]}{c_t^{-\eta}\delta_j} \right) \delta_j \tag{7}$$

where $\delta_j \equiv \mathbb{E}[\delta_{sj}]$ is the expected marginal per-capita benefit of the policy at j. From (7), we can see how the value at t of the sum of future risky benefits δ_{sj} can be rewritten as the discounted sum of future *expected* benefits δ_j, with an associated discount factor $DF_j^A = \beta^{j-t}\frac{L_j}{L_t}\left(\frac{\mathbb{E}[c_{sj}^{-\eta}\delta_{sj}]}{c_t^{-\eta}\delta_j}\right)$.

[8] The 'approximately equal' sign comes from the fact that the policy may be chosen ex-ante, i.e. before risk at time t is resolved. As a consequence, marginal mitigation costs at t may be subject to risk.

The discount factor measures the value in period t of \$1 received in a future period $j > t$. The larger the discount factors are in future periods, the larger is the present value of future marginal monetary benefits from carbon taxation today. Therefore, the discount factors encode the key determinants of the optimal carbon tax path.

In order to highlight these determinants, we can take a second-order Taylor approximation of the discount factor around the expected consumption c_j and expected damages δ_j, for any $j > t$:

$$\ln DF_j^A \simeq (j-t)\ln\beta + \ln\frac{L_j}{L_t} \underbrace{-\eta\ln\frac{c_j}{c_t}}_{\text{growth effect}} + \underbrace{\frac{1}{2}\eta(\eta+1)V\left(\frac{c_{sj}}{c_j}\right)}_{\text{precautionary effect}} \underbrace{-\eta Cov\left(\frac{c_{sj}}{c_j},\frac{\delta_{sj}}{\delta_j}\right)}_{\text{risk premium}} \quad (8)$$

where $V\left(\frac{c_{sj}}{c_j}\right)$ denotes the variance of consumption across the states of nature (normalized by average consumption across the states), and $Cov\left(\frac{c_{sj}}{c_j},\frac{\delta_{sj}}{\delta_j}\right)$ the covariance between normalized consumption and the normalized marginal benefits of the policy.

The discount factor is driven by 5 distinct forces. First, a stronger level of patience (higher β) naturally causes a higher discount factor. Similarly, a higher relative population size creates a higher discount factor since there are more people that need protecting in the future, which raises the gains from carbon taxation today.

Conversely, the term labeled "growth effect" shows that, ceteris paribus, higher future consumption growth creates a lower discount factor, and hence less taxation today. Intuitively, if we expect future consumption to be higher for exogenous reasons (e.g. technological growth), then there is less need to impose high carbon taxation today. Furthermore, the strength of this channel is mediated by η in its role as the reciprocal of the elasticity of intertemporal substitution in the time dimension: higher η implies that the planner would like a smoother consumption path over time, which is achieved via less taxation today given that future consumption is higher.

The last two terms in (8) identify the impact of risk on the discount factor. The term labeled "precautionary effect" captures the fact that parameter uncertainty creates risk on future consumption summarized by the variance term. This risk creates a higher discount factor, causing the planner to increase carbon taxation today as an insurance policy against this risk. Intuitively, the strength of this channel depends on the coefficient of relative prudence $(\eta + 1)$: higher prudence implies a stronger desire for insurance against future risk.

Finally, the term labeled "risk premium" captures the fact that parameter uncertainty may affect future consumption and future marginal benefits of carbon taxation differently, as summarized by the covariance term. In particular, a positive covariance results in a negative risk premium, a lower discount factor, and hence less carbon taxation today.[9] Intuitively, when future consumption and policy benefits are positively correlated, there is less payoff to taxing carbon today, since the future benefits of taxation realize in states of nature in which consumption is already higher. Therefore, taxation serves only to widen the gap between consumption in good states and bad states respectively. The opposite is true for a negative correlation, in which case taxation today serves to mitigate against future consumption risk.

The sign of the covariance term depends on the type of parameter uncertainty that we are considering. For example, uncertainties surrounding technological parameters (such as the rate of TFP growth and rate of the TFP convergence) induce a positive covariance between future consumption and future benefits of taxation, while uncertainty on parameters governing the specification of damages in the model (such as climate sensitivity and damage function coefficients) tend to induce a negative covariance (Dietz et al., 2018). These results will play out in our numerical simulations.

[9] The terminology derives from the consumption-based asset pricing literature (?). If the returns of an asset covary positively with consumption, expected returns must be higher to persuade an investor to hold the asset, i.e. they must feature a 'risk premium'.

3.2 Optimal Carbon Taxation in NICER

We now consider the optimal path of carbon taxes in the disaggregated NICER model, which has objective (5).

As in the DICER model, the introduction of a carbon tax in period t will impose some monetary costs for the individuals of generation t due to the forced reduction in emissions, and some monetary benefits for the individuals living in future generations thanks to the deceleration of global warming and the resulting reduction in climate damages. The optimal tax is such that, at the margin, total mitigation costs are equal to the value at t of the sum of future monetary benefits. The main difference with respect to the DICER model is the definition of 'individual': the social planner now cares about the entire distribution of marginal mitigation costs and marginal mitigation benefits, not just about their average values.

Given the definition of quintile-specific consumption (3), the marginal mitigation costs of individuals in quintile i, region r and generation t are given by

$$\lambda_{irt} \equiv \frac{\Lambda'_{rt}}{1-\Lambda_{rt}} Y_{rt} \left[q_{ir}(1+D_{rt}) - d_{ir} D_{rt} \right] \frac{5}{L_{rt}}$$

with $\Lambda'_{rt} = \frac{\partial \Lambda_{rt}}{\partial \tau_t}$, while the marginal mitigation benefits for future individuals in quintile i, region r, state of the world s and generation $j > t$ are equal to:

$$\delta_{irsj} \equiv -T'_{sj} \frac{D'_{rsj}}{1+D_{rsj}} Y_{rsj} d_{ir} \frac{5}{L_{rj}}$$

where $T'_{sj} = \frac{\partial T_{sj}}{\partial \tau_t}$ and $D'_{rsj} = \frac{\partial D_{rsj}}{\partial T_{sj}}$. Note that, while the distribution of marginal mitigation costs depend on the post-damage distribution of consumption at t (identified by the term in square brackets), the distribution of marginal benefits depend exclusively on the damage share of each quintile, d_{ir}. The monetary returns at j for individuals characterized by the triplet (i, r, s) are evaluated according to their marginal welfare impact at j, $c_{irsj}^{-\eta}\delta_{irsj}$, and then translated into monetary values at period t by dividing them by the marginal consumption value at t. Since there is inequality at t, the marginal consumption value will be equal to $\frac{1}{5}\sum_{r,i}\frac{L_{rt}}{L_t}c_{irt}^{-\eta}$. Furthermore, since the mitigation costs are unequally distributed, the incentives to set a high carbon tax will be sensitive also to the progressivity or regressitivy of the tax itself, i.e. to which quintiles are going to pay most of it. At the optimum, the carbon tax satisfies the following condition (see Appendix B for a detailed description of the solution):

$$\lambda_t \simeq \sum_{j>t}\beta^{j-t}\frac{L_j}{L_t}\left(\frac{\hat{\mathbb{E}}\left[c_{irsj}^{-\eta}\delta_{irsj}\right]}{\hat{\mathbb{E}}\left[c_{irt}^{-\eta}\right]\delta_j}\right)\left(\frac{\hat{\mathbb{E}}\left[c_{irt}^{-\eta}\right]\hat{\mathbb{E}}[\lambda_{irt}]}{\hat{\mathbb{E}}\left[c_{irt}^{-\eta}\lambda_{irt}\right]}\right)\delta_j \tag{9}$$

where $\lambda_t = \hat{\mathbb{E}}\lambda_{irt}$, $\delta_j = \hat{\mathbb{E}}\delta_{irsj}$, and $\hat{\mathbb{E}}$ computes the average value of a variable along all its dimensions i, r, s (for example, $\hat{\mathbb{E}}\delta_{irsj} = \frac{1}{5}\sum_{i,r}\frac{L_{rj}}{L_j}\mathbb{E}\delta_{irsj}$). Note that the definition of λ_t and δ_j is consistent with those in the DICER model, as in that case, $\hat{\mathbb{E}}$ and $\mathbb{E}$ coincide by definition, since there is no inequality.

Condition (9) states that the marginal aggregate costs of the policy should be equal to the discounted value of future marginal aggregate benefits, where the discount factor DF_j to be used for benefits accruing at j is inequality-and-risk dependent, $DF_j = \beta^{j-t}\frac{L_j}{L_t}\left(\frac{\hat{\mathbb{E}}\left[c_{irsj}^{-\eta}\delta_{irsj}\right]}{\hat{\mathbb{E}}\left[c_{irt}^{-\eta}\right]\delta_j}\right)\left(\frac{\hat{\mathbb{E}}\left[c_{irt}^{-\eta}\right]\hat{\mathbb{E}}[\lambda_t]}{\hat{\mathbb{E}}\left[c_{irt}^{-\eta}\lambda_{irt}\right]}\right)$. It is crucial to note that, with respect to the DICER case, inequality will affect only the evaluation of future policy returns via changes in the discount factor. In other words, for a given tax and emission path, inequality will not affect the marginal expected costs λ_t and benefits δ_j. Therefore, all the additional effects of inequality can be studied by comparing the discount factor to the discount factor in the DICER model.

By taking a second-order Taylor approximation of the discount factor around average expected consumption $c_j = \hat{\mathbb{E}}[c_{irsj}]$, average expected damages $\delta_j = \hat{\mathbb{E}}[\delta_{irsj}]$ and average mitigation cost λ_t, we get the following expression:

$$\ln DF_j \simeq \ln DF_j^A + \underbrace{\frac{1}{2}\eta(1+\eta)\left[\mathbb{E}V_s\left(\frac{c_{irsj}}{c_j}\right) - V\left(\frac{c_{irt}}{c_t}\right)\right]}_{\text{inequality effect}} + \\ - \underbrace{\eta\mathbb{E}Cov_s\left(\frac{c_{irsj}}{c_j}, \frac{\delta_{irsj}}{\delta_j}\right)}_{\text{benefit inequality premium}} + \underbrace{\eta Cov\left(\frac{c_{irt}}{c_t}, \frac{\lambda_{irt}}{\lambda_t}\right)}_{\text{cost inequality premium}} \quad (10)$$

where DF_j^A is the approximated discount factor in the DICER model.

Relative to the DICER case, the decomposition emphasizes how within-region inequality and its interaction with parameter uncertainty alter the discount factor.

Consider the term labeled "inequality effect", where $V_s\left(\frac{c_{irsj}}{c_j}\right)$ is the variance of the cross-section of individual consumption in state s, and $\mathbb{E}V_s(\cdot)$ is the expected degree of future variance in consumption.[10] In addition, $V\left(\frac{c_{irt}}{c_t}\right)$ is the variance of individual consumption in period t.[11]

[10] The state-dependent variance of individual consumption is computed as (by neglecting the denominator for simplicity):

$$V_s(c_{irsj}) = \frac{1}{5}\sum_{r,i}\frac{L_{rj}}{L_j}\left(c_{irsj} - c_{sj}\right)^2$$

where $c_{sj} \equiv \frac{1}{5}\sum_{r,i}\frac{L_{rj}}{L_j}c_{irsj}$ is the average consumption is state s. Note that this variance term can be further decomposed into a between-region variance and an average within-region variance, thereby disentangling the two sources of inequality:

$$V_s(c_{irsj}) = V_s(c_{rsj}) + \sum_r \frac{L_{rj}}{L_j}V_{rs}(c_{irsj}) = \sum_r \frac{L_{rj}}{L_j}(c_{rsj} - c_{sj})^2 + \sum_r \frac{L_{rj}}{L_j}\left(\frac{1}{5}\sum_i (c_{irsj} - c_{rsj})^2\right)$$

where $V_s(c_{rsj})$ is the variance of average regional consumption in state s, with $c_{rsj} = \frac{1}{5}\sum_i c_{irsj}$, which represents the degree of between-region inequality in state s; $V_{rs}(c_{irsj})$ is the variance of individual consumption in region r and state s, representing the degree of within region inequality, and $\sum_r \frac{L_{rj}}{L_j}V_{rs}$ is the average degree of within region inequality in state s.

[11] Note that also this variance term can be decomposed into a degree of between-region inequality and a degree of within-region inequality:

$$V(c_{irt}) = \frac{1}{5}\sum_{r,i}\frac{L_{rt}}{L_t}(c_{irt} - c_t)^2 = V(c_{rt}) + \sum_r \frac{L_{rt}}{L_t}V_r(c_{irt}) = \sum_r \frac{L_{rt}}{L_t}(c_{rt} - c_t)^2 + \sum_r \frac{L_{rt}}{L_t}\left(\frac{1}{5}\sum_i (c_{irt} - c_{rt})^2\right)$$

where $c_{rt} = \frac{1}{5}\sum_i c_{irt}$ is average consumption in region r.

The decomposition shows that higher relative expected future consumption inequality raises the discount factor, and hence is a force for higher carbon taxation today. Intuitively, since the planner dislikes inequality, she will use carbon taxation to lower future damages, and hence lower future inequality.

The term labeled "benefit inequality premium" captures the fact that carbon taxation creates benefits that are unevenly distributed over the global population. In this term, $Cov_s\left(\frac{c_{irsj}}{c_j}, \frac{\delta_{irsj}}{\delta_j}\right)$ is the covariance between individual consumption and individual marginal benefits from the tax, conditional on the state of the world s, and $\mathbb{E}Cov_s(\cdot,\cdot)$ the expected covariance. When this covariance term is positive, the benefits of carbon taxation flow mainly to those individuals with higher consumption levels. This exacerbates existing inequality, and is therefore a force for a lower discount factor and less carbon taxation today.

Finally, the term labeled "cost-inequality premium" captures the analogous idea that the mitigation costs associated with carbon taxation may be unevenly distributed across the global population. When the covariance of consumption and mitigation costs is positive, individuals with higher consumption also pay larger amounts of the mitigiation costs. This acts to reduce inequality, and is hence a force for a higher discount factor and increased carbon taxation. Note that in our model, the cost-inequality premium is always positive since the marginal mitigation costs are proportional to both individual and regional consumption.[12]

3.3 Interaction of Risk and Inequality

We can now use the decompositions we have derived to shed light on the interaction between parameter uncertainty and inequality. In particular, we can ask how much the optimal carbon tax changes when we introduce the same type of parameter uncertainty into both the DICER and NICER models. If the change in tax is larger in the NICER model, then

[12] Indeed, from the definition of λ_{irt} and the fact that consumption will be a (almost constant) proportion of output, we have that $\lambda_{irt} \simeq \frac{\Lambda'_{rt}}{1-\Lambda_{rt}} c_{irt}$, and $\frac{1}{5}\sum_i \lambda_{irt} \simeq \frac{\Lambda'_{rt}}{1-\Lambda_{rt}} c_{rt}$.

we can conclude that inequality exacerbates the impact of risk, i.e. that risk and inequality have complementary effects on the optimal path of carbon taxes.

Given the commonality of terms in the decompositions for DICER and NICER respectively, and the fact that the term labeled "cost inequality premium" is invariant to parameter uncertainty, it is clear that any interaction between risk and inequality is captured by the terms labeled "inequality effect", and "benefit inequality premium".

Consider first the "inequality effect" term. In order to distinguish the interaction of risk and inequality, it is useful to write

$$\mathbb{E}V_s\left(\frac{c_{irsj}}{c_j}\right) - V\left(\frac{c_{irt}}{c_t}\right) = \underbrace{V\left(\frac{c_{irj}}{c_j}\right) - V\left(\frac{c_{irt}}{c_t}\right)}_{\text{Pure Inequality}} + \underbrace{\mathbb{E}V_s\left(\frac{c_{irsj}}{c_j}\right) - V\left(\frac{c_{irj}}{c_j}\right)}_{\text{Interaction}} \tag{11}$$

where c_{irj} is consumption in the economy without parameter uncertainty.

This expansion breaks the overall "inequality effect" into two components. The difference $V\left(\frac{c_{irj}}{c_j}\right) - V\left(\frac{c_{irt}}{c_t}\right)$ measures the inequality in consumption in the economy without risk, and so captures the pure inequality effect of a carbon tax policy. The difference $\mathbb{E}V_s\left(\frac{c_{irsj}}{c_j}\right) - V\left(\frac{c_{irj}}{c_j}\right)$ then captures the interaction of inequality with risk, where we recall that the pure risk effect is already present in the $\ln DF_j^A$ component of the discount factor expansion.

Therefore, risk and inequality interact to increase the optimal carbon tax in period j when

$$\mathbb{E}V_s\left(\frac{c_{irsj}}{c_j}\right) > V\left(\frac{c_{irj}}{c_j}\right) \tag{12}$$

i.e. when expected variance of consumption inequality in the economy with parameter uncertainty in period j is larger than the variance of consumption in the economy without risk.

Intuitively, in the presence of parameter uncertainty, the inequality-averse planner cares about the average level of consumption inequality across states of the world when setting carbon tax policy. When this expected value is higher than the variance of consumption in the economy without risk, it is optimal to impose a higher carbon tax in order to combat this increase in inequality.

To further understand this condition, suppose that there are two states of the world, one with high aggregate consumption and one with low aggregate consumption, and consider how the variance of consumption at the quintile-region level changes across states of the world. If changes in aggregate consumption are equally distributed across the global population, then the variance will remain at its value in the economy without risk. In this case, there is no effect of the interaction between risk and inequality on the optimal carbon tax.

However, suppose that individuals with relatively high consumption levels are mainly exposed to the state of nature with high aggregate consumption while individuals with relatively low consumption are mainly exposed to the state with low aggregate consumption. This exposure pattern implies that the average variance of consumption is higher than in the economy without risk. As such, the planner chooses higher carbon taxes in order to provide additional insurance against this uneven exposure to parameter uncertainty.

Now consider the "benefit inequality premium" term. A similar decomposition yields that risk and inequality interact to increase the optimal carbon tax in period j when

$$\mathbb{E}Cov_s\left(\frac{c_{irsj}}{c_j}, \frac{\delta_{irsj}}{\delta_j}\right) < Cov\left(\frac{c_{irj}}{c_j}, \frac{\delta_{irj}}{\delta_j}\right) \tag{13}$$

i.e. when the expected covariance between individual consumption and benefits from carbon taxation in period j is less than the covariance in the deterministic economy.

Intuitively, this condition holds when individuals with relatively high consumption levels are mainly exposed to downside parameter risk (states of nature with lower benefit levels) while individuals with relatively low consumption are mainly exposed to upside risk (states of nature with higher benefit levels). Given this exposure pattern, the planner chooses higher carbon taxes in order to exploit the fact that the benefits of taxation tend to reduce consumption inequality.

In conclusion, our analysis has shown that the interaction of risk and inequality calls for higher carbon taxation when parameter uncertainty increases the average variance of consumption across states of the world, and decreases the average covariance between individual consumption and the individual marginal benefits of taxation. In these cases, the optimal path of taxes has a superadditive property: the effect of both risk and inequality on the optimal carbon tax is larger than the sum of the effects studies in isolation.

In order to determine the quantitative magnitude of this superadditivity, we now turn to numerical simulations of the model.

4 Numerical Exercises

Having analyzed the key determinants of optimal carbon taxes, we now compute optimal carbon taxes in the presence of risk and inequality.

4.1 Calibration

Parameter Uncertainty We consider the effects of uncertainty for 4 different parameters: the growth rate of aggregate TFP (the rate at which all regions' TFPs grow in the long run), the convergence rate of regional TFPs (how quickly it takes regions' TFPs to converge to the aggregate level), climate sensitivity, and the quadratic coefficient of the damage function for each region.[13]

In order to create parameter uncertainty, we model each parameter as a random variable with an associated distribution. In order to keep our results comparable to the existing literature on parameter uncertainty, we use standard calibrations of risk for each of the

[13] There are technically 12 of these parameters, but we vary them in a perfectly correlated manner and so treat them as 1 in the exposition.

Parameter	Distribution	Mean	Std. Dev.	Source
Long run TFP growth rate	Normal	1.5%	0.6%	Dietz and Asheim (2012)
TFP Convergence Rate	Beta(2,18)	0.1	0.065	Own calibration
Climate Sensitivity	Log-Logistic	3	1.4	Dietz et al. (2018)
Quadratic Damage Coefficients	Normal	2%	1%	Nordhaus (2018)

Table I: Parameterization of uncertainty.

parameters. Table I summarizes the relevant probability distributions and Appendix C provides the details on the calibrations.

For clarity, we examine the effects of uncertainty for each parameter separately. In order to map our distributional assumptions into the model, we use the set of states of nature $\mathbb{S} = \{1, .., S\}$. Formally, we assign each state s an equal probability, and let the realized parameter value in state s correspond to the s^{th} quintile of the chosen distribution. We tested different sizes of the state space (up to a maximum of $S = 1000$), and they all produce comparable outcomes. The results presented below are based on 10 draws for each parameter.

Inequality There is little existing evidence with which to calibrate the elasticity of damages to income parameter ξ. Therefore, we consider two different values, $\xi \in \{0, 1\}$. When $\xi = 1$, damages are spread proportionally across quintiles and have no impact on consumption inequality. In contrast, when $\xi = 0$, damages are equally divided across quintiles and hence worsen consumption inequality as the poor segments of the population bear disproportionally more damages.

Standard Parameters We set the annual rate of impatience to $\rho = 1.5\%$, corresponding to an annual discount factor of $\beta = 0.985$. We set the coefficient of relative risk (and inequality) aversion to $\eta = 2$.

Table II: Optimal taxes in 2015 ($ per ton of CO_2).

Scenario	DICE	NICER ($\xi = 1$)	NICER ($\xi = 0$)
Deterministic	34.7	31.5	138.2
TFP growth risk	35.4	32	140.2
TFP convergence rate risk	35.7	34.2	149.3
Climate sensitivity risk	37.5	33.6	149.7
Damage coefficient risk	39.3	33.5	151.9

All remaining parameters that govern production technology and the climate system are set in line with the existing integrated assessment modeling literature.

4.2 Optimal Carbon Taxation

Table II displays the optimal carbon tax in period $t = 0$ for the DICER and NICER models, both in the deterministic case and for each type of parameter uncertainty. In the deterministic scenario, the risky parameters are fixed at their respective mean values.

Consider first the DICER model. Reading down the relevant column, we see how the presence of parameter uncertainty raises the short term carbon tax with respect to the deterministic scenario for each type of risk that we consider. Quantitatively, the largest increase occurs in the case of damage function risk.

In the case of consumption inequality, the optimal taxes depend crucially on the value of the elasticity of damages with respect to income ξ. When $\xi = 1$, damages are proportional to consumption and so are unaffected by carbon taxation. In this case, the optimal taxes are very similar to the DICER case in both the deterministic and stochastic scenarios.[14]

[14] The taxes in NICER with $\xi = 1$ are slightly lower than in the DICER model. This occurs since the planner would like to raise the mean consumption level of the economy in order to combat inequality, and does so by slightly reducing the carbon tax.

Table III: Change in optimal taxes in 2015 ($ per ton of CO_2) relative to deterministic DICE model.

Scenario	DICE	NICER ($\xi = 0$)	Additive Effect
Deterministic	0	103.4	103.4
TFP growth risk	0.7	105.5	104.1
TFP convergence rate risk	1	114.5	104.4
Climate sensitivity risk	2.8	114.9	106.2
Damage coefficient risk	4.6	117.2	108

When $\xi = 0$ however, damages fall disproportionally on individuals with lower consumption. As such, the planner has an incentive to increase the carbon tax to reduce inequality due to climate damages. Quantitatively, the optimal carbon taxes exhibit significant deviations from the DICER values. The optimal taxes increase by approximately $100 for all risks, and the taxes display substantial dispersion across the different parameter risks we consider. In particular, uncertainty on the damage function coefficients results in an optimal carbon tax of $167 per ton of CO_2, which is around 5 times larger than optimal taxes in the DICER economy.

Super-additivity In order to examine whether the interaction of risk and inequality results in super-additive optimal taxes, we study whether the total change in the optimal tax due to risk and inequality is larger than the sum of the changes due to each separate effect. The results are summarized in table III, which computes the tax changes relative to the deterministic economy without inequality.

The DICE column computes the change in the optimal tax due to parameter uncertainty alone. The first row of the NICER column shows that inequality (but no risk) causes the tax to increase by $103 relative to the deterministic economy without inequality. Therefore, we can compute the "additive effect" of risk and inequality by summing the DICE column and the deterministic NICER effect. These sums are computed in the final column.

We can then compare the additive effects to the NICER column, which computes the tax change due to the joint effect of risk and inequality. This comparison provides clear evidence of the "super-additivity" property we illuminated analytically. For example, damage risk alone causes the optimal tax to increase by almost \$5, while inequality alone causes the tax to increase by \$103. Therefore, the sum of the separate effects implies an increase of \$108, while the joint effect is an increase of \$117. Therefore, the joint interaction of damage risk and inequality generates almost 10% of the total increase in the optimal carbon tax.

5 Conclusion

We study the interaction of risk and inequality, and examine their joint impact on optimal carbon taxation. Our analytical and numerical results demonstrate a strong interaction effect, that justifies our approach of analyzing both features of the world simultaneously.

Our numerical results are of course sensitive to our calibrations of parameter uncertainty, unequal damage exposure (the ξ parameter), and our choice of objective function. Our calibration of risk is somewhat conservative, and ignores extreme tail events or "disaster" risk, which likely understates the true extent of its interaction with inequality.

Similarly, we have assumed equal (and conservative) values for risk and inequality aversion. Exploring more nuanced approaches to evaluating social welfare in this context would be a worthwhile exercise.

A Tables and Graphs

Figure 1: Taxes across all risks, DICE

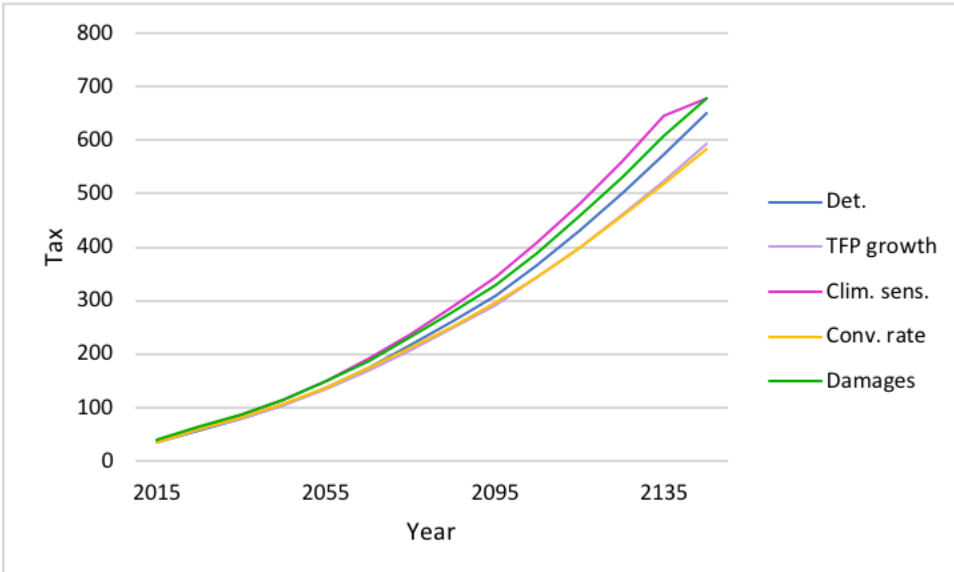

Figure 2: Taxes across all risks, NICER $\xi = 1$

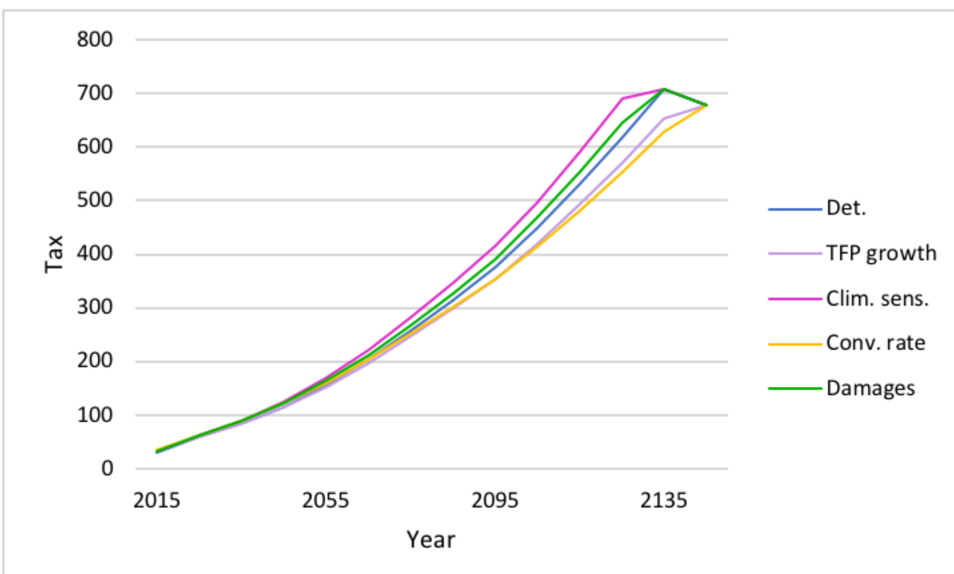

Figure 3: Taxes across all risks, NICER $\xi = 0$

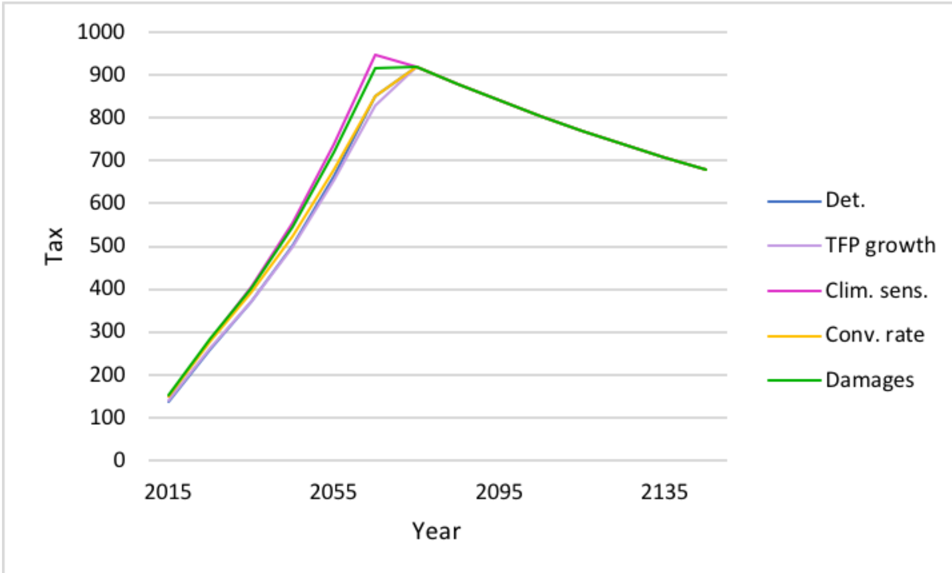

B The NICER model

Let i, r, s, t denote, respectively, quintile, region, state of the world and generation (or time period). As in Nordhaus' RICE model, we assume that the world is populated by 12 autarkic, single-good production regions, $r = 1, .., 12$, each of them characterized by a regional production function $Q_{rst} = F_{rst}(K_{rst}, L_{rt})$, which depends on (exogenous) regional population L_{rt} and (endogenous) regional capital stock K_{rst}.

Regional emissions are proportional to gross output. The mitigation rate μ_{rst} is the proportion of emissions reduced with respect to the level that would be emitted in the business as usual scenario, $E_{rst} = \sigma_{rt}(1 - \mu_{rst})Q_{rst}$, where σ_{rt} represents the exogenous emissions to output ratio. Total period t emissions add to the stock of carbon in the atmosphere, which induces global warming. Let T_{st} be the average increase in atmospheric temperature with respect to the pre-industrial level. Due to the inertia of the climate system and the long-lived nature of carbon, emissions produced at t will affect the increase in temperature for all periods $j \geq t$. The climate module, which describes how emissions contribute to warming, follows exactly the one in the RICE model (see the online Appendix of Nordhaus (2010)).

The increase in temperature produces damages in terms of production loss. Let Y_{rst} be the regional output net of climate damages D_{rst} and mitigation costs Λ_{rst}, such that:

$$Y_{rst} = \frac{1 - \Lambda_{rst}}{1 + D_{rst}} Q_{rst}$$

The damage function is quadratic, $D_{rst} = \psi_{rs} T_{st} + \phi_r T_{st}^2$, while the mitigation cost is a convex function of the mitigation rate, $\Lambda_{rst} = \theta_{rt}^1 \mu_{rst}^{\theta_2}$, where θ_{rt}^1 and θ_2 are exogenous parameters whose values depend on the price of a completely green backstop technology (see Nordhaus (2010) for details).[15] Finally, each region faces a budget constraint, $C_{rst} = Y_{rst} - K_{rst+1}$, where C_{rst} denotes total regional consumption.

We assume that four types of exogenous parameters are subject to uncertainty: initial TFP growth rates (12 parameters as there are 12 regions), TFP convergence rate, climate sensitivity and the linear coefficient of the damage function ψ_{rs} (again, 12 parameters as there are 12 regions). Uncertainty over the TFP parameters lead to parametric uncertainty over the regional production functions F_{rst}, while climate sensitivity risk induces uncertainty over the increase in average atmospheric temperature T_{st}.

In the NICER model (and in its ancestor NICE), it is assumed that regions select optimal mitigation, consumption and capital levels given a uniform global tax on emissions $\vec{\tau} = (\tau_0, ..., \tau_T)$. Regions maximize their own ex-ante expected inter-temporal utility function $U_r(\vec{\tau})$:

$$U_r(\vec{\tau}) = \sum_t \beta^t \frac{L_{rt}}{5} \sum_i \frac{\mathbb{E}[c_{irst}(\vec{\tau})^{1-\eta}]}{1 - \eta}$$

[15] In RICE, $\theta_2 = 2.8$, while θ_{rt}^1 are decreasing over time at an exogenous rate and calibrated such that, at the beginning of the simulation period (2005), the marginal cost of the last unit of mitigation is equal to the price of the backstop technology.

where $c_{irst}(\vec{\tau}) = C_{rst}(\vec{\tau})\left(q_{ir}(1+D_{rst}(\vec{\tau})) - d_{ir}D_{rst}(\vec{\tau})\right)\frac{5}{L_{rt}}$, as explained in the main text (equation 3). The expectation operator $\mathbb{E}$ reflects the amount of information available today (i.e. to generation $t = 0$) about the true value of the unknown parameters. As there is no learning, the information set does not change over time, and $\mathbb{E}$ captures expectations over all the possible future realizations of the uncertain parameters. For a given carbon tax trajectory $\vec{\tau} = (\tau_0, ..., \tau_T)$, the maximization will result in a stream of optimal mitigation levels $\{\mu^*_{rst}(\vec{\tau})\}_{\forall r,s,t}$, optimal capital levels $\{K^*_{rst}(\vec{\tau})\}_{\forall r,s,t}$, and optimal consumption levels $\{C^*_{rst}(\vec{\tau})\}_{\forall r,s,t}$, given a path for the carbon tax $\vec{\tau}$. Following Nordhaus, optimal regional mitigation rates are taken to be:

$$\mu^*_{rt}(\vec{\tau}) = \left(\frac{\tau_t \sigma_{rt}}{\theta^1_{rt}\theta_2}\right)^{\frac{1}{\theta_2 - 1}} \tag{14}$$

As a consequence, mitigation rates are state-independent, which implies that also the mitigation costs Λ_{rt} are state-independent. In other words, as long as the tax is uniform across states of the world, also the mitigation rates are going to be uniform across states of the world, although potentially region-specific. Moreover, as already pointed out in the literature, the integrated assessment models belonging to the DICE family imply a saving rate that is largely independent of the climate (Golosov et al., 2014; Traeger, 2015; Gerlagh and Liski, 2018), and almost constant over time. That is true also in our case; as a consequence, in the simulation, when determining the optimal tax trajectory, we approximate the savings rates to an exogenous constant level S, such that $C^*_{rst}(\vec{\tau}) \simeq (1-S)Y^*_{rst}(\vec{\tau})$, where the net output level depends on the optimal mitigation rates $\{\mu^*_{rj}(\vec{\tau})\}_{\forall j,r}$.

The tax is chosen by a benevolent policy-maker who takes into account regions' reaction, and who maximizes the welfare function W:

$$W(\vec{\tau}) = \sum_r U_r(\vec{\tau}) = \sum_{t,r,i} \beta^t \frac{L_{rt}}{5} \mathbb{E}\left[\frac{c_{irst}(\vec{\tau})^{1-\eta}}{1-\eta}\right]$$

By taking into account regions' optimal choice of mitigation, capital and consumption for a given tax path, a first order derivative of the previous expression with respect to τ_t, for all t, yields the following condition:

$$-\beta^t \mathbb{E}\left[\sum_{r,i} L_{rt} c_{irst}^{-\eta} \frac{\Lambda'_{rt}}{1-\Lambda_{rt}} \frac{Y_{rst}}{L_{rt}} \gamma_{irst}\right] + \sum_{j \geq t} \beta^j \mathbb{E}\left[\sum_{r,i} L_{rj} c_{irsj}^{-\eta} \frac{-D'_{rsj}}{1+D_{rsj}} d_{ir} \frac{Y_{rsj}}{L_{rj}}\right] = 0$$

with $\Lambda'_{rt} = \frac{\partial \Lambda_{rt}}{\partial \tau_t}$, $\gamma_{irst} = q_{ir}(1+D_{rst}) - D_{rst} d_{ir}$, and $D'_{rsj} = \frac{\partial D_{rsj}}{\partial \tau_t}$. The condition is derived by replacing the expression for optimal mitigation rates (14) in the welfare function, and differentiating with respect to τ_t, while taking into account the optimal regional responses in terms of consumption $\{C^*_{rsj}(\vec{\tau})\}_{\forall r,s,j}$ and capital $\{K^*_{rsj}(\vec{\tau})\}_{\forall r,s,j}$. By the envelope theorem, only the impact of $\vec{\tau}$ on mitigation rates matters when choosing the optimal tax profile.[16] The tax will increase mitigation costs at t (the first term in the previous condition), and will reduce emissions E_{rst}, which will induce a lower increase in atmospheric temperature for all future periods $j \geq t$, and, as a result, a reduction in future damages (the second term in the expression).

Let $\lambda_{irst} \equiv \frac{\Lambda'_{rt}}{1-\Lambda_{rt}} \frac{Y_{rst}}{L_{rt}} \gamma_{irst}$ be the marginal mitigation cost (in terms of lost output) paid by individuals in quintile i, region r and state of the world s, and $\delta_{irsj} \equiv \frac{-D'_{rsj}}{1+D_{rsj}} d_{ir} \frac{Y_{rsj}}{L_{rj}}$ the marginal returns of the tax (in terms of gained output) for individuals in quintile i, region r, state s and generation j. Then, the previous condition is equal to:

$$-\beta^t \mathbb{E}\left[\sum_{r,i} L_{rt} c_{irst}^{-\eta} \lambda_{irst}\right] + \sum_{j \geq t} \beta^j \mathbb{E}\left[\sum_{r,i} L_{rj} c_{irsj}^{-\eta} \tilde{\delta}_{irsj}\right] = 0$$

Note that this condition takes into account the possibility that the tax at t is chosen before knowing the realization of the risk at t. The expression can be rewritten so as to highlight

[16] For the sake of simplicity, we skip the derivation of the regional first order conditions with respect to capital and consumption, as, by the envelope theorem, we can look just at the direct effects of the tax on total welfare, while neglecting the indirect effects through the optimal levels of investment and consumption.

the impact of the risk at t on the tax choice. Let $c_{irt} = \mathbb{E}[c_{irst}]$ and $\lambda_{irt} = \mathbb{E}[\lambda_{irst}]$. Then:

$$-\beta^t \sum_{r,i} L_{rt} c_{irt}^{-\eta} \lambda_{irt} + \beta^t \sum_{r,i} L_{rt} \left(c_{irt}^{-\eta} \lambda_{irt} - \mathbb{E}[c_{irst}^{-\eta} \lambda_{irst}] \right) + \sum_{j \geq t} \beta^j \mathbb{E} \left[\sum_{r,i} L_{rj} c_{irsj}^{-\eta} \tilde{\delta}_{irsj} \right] = 0$$

Let $\hat{\mathbb{E}}$ denote the expectation operator that computes the average value of a variable with respect to both states of the world and individuals (regions+quintile): $\hat{\mathbb{E}} x_{irst} = \sum_{r,i} \frac{L_{rt}}{L_t} \frac{1}{5} \mathbb{E} x_{irst}$, for any random variable x and period t, with $L_t = \sum_r L_{rt}$ denoting total population at t. If we multiply and divide the first term of the previous expression by $\hat{\mathbb{E}}[\lambda_{irt}] \hat{\mathbb{E}}[c_{irt}^{-\eta}]$, after some re-adjustments we obtain the following condition:[17]

$$\lambda_t = \sum_{j \geq t} \beta^{j-t} \frac{L_j}{L_t} \left(\frac{\hat{\mathbb{E}}\left[c_{irsj}^{-\eta} \delta_{irsj}\right]}{\hat{\mathbb{E}}\left[c_{irt}^{-\eta}\right] \delta_j} \right) \left(\frac{\hat{\mathbb{E}}\left[c_{irt}^{-\eta}\right] \hat{\mathbb{E}}\left[\lambda_{irt}\right]}{\hat{\mathbb{E}}\left[c_{irt}^{-\eta} \lambda_{irt}\right]} \right) \delta_j + \chi_t \tag{15}$$

where $\lambda_t \equiv \hat{\mathbb{E}} \lambda_{irt}$, and $\delta_j \equiv \hat{\mathbb{E}}[\delta_{irsj}]$ represents the average expected marginal returns from the policy. The residual term χ_t adjusts for the presence of risk at t:

$$\chi_t = \sum_{j \geq t} \beta^{j-t} \frac{L_j}{L_t} \left(\frac{\hat{\mathbb{E}}\left[c_{irsj}^{-\eta} \delta_{irsj}\right]}{\hat{\mathbb{E}}\left[c_{irt}^{-\eta}\right]} \right) \left(\frac{\hat{\mathbb{E}}\left[c_{irt}^{-\eta}\right] \hat{\mathbb{E}}\left[\lambda_{irt}\right]}{\hat{\mathbb{E}}\left[c_{irt}^{-\eta} \lambda_{irt}\right]} \right) \left(\frac{\hat{\mathbb{E}}\left[c_{irt}^{-\eta} \lambda_{irt}\right]}{\hat{\mathbb{E}}\left[c_{irst}^{-\eta} \lambda_{irst}\right]} - 1 \right)$$

According to (15), the optimal carbon tax is such that, at the margin, the average expected cost imposed by the tax is equal to the present value of the future average expected benefits derived from the policy itself. The variable δ_j measures the expected average returns at the margin for generation j, while all the other terms denote the discount factor that is applied to those average expected benefits to make them comparable to the current costs λ_t.

Let DF_j be the discount factor applied to the expected average marginal returns of the policy, δ_j:

$$DF_j \equiv \beta^{j-t} \frac{L_j}{L_t} \left(\frac{\hat{\mathbb{E}}\left[c_{irsj}^{-\eta} \delta_{irsj}\right]}{\hat{\mathbb{E}}\left[c_{irt}^{-\eta}\right] \delta_j} \right) \left(\frac{\hat{\mathbb{E}}\left[c_{irt}^{-\eta}\right] \hat{\mathbb{E}}\left[\lambda_{irt}\right]}{\hat{\mathbb{E}}\left[c_{irt}^{-\eta} \lambda_{irt}\right]} \right)$$

[17]Note that $\hat{\mathbb{E}}[\lambda_{irt}] = \sum_{r,i} \frac{L_{rt}}{L_t} \frac{1}{5} \mathbb{E} \lambda_{irt} \equiv \sum_{r,i} \frac{L_{rt}}{L_t} \frac{1}{5} \lambda_{irst}$ as the variable λ_{irt} is independent of the state of the world. Similarly for $\hat{\mathbb{E}}[\delta_{irt}]$.

The discount factor can be decomposed as follows:

$$DF_j = (j-t)\ln\beta + \ln\frac{L_j}{L_t} \underbrace{-\eta\ln\frac{c_j}{c_t}}_{\text{growth effect}} + \underbrace{\ln\frac{\mathbb{E}\left[c_{sj}^{-\eta}\right]}{c_j^{-\eta}}}_{\text{precautionary effect}} + \underbrace{\ln\frac{\mathbb{E}\left[c_{sj}^{-\eta}\delta_{sj}\right]}{\mathbb{E}\left[c_{sj}^{-\eta}\right]\delta_j}}_{\text{risk premium}} +$$

$$+\underbrace{\ln\frac{\hat{\mathbb{E}}\left[c_{irsj}^{-\eta}\right]}{\mathbb{E}\left[c_{sj}^{-\eta}\right]} - \ln\frac{\hat{\mathbb{E}}\left[c_{irt}^{-\eta}\right]}{c_t^{-\eta}}}_{\text{inequality effect}} + \underbrace{\ln\frac{\hat{\mathbb{E}}\left[c_{irsj}^{-\eta}\delta_{irsj}\right]}{\hat{\mathbb{E}}\left[c_{irsj}^{-\eta}\right]}\frac{\mathbb{E}\left[c_{sj}^{-\eta}\right]}{\mathbb{E}\left[c_{sj}^{-\eta}\delta_{sj}\right]}}_{\text{benefit inequality premium}} - \underbrace{\ln\frac{\hat{\mathbb{E}}\left[c_{irt}^{-\eta}\lambda_{irt}\right]}{\hat{\mathbb{E}}\left[c_{irt}^{-\eta}\right]\lambda_t}}_{\text{cost inequality premium}}$$

A second order Taylor approximation of the previous expression around c_j, δ_j and λ_t yields expression (10).

A similar procedure is followed for the aggregate model á la DICE, only that in this case we consider a single policy-maker, who decides both how much each region has to mitigate, invest and consume given a global tax, and the optimal carbon tax over time. In that case, the policy maker maximizes the welfare function W^A:

$$W^A(\vec{\tau}) = \sum_t \beta^t L_t \frac{\mathbb{E}[c_{st}(\vec{\tau})^{1-\eta}]}{1-\eta}$$

where $c_{st}(\vec{\tau}) = \frac{1}{5}\sum_{r,i}\frac{L_{rt}}{L_t}c_{irst}(\vec{\tau})$ is the average per-capita consumption in state s. Taking into account the optimal levels of mitigation (still given by (14)), investment and consumption and the regional budget constraints (and the fact that, as in the dis-aggregate economy, the optimal mitigation rates are uniform across states of the world), by the envelope theorem the maximization with respect to the tax τ_t yields the following expression, for each t:

$$-\beta^t\mathbb{E}\left[c_{st}^{-\eta}\sum_r L_{rt}\frac{\Lambda'_{rt}}{1-\Lambda_{rt}}\frac{Y_{rst}}{L_{rt}}\right] + \sum_{j\geq t}\beta^j\mathbb{E}\left[c_{js}^{-\eta}\sum_r L_{rj}\frac{-D'_{rst}}{1+D_{rst}}\frac{Y_{rst}}{L_{rt}}\right] = 0$$

with $\Lambda'_{rt} = \frac{\partial \Lambda_{rt}}{\partial \tau_t}$, and $D'_{rsj} = \frac{\partial D_{rsj}}{\partial \tau_t}$, for all $j \geq t$. Note that the summation over quintiles i disappears from the first order condition in the aggregate case, since, by (3), $\frac{1}{5}\sum_i c_{irst} = \frac{C_{rst}}{L_{rt}} \sum_i [q_{ir}(1 + D_{rst}) - d_{ir} D_{rst}] = \frac{C_{rst}}{L_{rt}}$ as $\sum_i q_{ir} = \sum_i d_{ir} = 1$.

By defining $\lambda_{rst} = \frac{\Lambda'_{rt}}{1-\Lambda_{rt}} \frac{Y_{rst}}{L_{rt}}$ the marginal regional mitigation cost in terms of output loss, and $\delta_{rsj} = \frac{-D'_{rsj}}{1+D_{rsj}} \frac{Y_{rst}}{L_{rt}}$ the marginal regional returns of the tax in terms of output gained (note that $\delta_{rsj} = \frac{1}{5}\sum_i \delta_{irsj}$), the previous condition can be rewritten as:

$$-\beta^t \mathbb{E}\left[c_{st}^{-\eta} \sum_r L_{rt} \lambda_{rst}\right] + \sum_{j \geq t} \beta^j \mathbb{E}\left[c_{js}^{-\eta} \sum_r L_{rj} \delta_{rsj}\right] = 0$$

Once again, the optimality condition has been derived under the assumption that the policy at t is chosen before the realization of the risk at t. As before, we can rewrite the condition by isolating the impact of risk at t:

$$-\beta^t L_t c_t^{-\eta} \lambda_t + \beta^t L_t \left(c_t^{-\eta} \lambda_t - \mathbb{E}[c_{st}^{-\eta} \lambda_{st}]\right) + \sum_{j \geq t} \beta^j \mathbb{E}\left[c_{js}^{-\eta} \sum_r L_{rj} \delta_{rsj}\right] = 0$$

with $c_t = \mathbb{E}[c_{st}]$, and $\lambda_{st} = \sum_r \frac{L_{rt}}{L_t} \lambda_{rst}$. Let $\delta_{sj} = \sum_r \frac{L_{rj}}{L_j} \delta_{rsj}$ be the average marginal benefit from the policy for region r and state s, for all $j \geq t$. Furthermore, let us consider also the expected average marginal costs and returns of the policy: $\lambda_t = \mathbb{E}\lambda_{st}$ and $\delta_j = \mathbb{E}\delta_{sj}$, respectively. After some re-adjustments we obtain the following condition:

$$\lambda_t = \sum_{j \geq t} \beta^{j-t} \frac{L_j}{L_t} \left(\frac{\mathbb{E}\left[c_{sj}^{-\eta} \delta_{sj}\right]}{c_t^{-\eta} \delta_j}\right) \delta_j + \chi_t^A \tag{16}$$

where the residual term χ_t^A is defined as:

$$\chi_t^A \equiv \sum_{j \geq t} \beta^{j-t} \frac{L_j}{L_t} \frac{\mathbb{E}\left[c_{sj}^{-\eta} \delta_{sj}\right]}{c_t^{-\eta}} \left(\frac{\lambda_t c_t^{-\eta}}{\mathbb{E}[c_{st}^{-\eta} \lambda_{st}]} - 1\right)$$

The residual term takes into account the fact that the optimal tax at t may be chosen before knowing the realization of the risks at t. Note that for aggregate random variables such as c_{sj} and δ_{sj}, the expectation operators $\hat{\mathbb{E}}$ and $\mathbb{E}$ yield the same result.

Let DF_j^A be the discount factor for the expected marginal benefits occurring at period j, δ_j, under the assumption that consumption at t is observed:

$$DF_j^A \equiv \beta^{j-t}\frac{L_j}{L_t}\left(\frac{\mathbb{E}\left[c_{sj}^{-\eta}\delta_{sj}\right]}{c_t^{-\eta}\delta_j}\right)$$

The discount factor can be decomposed into its main determinants as follows:

$$\ln DF_j^A = (j-t)\ln\beta + \ln\frac{L_j}{L_t} \underbrace{-\eta\ln\frac{c_j}{c_t}}_{\text{growth effect}} + \underbrace{\ln\frac{\mathbb{E}\left[c_{sj}^{-\eta}\right]}{c_j^{-\eta}}}_{\text{precautionary effect}} + \underbrace{\ln\frac{\mathbb{E}\left[c_{sj}^{-\eta}\delta_{sj}\right]}{\mathbb{E}\left[c_{sj}^{-\eta}\right]\delta_j}}_{\text{risk premium}}$$

where $c_j = \mathbb{E}c_{sj}$. A second order Taylor approximation of the previous components around c_j and δ_j yields expression (8).

C Calibrations

The probability distributions for the initial growth rate of TFP and the climate sensitivity parameter are drawn by Dietz et al. (2016) and adjusted to meet the characteristics of our model. In that paper, the initial growth rate of TFP is assumed to follow a Normal distribution with standard deviation equal to 0.0059 and mean equal to 0.0084. In Dietz et al. there is a single representative agent, while we have multiple regions. Therefore, we replaced the aggregate mean with a set of regional means, representing the adjusted average in GDP per capita growth from 1995 to 2015: USA (0.0151); OECD Europe (0.0162); Japan (0.0138); Russia (0.026); Non-Russia Eurasia (0.0247); China (0.0714); India (0.0455);

Middlea-East (0.0235); Africa (0.0365); Latin America (0.0292); OHI (0.0188); Other non-OECD Asia (0.028). We keep the assumption that all regional TFP risks have the same standard deviation equal to 0.0059.

In Dietz et al., the climate sensitivity parameter has a loglogistic distribution with mean 2.9 and standard deviation 1.4, truncated from below at 0.75. We keep the same type of distribution, but we use a mean equal to 3, such as to make our work comparable to other studies.

The last parameter subject to uncertainty is the rate of convergence of regional GDP. In the RICE model (and original NICE), the TFP convergence rate is set at 10% per period. Lacking any empirical evidence, we decided to use a Beta distribution with shape parameters $\alpha = 2$ and $\beta = 18$. The result is a distribution with a mean value of 0.1 and a variance of 0.0043, which means that most of the mass is in a neighborhood of 0.1.

References

Cai, Y., K. L. Judd, and T. S. Lontzek (2013). The social cost of stochastic and irreversible climate change. Technical report, National Bureau of Economic Research. 4

Chichilnisky, G. and G. Heal (1994). Who should abate carbon emissions?: An international viewpoint. *Economics Letters 44*(4), 443–449. 9

Crost, B. and C. P. Traeger (2014). Optimal CO2 mitigation under damage risk valuation. *Nature Climate Change 4*(7), 631–636. 4

Daniel, K. D., R. B. Litterman, and G. Wagner (2016). Applying asset pricing theory to calibrate the price of climate risk. Technical report, National Bureau of Economic Research. 4

Dennig, F., M. B. Budolfson, M. Fleurbaey, A. Siebert, and R. H. Socolow (2015). Inequality, climate impacts on the future poor, and carbon prices. *Proceedings of the National Academy of Sciences 112*(52), 15827–15832. 2, 5, 6, 7, 8

Dietz, S. and G. B. Asheim (2012). Climate policy under sustainable discounted utilitarianism. *Journal of Environmental Economics and Management 63*(3), 321–335. 21

Dietz, S., C. Gollier, and L. Kessler (2018). The climate beta. *Journal of Environmental Economics and Management 87*, 258–274. 4, 13, 21

Gerlagh, R. and M. Liski (2018). Carbon prices for the next hundred years. *The Economic Journal 128*, 728–757. 28

Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski (2014). Optimal taxes on fossil fuel in general equilibrium. *Econometrica 82*(1), 41–88. 28

Hallegatte, S., M. Bangalore, M. Fay, T. Kane, and L. Bonzanigo (2015). *Shock waves: managing the impacts of climate change on poverty*. World Bank Publications. 2

Howarth, R. B. (2003). Discounting and uncertainty in climate change policy analysis. *Land Economics 79*(3), 369–381. 4

Hwang, I. C., F. Reynes, and R. Tol (2017). The effect of learning on climate policy under fat-tailed uncertainty. *Resource and Energy Economics 48*, 1–18. 4

Jensen, S. and C. Traeger (2021). Pricing Climate Risk. Technical report, working paper. 4

Jensen, S. and C. P. Traeger (2014). Optimal climate change mitigation under long-term growth uncertainty: Stochastic integrated assessment and analytic findings. *European Economic Review 69*, 104–125. 4

Kelly, D. L. and C. D. Kolstad (1999). Bayesian learning, growth, and pollution. *Journal of economic dynamics and control 23*(4), 491–518. 4

Kelly, D. L. and Z. Tan (2015). Learning and climate feedbacks: Optimal climate insurance and fat tails. *Journal of Environmental Economics and Management 72*, 98–122. 4

Leach, A. J. (2007). The climate change learning curve. *Journal of Economic Dynamics and Control 31*(5), 1728–1752. 4

Lemoine, D. (2021). The climate risk premium: how uncertainty affects the social cost of carbon. *Journal of the Association of Environmental and Resource Economists 8*(1). 4

Lemoine, D. and C. P. Traeger (2016). Economics of tipping the climate dominoes. *Nature Climate Change 6*(5), 514–519. 4

Lontzek, T. S., Y. Cai, K. L. Judd, and T. M. Lenton (2015). Stochastic integrated assessment of climate tipping points indicates the need for strict climate policy. *Nature Climate Change 5*(5), 441–444. 4

Nordhaus, W. (2018, August). Projections and uncertainties about climate change in an era of minimal climate policies. *American Economic Journal: Economic Policy 10*(3), 333–60. 3

Nordhaus, W. D. (2010). Economic aspects of global warming in a post-copenhagen environment. *Proceedings of the National Academy of Sciences 107*(26), 11721–11726. 2, 4, 5, 27

Sandsmark, M. and H. Vennemo (2007). A portfolio approach to climate investments: CAPM and endogenous risk. *Environmental and Resource Economics 37*(4), 681. 4

Traeger, C. P. (2015). Analytic integrated assessment and uncertainty. 28