

Risk, Inequality, and Climate Change^{*}

Joshua Bernstein[†], Francis Dennig[‡], Maddalena Ferranna[§], Marc Fleurbaey[¶]

April 13, 2017

Abstract

We study both analytically and numerically the impact of risk and inequality on the optimal carbon price, focusing in particular on the possibility that the world's poor will bear the brunt of climate damages. Building on an extension of the NICE model (Dennig et al. 2015) that features uncertainty over a set of key parameters, we analyze whether the presence of unequal damage impacts reinforces the incentives to hedge against future climate risks, and whether it affects the value of learning. We show that there is a positive complementarity between risk and inequality, especially when the risk concerns the damage function. In contrast, the value of learning is barely affected by the presence of inequality, due to the large inertia of the climate system and the fact that we are on a pre-learning optimal path. The analysis of an extreme form of wait-and-see strategy, “do nothing till we learn”, highlights how unequal damage incidence disproportionately increases the cost of delay.

^{*}This paper has benefited from a workshop in Princeton, and from funding by the Princeton Climate Futures Initiative via a PEI-ACEE grant.

[†]Princeton University

[‡]Yale-National University of Singapore College

[§]Princeton University

[¶]Princeton University

1 Introduction

A core challenge for climate policy is the great uncertainty surrounding the pace and impacts of climate change. In addition, impact damages are unlikely to be spread proportionally to income among the world population, because poor populations are more vulnerable to weather and health shocks. Such "unequal damage incidence" has been confirmed in a recent World Bank report (2015), and its implications for carbon pricing have been studied in a literature spanning from Anthoff et al. (2009), where the focus was on geographical impacts across countries of unequal levels of development, to Dennig et al. (2015), focusing on socio-economic inequalities. An overriding theme emerges from these analyses: when climate damages are sufficiently biased towards hurting the world's poor more than the rich, optimal tax policy significantly departs from the baseline in which damages are proportional to income.

In this paper, we study how unequal damage incidence interacts with structural uncertainty and re-examine the impact of uncertainty over optimal carbon prices, with or without the possibility of learning, following in the footsteps of Keller et al. (2006) and subsequent literature. In order to analyze the effects of unequal damage incidence under uncertainty, we extend the deterministic NICE model, developed by Dennig et al. (2015), to allow for uncertainty about model parameters. NICE builds on Nordhaus' RICE model¹, introducing within-region inequalities and letting these inequalities depend on the level of damages in each region, so that higher regional damage can increase regional consumption inequality (via unequal damage incidence). In this paper we extend the NICE framework to consider stochastic scenarios in which some model parameters are treated as random variables, and policy choices can branch into two paths from a point in time when learning occurs about which half of the state space we are in: the more optimistic parameter values, or the more pessimistic ones (similarly to Keller et al., 2006). We call this variant of the model NICER (Nested Inequalities Climate Economy with Risk).

We consider (separately) uncertainty on the parameters governing the regional total factor productivity (TFP) growth rates, the convergence rate of regional TFPs, climate sensitivity, and the linear parameters of the (quadratic) regional damage functions. We use standard calibrations for all parameters except those in the damage functions. Here, we provide a novel calibration, allowing for the non-monotonic relationship between climate and

¹See <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.

output summarised in Tol (2012). We draw three key conclusions: 1) In the short run, all forms of uncertainty imply higher optimal taxes than in the deterministic case, a result that confirms the existing literature; however, in the long run, “background” uncertainties (TFP, convergence rate) reduce the optimal tax whereas the other risks maintain the positive gap in taxes relative to the deterministic case. 2) Across all time horizons and risks, more unequal damage incidence leads to higher optimal taxes. This is in line with the results found in Dennig et al. (2015) for the model without parameter uncertainty. The reason for this is that the marginal burden to the social objective of an additional ton of carbon is higher when there is unequal damage incidence than when there is proportional damage incidence. 3) Unequal damage incidence reinforces the increase in taxes due to the presence of risk. The strength of this complementarity varies across forms of uncertainty. For example, the combined effect of unequal incidence and damage function risk on optimal taxes is roughly 1.5 times the sum of the separate effects, while the combination of unequal incidence and TFP growth rate risk is almost identical to the sum of the separate effects.

While our paper focuses on mitigation policy, it is worth observing that adaptation mechanisms that insure populations against climate damages can be interpreted, in our stylized model, as reducing the phenomenon of unequal damage incidence. In view of our results, such adaptation policies, if a long-term commitment to steadfastly implement them could be made, would make the optimal mitigation efforts less stringent.

In order to explain our results, we analytically determine the main drivers through which the presence of risk affects the optimal tax. We find that for our risk calibrations the technological risks translate into far greater risk in consumption than the risks related to the damage function. However, technological risks imply that mitigation produces its main benefits when the economy is growing faster, thereby creating a positive “climate risk premium”. Since the importance of risk reduces over time, while the positive correlation between the state of the economy and the returns of the policy increases, the presence of technological risks has a small effect on optimal carbon pricing, if not negative. Risks related to damage, instead, are highly negatively correlated with the economy, especially in the presence of unequal damage incidence. In that case, mitigation not only insures against future damage risks, but it also reduces future potential inequalities created by climate change.

We further explore the link between risk and unequal damage incidence by allowing the social planner to observe a precise signal about the value of the uncertain parameter at some future date. Specifically, at a point in time determined by the modeler, the social planner learns (with certainty) about which *half* of the state space she is in. This allows

her to adjust the optimal policy to the new information from that point in time onwards. We find that, regardless of when learning occurs, the single path before learning is (almost) unaffected, and the two paths after the learning node are also (almost) unaffected. Keller et al. (2006) find similar results. We also compute the value of learning, i.e., the amount society would be willing to pay in order to bring the learning node forward (thus be able to update and re-optimize earlier). We find that this is very small for all the parameters we consider, and, perhaps more surprisingly, unequal damage incidence has little additional effect on this value.

Note that learning at a later date is equivalent, for the optimization problem, to learning at an earlier date under the constraint of not using the information learned until the later date, i.e., retaining a policy that is the same in all states of the world. Therefore, in order to understand how the optimum changes when a constraint on the social planner is lifted at a point in time, we investigate a more extreme form of constraint: zero taxes (business-as-usual) for the first three periods (until the year 2045), and the continuation optimum from then on. We interpret this scenario as one in which policy makers choose to delay action either to wait and learn more about the climate, or due to political constraints. Unlike in the learning case, we find that the cost of delaying action is very sensitive to the strength of unequal damage incidence, and increases by a factor of 10 over a reasonable range. Further, optimal tax policy displays the same sensitivity to unequal damage incidence, and also exhibits compensating behavior, with taxes in the delay case overshooting those in the no-delay case once the delay period has ended. We decompose the cost of delay, and show that unequal damage incidence generates a novel and numerically significant term reflecting how inequality differs under the delay and no-delay scenarios.

We argue that the costs of delay are so much larger than the values of learning because of the evolution of climate variables in each case. In our study of learning, taxes are still optimized before the learning node, balancing the costs of being in either of the two halves of the state space. This results in significant emission reductions relative to business-as-usual (i.e., the delay case). This emission reduction results in significantly lower temperatures and, thus, much lower damages. The value of learning is small and insensitive to unequal damage because the policy before the learning node already provides much protection against the worst effects of climate damages, while a delay of 3 decades commits us to quite large damages, even if the optimal policy is chosen from then on.

Related Literature. Our paper is related to the literature on stochastic integrated assessment models, which studies how optimal mitigation policy depends on uncertainty (and eventually learning) about economic growth (Jensen and Traeger, 2014), damages (Daniel et al., 2016; Crost and Traeger, 2014; Cai et al., 2013), climate sensitivity and warming (Jensen and Traeger, 2016; Kelly and Tan, 2015; Hwang et al. 2014; Leach, 2007; Kelly and Kolstad, 1999), and tipping points (Lemoine and Traeger, 2016; Lontzek et al., 2015). Relative to those papers, our key contribution is to consider a more disaggregated model, with both regional and sub-regional income inequality, and the possibility of unequal distribution of climate impacts inside a given region. The paper is also related to the theoretical literature that analyzes the correlation between the risks on the benefits of the mitigation policy and the aggregate risk on consumption (Howarth, 2003; Sandsmark and Vennemo, 2007; Dietz et al., 2016; Lemoine, 2016). We contribute to this literature by quantifying the importance of this “climate risk premium” for each model and source of risk, and by analyzing its adjustment in the presence of inequality.

The paper proceeds as follows: Section 2 introduces the NICE model and reviews its key properties. Section 3 analyzes the interaction between risk and unequal damage incidence. This interaction is extended to the case of learning in section 4. Given the results about a low value of learning, section 5 explores the costs associated with delaying action. Finally, section 6 concludes.

2 The Model

The NICER Model (Nested Inequalities Climate Economy with Risk) introduces risk into the NICE model, which itself adds sub-regional inequality to the RICE model.² In addition to our optimisation results for NICER, we will also report results arrived at by aggregating the consumption data from NICER into per-period global per-capita consumption and maximising a representative agent’s ex-ante inter-temporal utility over this aggregate consumption stream. By this re-aggregation we recover a globally aggregated model à la DICE, which we hence refer to as DICE.³

In this section we review the key concepts associated with NICER, in particular the nature of the inequality in NICE and the extension to parameter uncertainty and the implications

²See footnote 1.

³Notice that this is not actually the same as the latest version of DICE published by William Nordhaus in 2013, but rather a re-aggregation of the model described here.

for the path of optimal carbon prices. Further detail about NICE can be found in Dennig et al. (2015).

2.1 Social Welfare and Unequal Damage Incidence

Our analysis is based on the maximisation of the social objective

$$W = \sum_t R^t \sum_r \frac{L_{rt}}{5} \sum_i \frac{\mathbb{E}_s[c_{irst}^{1-\eta}]}{1-\eta}$$

where W denotes social welfare, R the pure time discount factor, L population, and c per-capita consumption. The subscripts $i \in \{1, 2, \dots, 5\}$, $r \in \{1, 2, \dots, 12\}$, $s \in \{1, 2, \dots, S\}$, and $t \in \{1, 2, \dots, T\}$ denote regional quintile, region, state of the world, and time respectively. The parameter η governs both inequality aversion and risk aversion. The original RICE model consists of equations determining the evolution of regional per-capita consumption c_{rt} as a function of exogenous parameters and a tax policy vector. We extend this to two additional dimensions: sub-regional income groups and states of the world.

The basic NICE model described in Dennig et al. (2015) consists of the first extension and computes disaggregated per-capita consumption in quintile i , region r and period t as

$$c_{irt} = 5c_{rt} (q_{ir} + D_{rt} (q_{ir} - d_{ir})),$$

where D_{rt} is the damage in region r and period t , q_{ir} is the income share of the i th quintile in region r , and d_{ir} is the share of damage of the i th quintile in region r . In the absence of damages, $c_{irt} = 5c_{rt}q_{ir}$, so that q_{ir} is the fraction of total consumption allocated to quintile i in region r , thus reflecting exogenous consumption inequality (estimated using World Bank indicators — see Dennig et al. (2015) for details). By definition, $\sum_i q_{ir} = 1$. When damages are positive, the term $D_{rt}(q_{ir} - d_{ir})$ adjusts consumption in quintile i based on damages in the overall region. The parameter d_{ir} is computed as

$$d_{ir} = \frac{(q_{ir})^e}{\sum_i (q_{ir})^e}$$

where $e \in \mathbb{R}$ determines whether damages improve or worsen consumption inequality, and hence can be understood as governing the inequality of damage incidence. Following Dennig et al. (2015), we refer to e as the elasticity of damage with respect to income. For example, if $e = 1$, then $d_{ir} = q_{ir}$, damages are spread proportionally across quintiles, and hence do not

affect relative inequality. In this case NICE and RICE are very similar⁴. In contrast, if $e = 0$, then $d_{ir} = 1/5$ regardless of the distribution q_{ir} so that poorer quintiles lose a greater share of consumption due to damage than richer quintiles. In general, $e < 1$ implies that unequal damage incidence increases consumption inequality, while $e > 1$ means that unequal damage incidence lowers inequality. A key implication of this extension is that when $e < 1$, unequal damage incidence creates an additional force for carbon taxation since higher damages not only lead to lower total output in the future, but also higher inequality of consumption from today onwards, which lowers global social welfare for any level of positive inequality aversion, $\eta > 0$.

While there is not a wealth of evidence to support a particular value for e , we view $e < 1$ as the most plausible case empirically. Given poorer countries reliance on agriculture and proximity to the equator and coastal regions, it is very likely that lower income groups will face a disproportionate portion of the damages due to climate change (Anthoff and Tol, 2014). Furthermore, climate-related shocks such as natural disasters, spread of disease, and crop failure, are all more likely to be felt more intensely by lower income groups, as emphasized by the World Bank *Shockwaves* report (World Bank, 2016). Throughout the paper, we focus on the cases $e = 1$ and $e = 0$ to emphasize the role that unequal damage incidence plays in determining optimal carbon policies.

In addition to sub-regional inequality, in NICER we also add uncertainty to a number of the exogenous parameters of the RICE model. This is done by creating a state space $s \in \{1, 2, \dots, S\}$ of equi-probable outcomes, and drawing a value for the corresponding exogenous parameters in each state according to a distribution. The uncertainty propagates to per-capita income to yield a consumption variable c_{irst} . The parameters we consider uncertain are the TFP growth rate in each region, TFP convergence across regions, climate sensitivity, and the linear damage function coefficients. The details concerning the distributions of these parameters can be found in Appendix C.

In order to compare the effect of uncertainty in our model to the effect it would have without any inequality, we re-aggregate consumption to

$$c_{st} = \sum_r \frac{L_{rt}}{L_t} \frac{1}{5} \sum_i c_{irst}$$

⁴More precisely, the optimal policy computed with RICE is the same as in NICE if either $q_{ir} = \frac{1}{5}$ for all quintiles, or if $e = 1$ and the elasticity of marginal utility $\eta = 1$. The first equivalence is obvious and the second equivalence is a consequence of the invariance to proportional changes in distribution of the marginal utility of consumption when $\eta = 1$

The corresponding social objective is the population-weighted social welfare function

$$W^A = \sum_t R^t L_t \frac{\mathbb{E}_s[c_{st}^{1-\eta}]}{1-\eta}$$

2.2 The Optimal Carbon Tax

The NICER model is solved by letting a benevolent policy-maker choose the globally uniform tax path that maximizes social welfare, subject to constraints describing consumption and savings behavior, technology, and climate-economy interactions. Appendix B describes how to solve the model, and derives the optimal carbon tax. We now describe the key differences between the optimal taxes in the aggregate model à la DICE and in NICER.

Let $\lambda_{rt} \equiv \frac{\partial \Lambda_{rt}}{\partial \tau_t} \frac{1}{1-\Lambda_{rt}}$ be the proportional change in mitigation costs of region r at time t , Λ_{rt} , due to an increase in the carbon tax τ_t at time t . The increase in tax induces more mitigation, which reduces regional output by a proportion λ_{rt} . At the same time, the tax reduces total emissions E_t at time t , future atmospheric carbon stock M_j , for all $j \geq t$, and, through radiative forcing and feedbacks, future temperatures $\{T_j\}_{j \geq t}$. The change in future atmospheric temperature affects future regional output, which will increase by a proportion $\delta_{rsj} \equiv -\frac{1}{1+D_{rsj}} \frac{\partial D_{rsj}}{\partial M_{sj}} \frac{\partial M_{sj}}{\partial E_{st}} \frac{\partial E_{st}}{\partial \tau_t}$, where D_{rsj} denotes the damages suffered by region r in state s and time j .

In the standard aggregate model, for each period t , the optimal tax τ_t^A satisfies the following condition:

$$\tau_t^A : \quad \lambda_t = \sum_{j=t}^T R^{j-t} \frac{L_j}{L_t} \xi_j \left(\frac{\mathbb{E}_s[c_{sj}^{1-\eta} \delta_{sj}]}{\mathbb{E}_s[c_{st}^{1-\eta}] \delta_j} \right) \delta_j \quad (1)$$

where $\lambda_t = \sum_r \frac{L_{rt}}{L_t} \lambda_{rt}$ is the aggregate marginal, proportional mitigation cost, $\delta_j = \sum_r \frac{L_{rj}}{L_j} \mathbb{E}_s[\delta_{rsj}]$ the aggregate marginal, expected proportional damage, $\delta_{sj} = \sum_r \frac{L_{rj}}{L_j} \delta_{rsj}$ the aggregate marginal, proportional damages in state s , and $c_{sj} = \sum_r \frac{L_{rt}}{L_t} c_{rsj}$ the aggregate consumption at time j in state s . In every period, the optimal carbon tax equalizes its marginal cost λ_t to its marginal benefit, which is the discounted sum of future damage reductions δ_j , where we interpret $R^{j-t} \frac{L_j}{L_t} \xi_j \left(\frac{\mathbb{E}_s[c_{sj}^{1-\eta} \delta_{sj}]}{\mathbb{E}_s[c_{st}^{1-\eta}] \delta_j} \right)$ as the discount factor. The variable ξ_j takes into account the distribution of mitigation costs and damages across regions of the world. In the standard DICE model, this term would be equal to unity; in our implementation, instead, it will assume a value close to one because of the way in which the aggregate version of NICER

is constructed, as previously explained. Appendix B contains a definition of the variable ξ . Neglecting the ξ term, the discount factor depends on time discounting, population and income growth, the evolution of consumption risk over time, and the potential correlation between consumption and marginal damages. Note that, at time 0 (when, by assumption, there is no risk), the right-hand side of (1) corresponds to the current Social Cost of Carbon per unit of consumption, evaluated along the optimal emission trajectory.

In NICE and NICER, the regional and sub-regional consumption inequality affects the optimal carbon taxes in a non-trivial manner. The optimal tax τ_t^{NICER} satisfies the following condition:

$$\tau_t^{NICER} : \quad \lambda_t = \sum_{j=t}^T R^{j-t} \frac{L_j}{L_t} \left(\frac{\hat{\mathbb{E}}_s[c_{irsj}^{1-\eta} \tilde{\delta}_{irsj}]}{\hat{\mathbb{E}}_s[c_{irst}^{1-\eta} \lambda_{rt}]} \frac{\lambda_t}{\tilde{\delta}_j} \right) \tilde{\delta}_j \quad (2)$$

where $\hat{\mathbb{E}}_s[x] = \sum_r \frac{L_{rj}}{L_j} \frac{1}{5} \sum_i \mathbb{E}_s[x]$ denotes the average expected value of variable x , $\tilde{\delta}_{irsj}$ the proportion of marginal damages borne by quintile i in region r and state s at time j :

$$\tilde{\delta}_{irsj} \equiv \frac{d_{ir} \delta_{rsj}}{q_{ir} + D_{rsj}(q_{ir} - d_{ir})}$$

and $\tilde{\delta}_j = \hat{\mathbb{E}}_s[\tilde{\delta}_{rsj}]$ the expected average regional marginal damage as a proportion of consumption. As in the aggregate model, the optimal tax equalizes its aggregate marginal cost λ_t to the discounted sum of future aggregate expected mitigation benefits. Note that the future marginal reduction in damages $\tilde{\delta}_j$ now depends also on the degree of unequal damage incidence. Moreover, the discount factor assigned to each future marginal reduction in damages $\tilde{\delta}_j$ takes into account not only on the evolution of regional consumption growth and risk, but also on both the evolution of within-region consumption inequality, and the correlation between quintile-specific consumption and damages. The next section provides results that inform our understanding of how these new forces affect the path of optimal carbon taxes.

3 The Risk-and-Equity-adjusted Carbon Price

In this section we study the impact of risk on the optimal climate policy and its interaction with the presence of regional and sub-regional inequality. We consider four different types of risk: risk on the initial growth rate of TFP, risk on the rate of convergence of regional TFPs, risk on the climate sensitivity parameter, and risk on the linear parameter of the damage function. We first review the impact of each type of risk on the optimal tax path in the aggregate model, and then compare the results to the NICER model, highlighting the impact

of unequal damage incidence. We provide both a theoretical explanation of why we expect the tax to differ across models and types of risk, and present some calibration exercises to support our conclusions. All the calibrations assume a rate of impatience $\rho = 1.5\%$ and a coefficient of relative risk (and inequality) aversion $\eta = 2$. Following Dietz et al. (2016), we assume a Normal distribution for the initial growth rate of TFP, and a loglogistic distribution for the climate sensitivity parameter. The linear damage parameter is calibrated to match the distribution of damages estimated by Tol (2012). Finally, the rate of regional TFP convergence is assumed to follow a Beta distribution. Details of calibrations can be found in Appendix C.

3.1 The Impact of Risk in DICE

In order to understand the role of risk, we analytically disentangle the reasons why risk matters for the definition of the optimal policy. In particular, we focus on the sign and size of the “risk premium” induced by each source of uncertainty, which arises from the potential correlation between consumption risk and the risk on the benefits of mitigation.⁵

It is useful to note that the optimality condition for the DICE tax (1) can be rewritten as

$$\lambda_t = \sum_{j=t}^T e^{-r_j^t(j-t)} \delta_j$$

where r_j^t is the discount rate for marginal damages δ_j occurring at time j and baseline year t ,

$$r_j^t = -\ln R - \frac{1}{j-t} \ln \frac{L_j}{L_t} - \frac{1}{j-t} \ln \frac{\mathbb{E}_s[c_{sj}^{1-\eta}]}{\mathbb{E}_s[c_{st}^{1-\eta}]} - \underbrace{\frac{1}{j-t} \ln \frac{\mathbb{E}_s[c_{sj}^{1-\eta} \delta_{sj}]}{\mathbb{E}_s[c_{sj}^{1-\eta}] \delta_j}}_{\text{risk premium}} - \frac{1}{j-t} \ln \xi_j \quad (3)$$

The third term reflects the evolution of expected utility over time, and will depend on consumption growth, and on the change of consumption risk over time. The fourth term of (3) represents the climate “risk premium”, which indicates whether the risk on marginal damages δ_j is correlated with the risk on consumption. The last term adjusts for the possible presence of inequality in mitigation costs and damages across regions. For the sake of simplicity, we will neglect it in this section from now on.

To better understand the components of the discount rate, and their relative importance, let us take a second order Taylor approximation of its different elements around the expected

⁵See Dietz et al (2016) for a recent discussion about the climate “risk premium”

consumption $c_j = \mathbb{E}_s c_{sj}$ and the expected damage $\delta_j = \mathbb{E}_s \delta_{sj}$, for all j , which yields

$$r_j^t \simeq \rho - g_L + (\eta - 1)g_c - \frac{0.5\eta(\eta - 1)}{j - t} \left[V\left(\frac{c_{sj}}{c_j}\right) - V\left(\frac{c_{st}}{c_t}\right) \right] + \underbrace{\frac{\eta - 1}{j - t} Cov\left(\frac{c_{sj}}{c_j}, \frac{\delta_{sj}}{\delta_j}\right)}_{\text{risk premium}} \quad (4)$$

where $\rho = -\ln R$, $g_L = \frac{1}{j-t} \ln \frac{L_j}{L_t}$ denotes the population growth rate, $g_c = \frac{1}{j-t} \ln \frac{c_j}{c_t}$ the growth rate of expected consumption, $V(\cdot)$ the variance of consumption risk, and $Cov\left(\frac{c_{sj}}{c_j}, \frac{\delta_{sj}}{\delta_j}\right)$ the covariance between consumption and marginal damages.

Crucially, the sign of the covariance term depends on the type of risk that we are considering. We have the following relationships for the 4 risks that we study,

$$x = \{\text{TFP growth, GDP convergence rate}\} : \frac{\partial c_{sj}}{\partial x} > 0, \frac{\partial \delta_{js}}{\partial x} > 0 \Rightarrow Cov\left(\frac{c_{sj}}{c_j}, \frac{\delta_{sj}}{\delta_j}\right) > 0$$

$$x = \{\text{climate sensitivity, damage parameter}\} : \frac{\partial c_{sj}}{\partial x} < 0, \frac{\partial \delta_{js}}{\partial x} > 0 \Rightarrow Cov\left(\frac{c_{sj}}{c_j}, \frac{\delta_{sj}}{\delta_j}\right) < 0$$

To understand why the covariance will be positive for TFP growth risk and the risk on the GDP convergence rate, first note that higher TFP growth and higher convergence rates are both associated with higher consumption. More subtly, note that when output is higher, emissions are higher for any given path of taxes,⁶ so that marginal damages, δ_j , are higher given the convexity of the damage function. Therefore, at higher levels of TFP growth and convergence rates, both consumption and marginal damages are higher resulting in a positive covariance. In contrast, when the risk concerns the climate sensitivity parameter or the linear damage function parameter, the covariance will be negative. Since damages are increasing and convex in temperature, increasing the value of either parameter raises the marginal damage (higher climate sensitivity leads to higher temperatures, while a higher linear term increases the slope of the entire function). Since neither parameter directly impacts consumption, the only effect is for the higher damages to lower consumption. Hence, a higher climate sensitivity or linear damage parameter lead to higher damages but lower consumption.

Therefore, as long as $\eta > 1$, the policy-maker has an incentive to increase the carbon tax in the presence of risk both when consumption risk is increasing over time (we care more

⁶Since $\frac{\partial E_{st}}{\partial \tau_t} = -\sum_r \frac{L_{rt}}{L_t} \sigma_{rt} \mu_{rt} Y_{rst} \frac{1}{\tau_t(\theta_2 - 1)}$, the higher the initial income Y_{rst} , the larger the decrease in initial emissions per unit of tax, and, consequently, the larger the future benefits of mitigation.

about the future generation because it faces a larger risk), and when this source of risk makes the benefits of mitigation negatively correlated with the consumption level.

Figure 1 depicts the optimal tax paths for the aggregate model with and without risks. We can draw two main conclusions: 1) In the short run, there is a small hedging effect under all types of risk. Indeed, the presence of uncertainty slightly raises taxes with respect to the deterministic case. 2) In the long run, there is a distinction between growth risks and risks linked to the damage function. In the presence of risk on TFP growth and GDP convergence, the optimal tax is reduced compared to the deterministic case. If the risk concerns either the climate sensitivity or the damage parameter, taxes are always larger in the stochastic case compared to the deterministic one. Moreover, the increase is much more pronounced for the climate sensitivity risk than for the damage risk.

Figure (4) depict the main components of formula (1) for each source of risk and for two different tax periods (2015 and 2065): the expected marginal benefits δ_j , the risk premium affecting the discount rate, and the evolution of the value of consumption risk over time (the third term in (4)). The plots have been created by fixing the optimal deterministic tax path, and by applying it in the first order conditions (1) for each selected time periods and risk. Although risks in TFP growth and TFP convergence induce the larger variability in consumption (last plot in Figure 4), they imply a small increase in mitigation benefits compared to the deterministic case and create a relatively large climate risk premium (Figure 4). Moreover, the climate risk premium becomes even more important for policies set in the future, which, combined with a reduction in variability of the growth rate of consumption, justifies the long run behavior of the corresponding carbon taxes. On the contrary, the presence of risks on the damage function induces a very small variability in consumption across time, but relatively large negative climate risk premia. Moreover, in the presence of risk on climate sensitivity, the larger increase in taxes compared to the deterministic case is mainly due to the high benefits from mitigation, which increase even more sharply for long run policies, thereby explaining the long run path of the corresponding optimal tax.⁷

3.2 The impact of risk in NICER

Figures 2 and 3 represent the optimal tax path under the different sources of risk and different values of the elasticity of damages to income. Recall that when $e = 1$, damages are spread

⁷The comparison between discount rates can be done only once we choose a specific tax path. As a consequence, it only partially informs us about the impact of risk, as we should account also for the dynamic impacts of the tax on future periods. However, the static analysis already points us to the main forces driving the tax results.

proportionally across quintiles, whereas when $e = 0$ damages are equally divided across quintiles. As a consequence, when $e = 1$ climate impacts have no effect on consumption inequality, while when $e = 0$ climate change worsens consumption inequality as the poor segments of the population bear proportionally more impacts.

As in the aggregate model, each source of risk increases taxes in the short run. On the contrary, in the long run the technological risks call for lower taxes than in the deterministic case, while the risks on damages call for larger taxes. Moreover, as we reduce the elasticity e , we find three main effects: 1) taxes are constantly higher and peak earlier; 2) the long run tax reduction due to technological risks becomes imperceptible; 3) the risk on the damage parameter appears to have the most prominent impact on taxes.

To understand these results, we can again decompose the optimal tax condition (2) into a discounted sum of future expected average mitigation benefits $\tilde{\delta}_j = \hat{\mathbb{E}}_s \tilde{\delta}_{rsj}$:

$$\tau_t^{NICER} : \quad \lambda_t = \sum_{j=t}^T e^{-\tilde{r}_j^t(j-t)} \tilde{\delta}_j$$

where the discount rate \tilde{r}_j^t will depend on the evolution of both risk and inequality over time, and on the correlation between the consumption risk and the risk concerning the individual returns of the mitigation policy:

$$\begin{aligned} \tilde{r}_j^t = & -\ln R - \frac{1}{j-t} \ln \frac{L_j}{L_t} - \frac{1}{j-t} \ln \frac{\hat{\mathbb{E}}_s[c_{irsj}^{1-\eta}]}{\hat{\mathbb{E}}_s[c_{irst}^{1-\eta}]} + \\ & \underbrace{-\frac{1}{j-t} \ln \frac{\hat{\mathbb{E}}_s[c_{irsj}^{1-\eta} \tilde{\delta}_{rsj}]}{\hat{\mathbb{E}}_s[c_{irsj}^{1-\eta} \tilde{\delta}_j]}}_{\text{equity-adjusted risk premium}} + \underbrace{\frac{1}{j-t} \ln \frac{\hat{\mathbb{E}}_s[c_{irst}^{1-\eta} \lambda_{rt}]}{\hat{\mathbb{E}}_s[c_{irst}^{1-\eta} \lambda_t]}}_{\text{cost inequality premium}} \end{aligned} \quad (5)$$

As in the aggregate model, the third term represents the evolution of expected utility. Contrary to the previous case, it will reflect consumption growth and the change of both consumption risk and consumption inequality over time. The fourth term denotes the “risk premium”, adjusted by the presence of inequality. In particular, it will describe both correlations between aggregate consumption risk and aggregate benefits from mitigation, and also correlations between the more disaggregated quantities (i.e. regional and sub-regional). Finally, the last term represents the potential inequality in mitigation costs held by different regions, due to the differences in technology.

As before, we can take a second order Taylor approximation of the different terms around the average expected consumption $c_j = \hat{\mathbb{E}}_s[c_{irsj}]$, the average expected damages $\tilde{\delta}_j = \hat{\mathbb{E}}_s[\tilde{\delta}_{irsj}]$, and the average mitigation costs $\lambda_t = \hat{\mathbb{E}}_s[\lambda_{rt}]$:⁸

$$\begin{aligned} \tilde{r}_j^t \simeq & \rho - g_L + (\eta - 1)g_c - 0.5 \frac{\eta(\eta - 1)}{t - j} \left[\hat{V} \left(\frac{c_{irsj}}{c_j} \right) - \hat{V} \left(\frac{c_{irst}}{c_t} \right) \right] + \\ & \underbrace{\frac{\eta - 1}{t - j} \hat{Cov} \left(\frac{c_{irsj}}{c_j}, \frac{\tilde{\delta}_{irsj}}{\tilde{\delta}_j} \right)}_{\text{equity-adjusted risk premium}} + \underbrace{\frac{\eta - 1}{t - j} \hat{Cov} \left(\frac{c_{rt}}{c_t}, \frac{\lambda_{rt}}{\lambda_t} \right)}_{\text{cost inequality premium}} \end{aligned} \quad (6)$$

where $\hat{V} \left(\frac{c_{irsj}}{c_j} \right) = \hat{\mathbb{E}}_s \left(\frac{c_{irsj} - c_j}{c_j} \right)^2$ denotes the total consumption variability at time j (due to the presence of both risk and inequality), and $\hat{Cov}_j = \hat{\mathbb{E}}_s \left(\frac{c_{irsj} - c_j}{c_j} \right) \left(\frac{\tilde{\delta}_{irsj} - \tilde{\delta}_j}{\tilde{\delta}_j} \right)$ represents the covariance between consumption and damages at the quintile level (which is also due to the presence of both risk and inequality). Finally, the last term summarizes the relationship between the mitigation costs borne by region r and its average consumption level.

Let $V_j = \mathbb{E}_s \left(\frac{c_{sj} - c_j}{c_j} \right)^2$ denote the variance of aggregate consumption risk $c_{sj} = \sum_r \frac{L_{rj}}{L_j} c_{rsj}$, where $c_{rsj} = \frac{1}{5} \sum_i c_{irsj}$ is the average consumption in region r and state s at time j . Moreover, let $V_{sj} = \sum_r \frac{L_{rj}}{L_j} \left(\frac{c_{rsj} - c_{sj}}{c_j} \right)^2$ denote the degree of inequality across regions in state s , and $V_{rsj} = \frac{1}{5} \sum_i \left(\frac{c_{irsj} - c_{rsj}}{c_j} \right)^2$ the degree of inequality within region r in state s and period j . Then, the total variability at time j , $\hat{V}_j \equiv \hat{V} \left(\frac{c_{irsj}}{c_j} \right)$, can be decomposed into three main drivers:

$$\hat{V}_j = V_j + \mathbb{E}_s V_{sj} + \mathbb{E}_s \sum_r \frac{L_{rj}}{L_j} V_{rsj}$$

Thus, total consumption variability depends on the risk on aggregate consumption, the expected across-regions inequality, and the expected average within region inequality.

A similar disaggregation can be performed for the covariance term. Let $Cov_j = \mathbb{E}_s \left(\frac{c_{sj} - c_j}{c_j} \right) \left(\frac{\tilde{\delta}_{sj} - \tilde{\delta}_j}{\tilde{\delta}_j} \right)$ be the covariance between aggregate consumption c_{sj} and aggregate marginal damages $\tilde{\delta}_{sj} = \sum_r \frac{L_{rj}}{L_j} \tilde{\delta}_{rsj}$, where $\tilde{\delta}_{rsj} = \frac{1}{5} \sum_i \tilde{\delta}_{irsj}$ is the average marginal damage in region r and state s . Furthermore, let $Cov_{sj} = \sum_r \frac{L_{rj}}{L_j} \left(\frac{c_{rsj} - c_{sj}}{c_j} \right) \left(\frac{\tilde{\delta}_{rsj} - \tilde{\delta}_{sj}}{\tilde{\delta}_j} \right)$ be the covariance between average regional consumption and average regional damage in state s , and $Cov_{rsj} = \frac{1}{5} \sum_i \left(\frac{c_{irsj} - c_{rsj}}{c_j} \right) \left(\frac{\tilde{\delta}_{irsj} - \tilde{\delta}_{rsj}}{\tilde{\delta}_j} \right)$ the covariance between consumption and damages inside region

⁸Note that, by definition, $\hat{\mathbb{E}}_s x = \mathbb{E}_s x$ for $x = \{c_{irsj}, \lambda_{rt}\}$.

r in state s . Then, total covariance \hat{Cov}_j can be decomposed as:

$$\hat{Cov}_j = Cov_j + \mathbb{E}_s Cov_{sj} + \mathbb{E}_s \sum_r \frac{L_{rj}}{L_j} Cov_{rsj}$$

In other words, the total relation between consumption and damages depends on the relation between aggregate consumption and aggregate damages, the expected correlation between average consumption and average damages (on average, do poor regions suffer more than rich regions?), and the expected average correlation between consumption and damages within a given region (on average, at the regional level, do poor people suffer more than rich people?).

Let us go back to the interpretation of condition (6). As in the aggregate model, the discount rate positively depends on the rate of pure time preference $\rho = -\ln R$ and on the growth rate of expected average consumption g_c , while it decreases as population growth rate g_L increases. Moreover, it is negatively affected by increases in aggregate consumption risk and increases in inequality across time, either within regions or across regions. Finally, the incentives to mitigate increase if the damages tend to fall on poor regions and/or poor layers of the population and if they occur when the aggregate economy is in a bad state.

In a nutshell, the approximation (6) is telling us that, in a disaggregated world, we should care both about the risk faced by the future generations, but also about the level of inequality forecast for the future. Moreover, we should not only care whether the mitigation policy has an insurance role (i.e. if it displays larger benefits in the bad states of nature), but also if mitigation has an impact on inequality, either across regions or within regions. If we find out that the poorest population groups will suffer the most from climate change, climate policy has the ancillary benefit of improving their quality of life, thereby reducing worldwide inequality. The incentives to set a high tax are mitigated if poor regions pay the largest share of mitigation costs (last term of 6).

Thus, we are interested, first of all, in the sign of the covariance terms and in the evolution of inequality in a deterministic scenario, and then how those signs change once we introduce risk. Figures 5 and 6 portray the main elements describing the optimal carbon tax when the damage elasticity is, respectively, 1 and 0. The exercise is similar to the one realized in the aggregate model. For each model, we have fixed the optimal deterministic carbon tax, and we have applied it to the first order condition (2) for two time periods (2015 and 2065), and each type of risk. In each plot, the first line portrays the expected reduction in damages for the different assumptions about risk. Note that these benefits are comparable to those in the aggregate case when the elasticity is 1, while they are a bit larger with a 0 elasticity because of

the greater benefit on the poor. The last line represents the evolution of the value of risk (as in the aggregate case) and the evolution of the value of inequality (their sum corresponds to the fourth term in 6). Total inequality is decreasing over time independently of the elasticity, while the value of risk depends on the source of variability, with technological risks playing a major role. Finally, the second line depicts the equity-adjusted risk premium, which describes whether mitigation has an insurance and/or inequality contraction role. Contrary to the aggregate case, this premium is negative for all types of risk, thereby indicating that inequality matters for the results. Moreover, the size of this premium is ten times larger when the elasticity is 0 with respect to the other case.

Figures 7 and 8 decompose the equity-adjusted risk premium into two terms, representing, respectively, the insurance component of mitigation (the pure climate risk premium, which depends exclusively on Cov_j) and its role in reducing worldwide inequality (an “inequality premium”, which depends by the sign of the remaining covariance terms). If the inequality premium is negative, it means that a stringent mitigation target has the advantage of reducing inequality. The calibration exercise highlights the veracity of that argument. Inequality reduces a little bit in the NICER model with elasticity equal to 1, and by a fairly large amount when the elasticity is 0. Moreover, that is valid also in a deterministic world.

Figure 2 shows that, despite the negative inequality premium, taxes are not substantially altered by the presence of risk when the damage elasticity is equal to 1, similarly to the aggregate model. This is due to the decline of inequality over time and the fact that risk effects are comparable to the aggregate model. In contrast, the stronger reaction in the case of 0 damage elasticity and damage risk is explained by the larger inequality premium induced by this type of uncertainty. Under unequal damage incidence, a risk on the damage parameter implies that there are some states of nature where the poor are particularly badly hit by the negative consequences of climate change. Therefore, the mitigation policy is going to prevent especially those very bad states of nature. The consequence is a higher carbon tax and a larger expected effect on worldwide inequality.

3.3 Risk and Inequality Compared and Combined

In the previous sections, we have analyzed the impact of different sources of risk given a pre-determined level of aggregation. Here, we focus on the combined effect of risk and inequality on the optimal tax path.

First of all, by comparing Figure 1 and Figure 2, we can see how the combination of more disaggregation and proportional damages has a negligible effect on the optimal policy, and,

more importantly for our analysis, on the impact of risk. Figure 9 compares the optimal trajectory path in the deterministic case across the three models that we have considered so far. Moving from the aggregate model to NICER with elasticity equal to 1 does increase taxes, mainly in the long run. Short run effects, instead, are almost imperceptible. The same pattern appears once we introduce risk: taxes are larger in the NICER model, but the distance with respect to the aggregate one is the same as in the deterministic case.

Once we allow for lower levels of the elasticity parameter e , we find that taxes are much larger compared to the aggregate model both in the deterministic and stochastic cases (see, e.g., Figure 9 for the deterministic path). Moreover, the presence of a zero elasticity magnifies the tax increase due to risks on the damage function and nullifies the reduction induced by technological risks. Taking as a reference the deterministic tax path of the aggregate model, Table 1 describes by how much the optimal tax increases when we combine risk and inequality with respect to the situation in which we take the two phenomena separately and we simply sum the respective increases in tax. In other words, for each year t , we have first computed the change in tax in the aggregate model when we have risk with respect to no risk case. Then, we have computed the change in tax in the deterministic NICE model with respect to the deterministic aggregate one. Finally, the change in tax in the stochastic NICER model with respect to the deterministic aggregate one. For each type of risk, Table 1 shows the difference between the last change in tax and the sum of the former two. Values are presented only for two years, 2025 and 2055 (period close to the peak in the NICER model with elasticity 0). When the value is positive, it means that we have positive complementarity between risk and inequality: the tax increases more when the two sources of variability are combined than when taken separately. The results confirm that, when the elasticity is equal to 1 risk has approximately the same impact in NICER as in the aggregate model. In contrast, once we introduce unequal damage incidence, the impact of risk sharply increases, especially in the presence of damage risk.

4 Learning

Although uncertainty is large over many important parameters that affect climate policy, much money is spent in order to improve our understanding of the mechanisms at work, so uncertainty is likely to decrease over time. In order to model the effect of learning about parameters on optimal policy we introduce a learning node at which the social planner's information set is split into two equi-probable halves. That is, at some future date the

social planner knows if an uncertain parameter is drawn from the lower half of its ex-ante probability mass or from the greater half. (This is similar to what is done in Keller et al. 2006). The social planner thus chooses two tax paths which coincide before the learning node and diverge after it, reflecting the more precise information she has at her disposal from then onwards. We consider two different learning periods, today and 2045. We will also compare the results to the no-learning case studied in the previous section.

Figures 10 and 11 display the results of the learning exercise for each type of risk, and for the damage elasticity equal to one and zero. When the elasticity equals 1, the optimal taxes display a pattern similar to that documented by Keller et al. (2006): upon learning, the taxes immediately jump to the paths that are optimal for the case in which we learnt at the start of the model. Perhaps surprisingly, the same results go through for the case in which the damage elasticity equals zero. Although the taxes are generally larger in magnitude (as we found in the previous section), the taxes exhibit the same jump behavior, indicating that there is little interaction between learning and unequal damage incidence, and that both effects manifest as separate phenomena.

However, the figures also reveal again an interesting difference between the TFP and convergence rates on the one hand, and the climate sensitivity and damage risks on the other hand. It is particularly visible when e equals zero. For the latter risks, the bad news scenarios (high climate sensitivity, high damages) require a greater tax over the whole period, so that when this is learned in 2045 the tax jumps to that level, after having followed the path of the tax that is optimal when learning never occurs. In contrast, for the former risks, the taxes when learning occurs immediately have crossing paths, with the tax for the high-growth scenario being postponed compared to the low-growth scenario (due to the fact that the future is richer and needs and is able to pay for greater mitigation). On the figures, the crossing happen to occur before learning for TFP growth, and at the same time for the convergence rate.

We also compute the value of learning in terms of the proportion of 2005 consumption the social planner would consider equivalent, in absence of learning, to the ability to learn earlier. We compare learning now to learning in 2045, and learning in 2045 to never learning.

Formally, let W_t^{nl} and W_t^l be, respectively, the aggregate welfare at time t when there is no or late learning (nl) and when there is (early) learning (l). Recall that in the NICER model, $W_t = R^t \sum_{r=1}^{12} \frac{L_{rt}}{5} \sum_{i=1}^5 \frac{c_{irt}^{1-\eta}}{1-\eta}$, where c_{irt} is the consumption of quintile i in region r at

time t . Social welfare is equal to $\mathbb{E}_s \left[\sum_{t=0}^T W_t \right]$. Then, the value of learning α_L satisfies:

$$(1 + \alpha_L)^{1-\eta} W_0^{nl} + \mathbb{E}_s \left[\sum_{t=1}^T W_t^{nl} \right] = \mathbb{E}_s \left[\sum_{t=0}^T W_t^l \right]$$

Solving with respect to α_L yields:

$$\alpha_L = \left(1 + \frac{\mathbb{E}_s \left[\sum_{t=0}^T W_t^l \right] - \mathbb{E}_s \left[\sum_{t=0}^T W_t^{nl} \right]}{W_0^{nl}} \right)^{\frac{1}{1-\eta}} - 1 \quad (7)$$

The results for α_L (in percent) are reported in Table 2. In general, unequal damage incidence increases the value of learning. However, except for damage uncertainty and inversely proportional damage ($\eta=-1$), where the value is almost 4% of consumption, the values are quite small. For proportional damage the values are negligible, which is in line with the existing literature. This reflects the fact that the differences between optimal taxes under learning and optimal taxes without learning do not result in significant changes in climate or inequality. However, this result hinges on the assumption that before any learning takes place, taxes are still optimized, and importantly, are well above zero. In the next section, we investigate the significant consequences of relaxing this assumption.

It must be emphasized that these results depend on the particular calibration of the risks. If we introduced more dramatic catastrophic risks in the analysis, the value of learning could be much higher.

5 The Cost of Delaying Action

An alternative strategy to taking action today and adjusting in the face of learning, is to simply wait and see how the climate evolves, delaying action until some later date when more information has been collected. In this section, we explore the social welfare cost of delaying action for the next 30 years, i.e. continuing along a business as usual path till 2045, and then introducing a globally optimal carbon price.

5.1 Computing the Cost of Delay

Figures 12 and 13 show the optimal tax paths in the deterministic scenarios for each NICE model when the taxes are constrained to 0 for the first three decades, and compares them to

the unconstrained solution. The graphs show that the optimal policy is almost insensitive to the initial constraints in the NICE model with elasticity equal to 1. Indeed, taxes jump to the unconstrained path immediately after the delay period, and there is only a very small compensation effect afterwards (post-delay taxes are imperceptibly larger than the unconstrained taxes). In the zero elasticity case, however, the optimal constrained policy compensates the delay by setting visibly higher taxes than in the case without delay. The divergence in effects hinges on the contrasting consequences of delay for the climate and for inequality respectively. It is well known that the climate system is calibrated to be very inert, with effects taking many periods to accumulate. Therefore, while three decades of delay does lead to increased emissions, this additional flow does not significantly affect the stock of CO₂ in the atmosphere and hence does not increase damages by much at all. Hence, delay proves to be fairly inconsequential for the climate. However, the increase in emissions that results from delay immediately translates into higher inequality when the elasticity of damage is less than 1. When the elasticity is equal to zero, inequality is significantly more sensitive to the increase in damages than climate variables are. Therefore, post-delay, optimal policy features overshooting behavior to counteract the increase in consumption inequality brought about by the period of delay. Although not reported, we find exactly the same tax pattern when we introduce parameter uncertainty.

In order to properly quantify the cost of delay, we use the consumption streams implied under each tax policy. Define the cost of delay as the percentage of current consumption that the present generation should receive to compensate them for the social welfare loss due to delaying action. Formally, let W_t^d and W_t^{nd} be, respectively, the aggregate welfare at time t when there is delay (d) and when there is no delay (nd). As in the previous section about learning, we can compute the cost of delay as the value of α_D that satisfies:

$$(1 + \alpha_D)^{1-\eta} W_0^d + \mathbb{E}_s \left[\sum_{t=1}^T W_t^d \right] = \mathbb{E}_s \left[\sum_{t=0}^T W_t^{nd} \right]$$

Solving with respect to α_D yields:

$$\alpha_D = \left(1 + \frac{\mathbb{E}_s \left[\sum_{t=0}^T W_t^{nd} \right] - \mathbb{E}_s \left[\sum_{t=0}^T W_t^d \right]}{W_0^d} \right)^{\frac{1}{1-\eta}} - 1 \quad (8)$$

Table 4 summarizes the cost of delay for the DICE and NICE models with different elasticities. It is clear that as the elasticity goes down, i.e., when the poor bear the larger share

of the climate damages, the cost of delay increases exponentially. For example, when the elasticity is equal to 0, i.e., when damages are equally distributed across the population, the cost of delaying action is ten times larger than the corresponding cost in the DICE model (and NICE model with elasticity equal to 1). On the contrary, the cost in the NICE model with elasticity equal to 1 is slightly lower than the cost in DICE. We therefore conclude that it is not the introduction of consumption inequality alone that makes the waiting strategy more costly, but crucially the combination of consumption inequality and unequal damage incidence that drives significant costs of delay. Moreover, notice how these conclusions hold for both the deterministic and the uncertainty cases. Although the cost of delay tends to be larger in the presence of risk than in the deterministic case for all models, it dramatically increases once we lower the elasticity.

In order to understand the main drivers of the cost of delay, and why it differs across models, we can look at its first order approximation, which expresses the cost of delay as the net present value of future consumption gains/losses when delaying, where each future consumption variation is multiplied by an appropriate discount factor.⁹ Appendix D contains the details of the approximation and its disaggregation. Let us rewrite total welfare at time t as $W_t = R^t L_t u(EEDE_t)$, where EDE_t represents the expected equally distributed amount of consumption that gives the same welfare as the set of consumptions uncertainty and unequally distributed both across and within regions¹⁰. By employing the EEDE concept, the approximated cost of delay is equal to:

$$\alpha_D \simeq - \sum_{t=0}^T R^t \frac{L_t}{L_0} \left(\frac{EED E_t^{nd}}{EED E_0^{nd}} \right)^{1-\eta} \frac{EED E_t^d - EED E_t^{nd}}{EED E_t^{nd}} \quad (9)$$

where $EED E_t^d$ is the equally distributed equivalent when delaying, and $EED E_t^{nd}$ when not delaying. The differences in cost of delay across models will depend both on differences in consumption variations due to delay and on differences in the size and evolution of the discount factor. Figures 14 and 15 represent the temporal evolution of the proportional change in consumption due to delay (the last term in 9) in the deterministic case for the NICER model with elasticities equal to 1 and to 0. Each plot depicts both the total proportional variation in consumption (the gray line) and the evolution of its main components (see Appendix D for details). Indeed, the proportional change in EEDE due to delay can be

⁹The approximation slightly underestimates the cost of delaying the policy. See Gollier (2011) for a quantification of the error introduced when using the approximation instead of the true welfare change.

¹⁰By definition, the expected equally distributed consumption at time t is such that: $u(EDE_t) = \mathbb{E}_s \sum_r \frac{L_{rt}}{L_t} \frac{1}{5} \sum_i u(c_{irst})$

decomposed into a change in average expected consumption and a change in the size of both risk and inequality. Average expected consumption changes mainly for two reasons: climate damages increase due to the higher carbon concentration, which will reduce consumption (green line); delay reduces mitigation costs in the short run, and, potentially, increases them in the long run due to the higher effort required once the policy is introduced (red line). In the DICE model, these two curves explain the total consumption variation as there is no inequality by construction. In the NICER model, instead, delay will affect also the degree of inequality both across regions and within each region (and the size of the risk if we were considering one of the stochastic scenarios). In the NICER model with elasticity equal to 1, only the inequality between regions is worsened by delay (the purple line): given that poor regions are more exposed to the adverse consequences of climate change, delay will further increase regional differences in consumption possibilities. The inequality within regions is not affected since, by construction, damages are distributed proportionally to income inside each region, i.e. climate impacts have no effect on consumption inequality at the regional level. On the contrary, in the NICER model with elasticity equal to 0, delay increases also inequality within regions on average, since, by construction, the poorest segments of the population bear the brunt of climate change.

Why is the cost of delay much higher in the NICER model with elasticity equal to 0? As Figure 15 shows, delaying even for only a few decades has huge impacts on worldwide inequality: the change in inequality explains at least three fourths of the proportional change in consumption. Moreover, note that the variations in damages and in mitigation costs are magnified by the presence of unequal damage incidence. The larger change in mitigation costs is due to the fact that the policy is more stringent in the NICE model with elasticity equal to 0 (see Figures 12 and 13); similarly, the slightly larger increase in damages is due to the fact that we are far away from the optimal policy.

So far we have discussed the components of the cost of delay in the deterministic case. If we introduce parametric uncertainty, delaying the policy will affect not only the degree of inequality, but also the size of the risk. Table 4 shows that the presence of risk increases the cost of delay, especially in the presence of risks related to the climate module. Figures 16 and 17 depict the decomposition of the change in EEDE due to delay for the NICER model with elasticity 0 and, respectively, TFP risk and damage risk. The main difference between the two plots concerns the evolution of the aggregate risk component. As previously seen, damages are positively correlated with the aggregate state of the economy when we have a technological risk. As a consequence, delaying the policy reduces the aggregate consumption

risk. On the contrary, when the source of risk is related to the damage function, mitigation has an insurance component. Therefore, delay increase the aggregate risk faced by the population. The larger cost of delay with damage risk is due exactly to the loss of insurance.

5.2 Regional cost of delay

Another interesting exercise is to look at which regions suffer more from delay, and whether there are some regions that are actually better off. We define the regional cost of delay as the percentage amount of consumption α_D^r that we should give to the current generation to compensate for the future losses due to delay:

$$\alpha_D^r = \left(1 + \frac{\sum_{t=0}^T \mathbb{E}_s U_t^{nd} - \sum_{t=0}^T \mathbb{E}_s U_t^d}{U_0^d} \right)^{\frac{1}{1-\eta}} - 1 \quad (10)$$

with $U_t = R^t L_{rt}^{\frac{1}{5}} \sum_{i=1}^5 \frac{c_{irst}^{1-\eta}}{1-\eta}$. Figure 24 shows the regional costs of delay for the NICE model with elasticity equal to 1 and to 0.¹¹ As the elasticity decreases, the cost rapidly increases for Africa, Europe and Middle East. At the same time, some countries are clearly better off by delaying as the elasticity reduces, in particular Russia, Non-Russia Eurasia and India. We have performed an analysis similar to the one in the aggregate case by decomposing the cost of delay into its main drivers, i.e. the change in regional mitigation costs, regional damages, within region inequality and regional risk. Figures from 18 to 23 show the results for three regions in the deterministic case: USA, India and Africa. For example, the huge increase in costs for Africa is due to the huge impact on the inequality index. The benefit of delaying for India, instead, depends on a combination of very high savings in mitigation costs and relatively low impact on inequality.

5.3 Why is the cost of delay much larger than the value of learning?

It is clear that unequal damage incidence is a powerful mechanism to increase the cost of delaying action against climate change. It is therefore puzzling that the same mechanism is so much more muted in the context of learning. In particular, it is potentially unexpected that unequal damage incidence does not lead to significantly higher values of learning than

¹¹Note that the aggregate cost of delay differs from the weighted sum of its regional counterparts because of the concavity of the utility function. See the Appendix for details.

in the case of equal incidence. However, this puzzle can be resolved by examining how the climate itself evolves in each case.

Figure 25 shows the paths of temperatures (measured as the difference from the pre-industrial era) under TFP risk, and in the case of delay and learning respectively. For each case, 3 curves are plotted, representing 3 of the possible uncertainty realizations (deciles 1, 5, and 10) for TFP growth. Two key results stand out. First, all the “delay” curves lie above the “learning” curves, indicating that optimal policy under delay always leads to higher temperatures than optimal policy under learning does. Second, the gap in corresponding curves peaks at around 0.6 degrees, which is a substantial difference. Furthermore, this gap is very persistent, which means that damages under delay will be significantly higher given the convexity of the damage function.

In light of this analysis, it becomes clear why delay is so much more empirically significant than learning. Under learning, the optimized taxes in the pre-learning period result in much lower temperatures over the whole model time frame than in the delay case which features zero taxes for three periods. These lower temperatures mean that damages are much more benign, and indeed much closer to the damages inflicted in the full information case, which results in the value of learning being very small.

6 Conclusion

In this paper, we have re-examined three key issues related to the uncertainty surrounding climate change through the lens of a model that accounts for sub-regional inequality and unequal damage incidence. Our analysis emphasizes the importance of accounting for unequal damage incidence, given its empirical likelihood, and the significant interactions with climate mechanisms that we have uncovered here.

The issues we discuss stem from the more basic question of assessing the social welfare cost of delaying mitigation, an issue of growing importance given the sluggish international process of the UNFCCC and the volatile politics of some countries. We hope that our results provide further evidence that optimal responses to climate change need to be formulated soon rather than later, and should take account of the large disparities in who bears the burden of climate damages.

It seems particularly important that delaying is much more costly for social welfare than optimizing under ignorance of the value of parameters. Even if ignorance makes the carbon price differ substantially from its optimal value, the difference remains considerably smaller

than the difference between the optimal price and business as usual. This conclusion is however contingent on specific calibrations of the risks, and specific forms of learning. The presence of catastrophic risks and more effective types of learning (narrowing the range of possible parameter values more than in our simulations) would raise the value of learning. Our model is designed as a tool that can be easily shared to explore such issues more extensively.

The question of insurance mechanisms that help populations to adapt to climate change impacts could be studied in more detail. Our results only indirectly provide evidence of a trade-off between adaptation and mitigation through the fact that unequal damage incidence raises the optimal carbon prices. Insofar as adaptation protects the most vulnerable populations, it reduces the phenomenon of unequal damage incidence. But a more explicit study of insurance mechanisms would require introducing stochastic shocks in the model instead of, or beside, structural uncertainty.

Another direction of research would explore variations in the social objective, and examine how dependent our results are on ethical assumptions embodied in the social welfare function. In particular, a lower pure time preference and a social objective that separates inequality aversion and risk aversion would be worth exploring.

References

- [1] Anthoff D, Hepburn C, Tol RSJ (2009), Equity weighting and the marginal damage costs of climate change, *Ecological Economics*, 68: 836-849
- [2] Anthoff D, Tol RSJ (2014), The income elasticity of the impact of climate change. Is the Environment a Luxury?: An Inquiry into the Relationship Between Environment and Income, eds Tiezzi S, Martini C (Routledge, London), pp 34-47.
- [3] Cai, Y., K.L. Judd and T. S. Lontzek (2013), The Social Cost of Stochastic and Irreversible Climate Change, Working Paper 18704, National Bureau of Economic Research
- [4] Crost, B. and C. Traeger (2014), Optimal CO2 Mitigation under Damage Risk Valuation, *Nature Climate Change*, 4:631-636
- [5] Daniel, K.D., R.B. Litterman and G. Wagner (2016), Applying Asset Pricing Theory to Calibrate the Price of Climate Risk, NBER Working Paper n. 22795
- [6] 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 of the United States of America*, 112(52): 15827:32
- [7] Dietz, S., C. Gollier and L. Kessler (2016), The Climate Beta, Working Paper No. 190, Grantham Research Institute on Climate Change and the Environment
- [8] Gerlagh, R. and M. Liski (2017), Carbon Prices for the Next Hundred Years, *The Economic Journal*, doi:10.1111/eoj.12436
- [9] Gollier, C. (2011), Discounting and Risk-Adjusting Non-Marginal Investment Projects, *European Review of Agricultural Economics*, 38(3): 325-334
- [10] Golosov, M., J. Hassler, P. Krusell and A. Tsyvinski (2014), Optimal Taxes on Fossil Fuel in General Equilibrium, *Econometrica*, 82:41-88
- [11] Hallegatte, Stephane, Mook Bangalore, Laura Bonzanigo, Marianne Fay, Tamaro Kane, Ulf Narloch, Julie Rozenberg, David Treguer, and Adrien Vogt-Schilb, (2016) Shock Waves: Managing the Impacts of Climate Change on Poverty, Climate Change and Development Series. Washington, DC: World Bank. doi:10.1596/978-1-4648-0673-5. License: Creative Commons Attribution CC BY 3.0 IGO

- [12] Hwang, I.C., F. Reynés and R.S.J. Tol (2014) The Effect of Learning on Climate Policy under Fat-Tailed Uncertainty, Munich Person RePEc Archive Paper 53681.
- [13] Howarth, R. B. (2003), Discounting and Uncertainty in Climate Change Policy Analysis, *Land Economics*, 79, 369-381.
- [14] Intergovernmental Panel on Climate Change (2014), Climate Change 2014: Mitigation of Climate Change, Working Group III contribution to the Fifth Assessment Report (IPCC WG III AR5)
- [15] Jensen, S. and C. Traeger (2014), Optimal Climate Change Mitigation under Long-Term Growth Uncertainty: Stochastic Integrated Assessment and Analytic Findings, *European Economic Review*, 69(C): 104-25
- [16] Jensen, S. and C. Traeger (2016), Pricing Climate Risk, mimeo
- [17] Kelly, D.L. and C.D. Kolstad (1999), Bayesian Learning, Growth and Pollution, *Journal of Economic Dynamics and Control*, 23(4): 491-518
- [18] 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
- [19] Leach, A.J. (2007), The Climate Change Learning Curve, *Journal of Economic Dynamics and Control*, 31(5): 1728-52
- [20] Lemoine, D. (2016), The Climate Risk Premium: How Uncertainty Affects the Social Cost of Carbon, University of Arizona Working Paper 15-01
- [21] Lemoine, D. and C.P. Traeger (2016), Economics of Tipping the Climate Dominoes, *Nature Climate Change*, 6:514-20
- [22] 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: 441-444.
- [23] Nordhuas, W.D. and P. Sztorc (2013), DICE 2013R: Introduction and User's Manual
- [24] Sandsmark, M. and H. Vennemo (2007), A portfolio Approach to Climate Investments: CAPM and Endogenous Risk, *Environmental and Resource Economics*, 37, 681-695.

- [25] Traeger, C. (2015), Analytic Integrated Assessment and Uncertainty, mimeo, Department of Agricultural and Resource Economics, UC Berkeley

A Tables and Graphs

Table 1: Complementarity between risk and inequality (%)

Year	NICER $e = 1$				NICER $e = 0$			
	TFP	Conv. Rate	Clim. Sens.	Dam.	TFP	Conv. Rate	Clim. Sens.	Dam.
2025	-8.65	-2.8	-9.62	1.57	45.63	75.4	76.85	193.65
2055	-3.88	-1.08	-3.27	1.74	37.62	54.02	92.33	196.96

Each number equals the difference between (a) the % change in the optimal tax when we move from DICE deterministic to NICER and (b) the sum of the separate effects of each move (from DICE deterministic to NICE deterministic + from DICE deterministic to DICE stochastic). This is done for elasticity of damage equal to either 1 or 0.

Table 2: Value of learning (%)

Elasticity	Uncertainty	V2015to2045	V2045to2325
1.0	TFP	0.0	0.004
1.0	Clim Sense.	0.005	0.019
1.0	Conv. Rate	0.001	0.003
1.0	Damage	0.001	0.001
0.0	TFP	0.001	0.008
0.0	Clim Sense.	0.062	0.088
0.0	Conv. Rate	0.011	0.002
0.0	Damage	0.667	0.513
-1.0	TFP	0.005	0.0
-1.0	Clim. Sense.	0.331	0.051
-1.0	Conv. Rate	0.037	0.0
-1.0	Damage	3.916	0.53

The value of learning is computed as the increase % in 2005 consumption the social planner would consider equivalent, in absence of learning, to the ability to learn earlier.

Table 3: Value of learning, ten branches (%)

Elasticity	Uncertainty	V2015to2045	V2045to2325
1.0	TFP	0.0	0.008
1.0	Clim. Sense.	0.009	0.031
1.0	Conv. Rate	0.002	0.006
1.0	Damage	0.05	0.071
0.0	TFP	0.002	0.011
0.0	Clim. Sense.	0.099	0.148
0.0	Conv. Rate	0.015	0.004
0.0	Damage	1.004	0.67
-1.0	TFP	0.007	0.0
-1.0	Clim. Sense.	0.502	0.141
-1.0	Conv. Rate	0.054	0.0
-1.0	Damage	4.665	0.877

The value of learning is computed as the increase % in 2005 consumption the social planner would consider equivalent, in absence of learning, to the ability to learn earlier. These are the values when, upon learning, the social planner learns about the which decile the parameter is drawn from (rather than which half of the parameter space).

Table 4: Cost of delaying action until 2045 (%)

	DICE	NICER			
		$e = 1$	$e = 0.5$	$e = 0$	$e = -1$
Deterministic	0.15	0.13	0.44	1.45	7.98
TFP risk	0.2	0.16	0.53	1.81	11.85
Clim. Sens. risk	0.22	0.17	0.59	2.08	22.14
Damage risk	0.16	0.14	0.54	2.9	61.62
Conv. Rate risk	0.2	0.17	0.6	2.02	13.28

The cost of delay is computed as percentage of 2005 consumption that the present generation should receive to compensate them for the social welfare loss due to delaying action.

Figure 1: Taxes across all risks, aggregate




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




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




Figure 4: Tax decomposition, aggregate

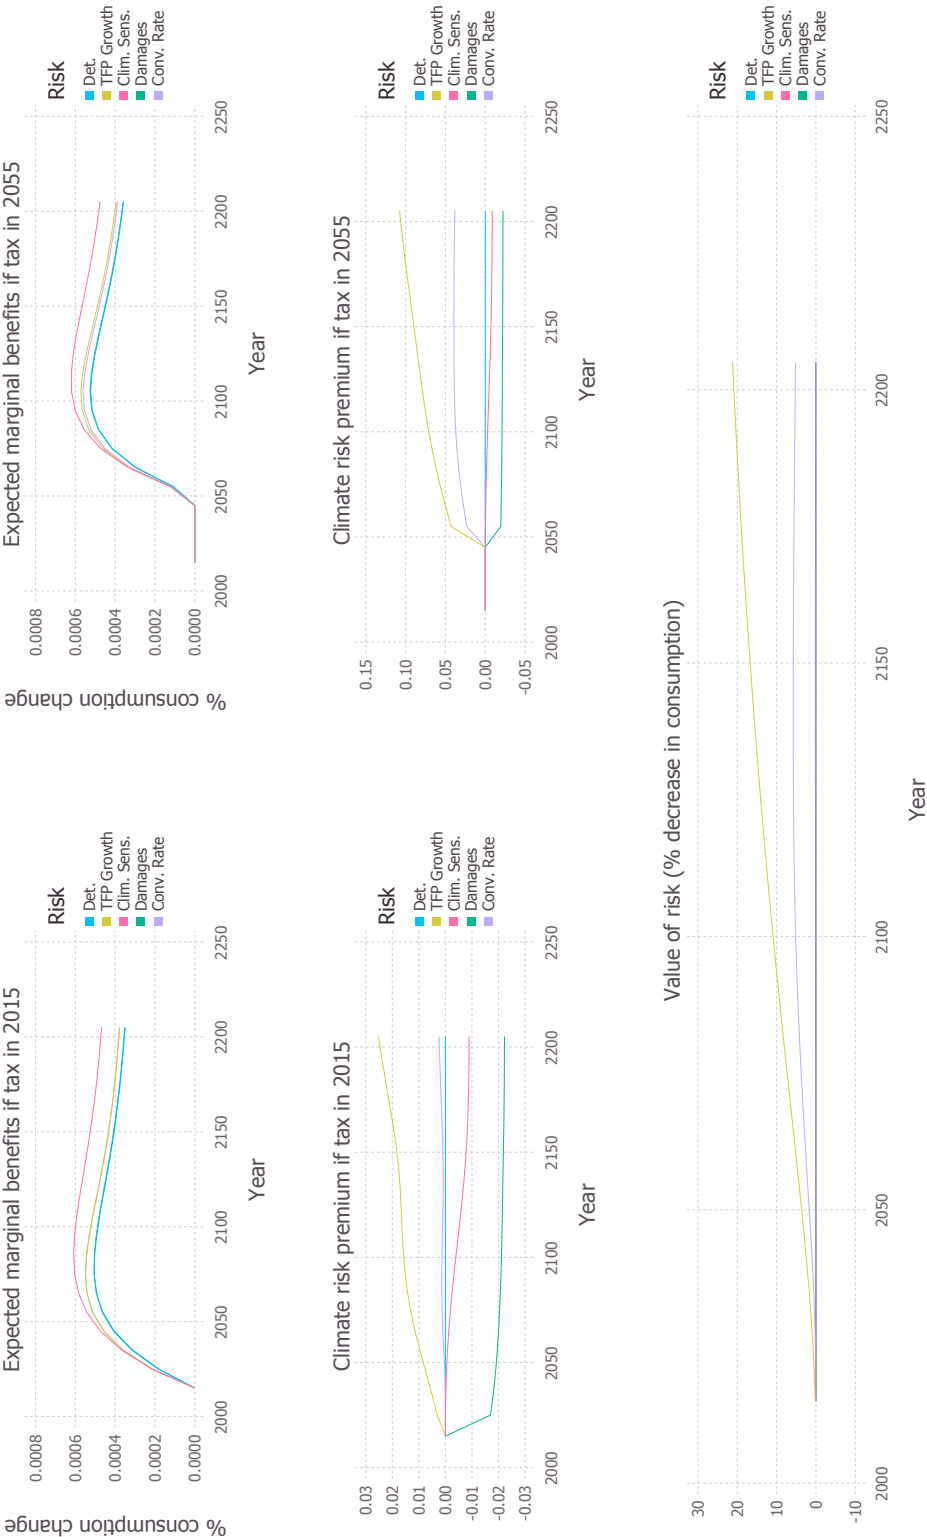


Figure 5: Tax decomposition, NICER $e = 1$

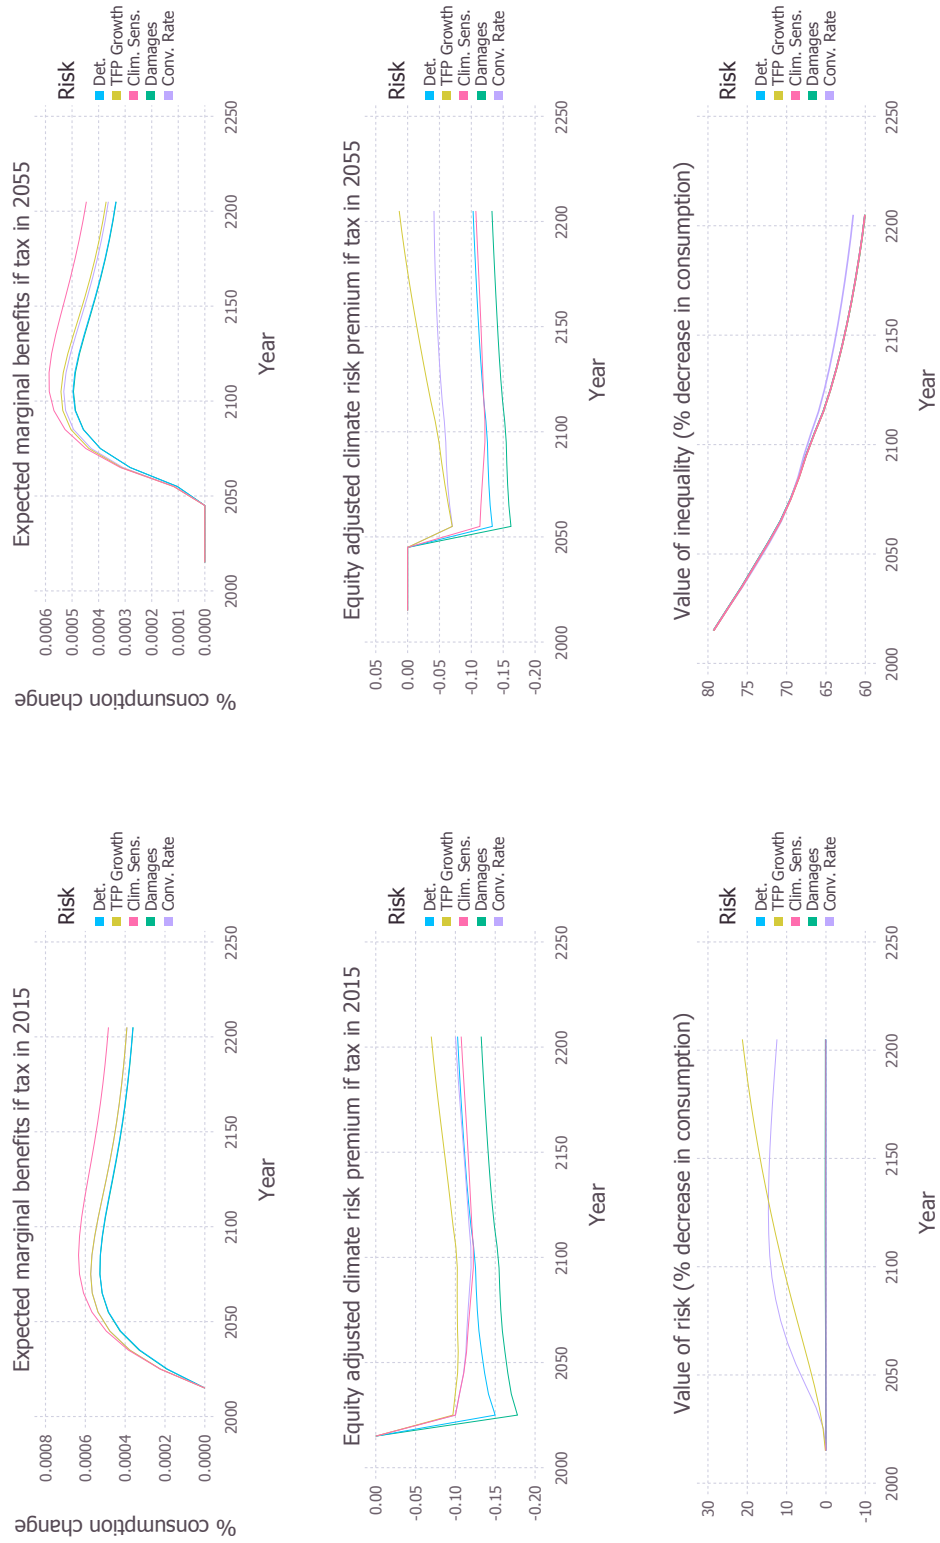


Figure 6: Tax decomposition, NICER $e = 0$

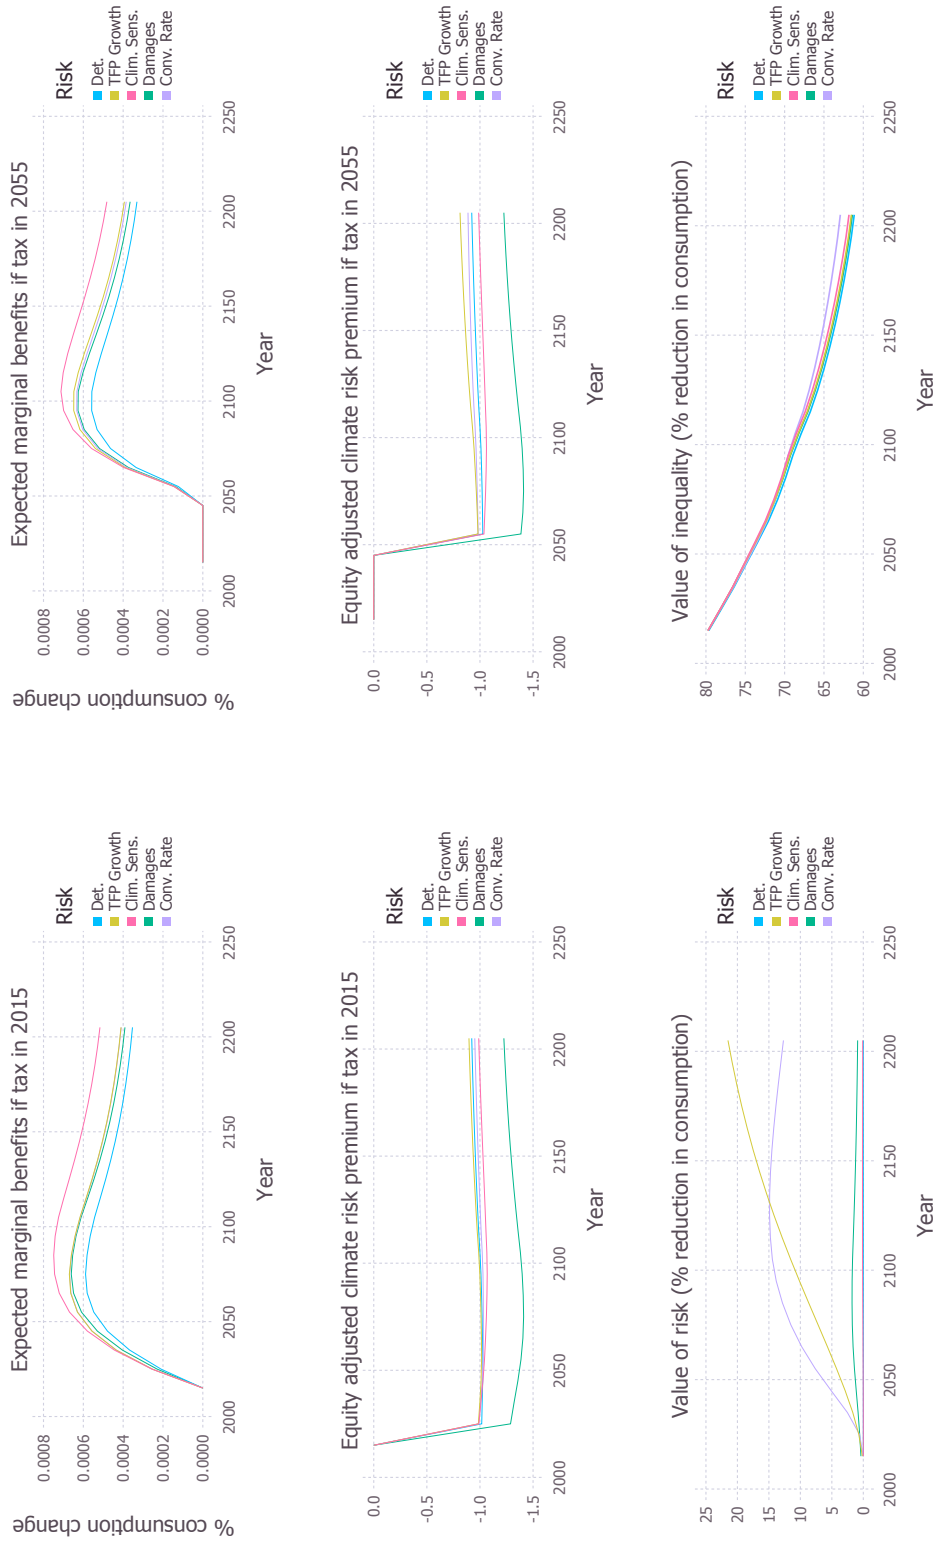


Figure 7: Decomposition equity-adjusted risk premium, NICER $e = 1$




Figure 8: Decomposition equity-adjusted risk premium, NICER $e = 0$

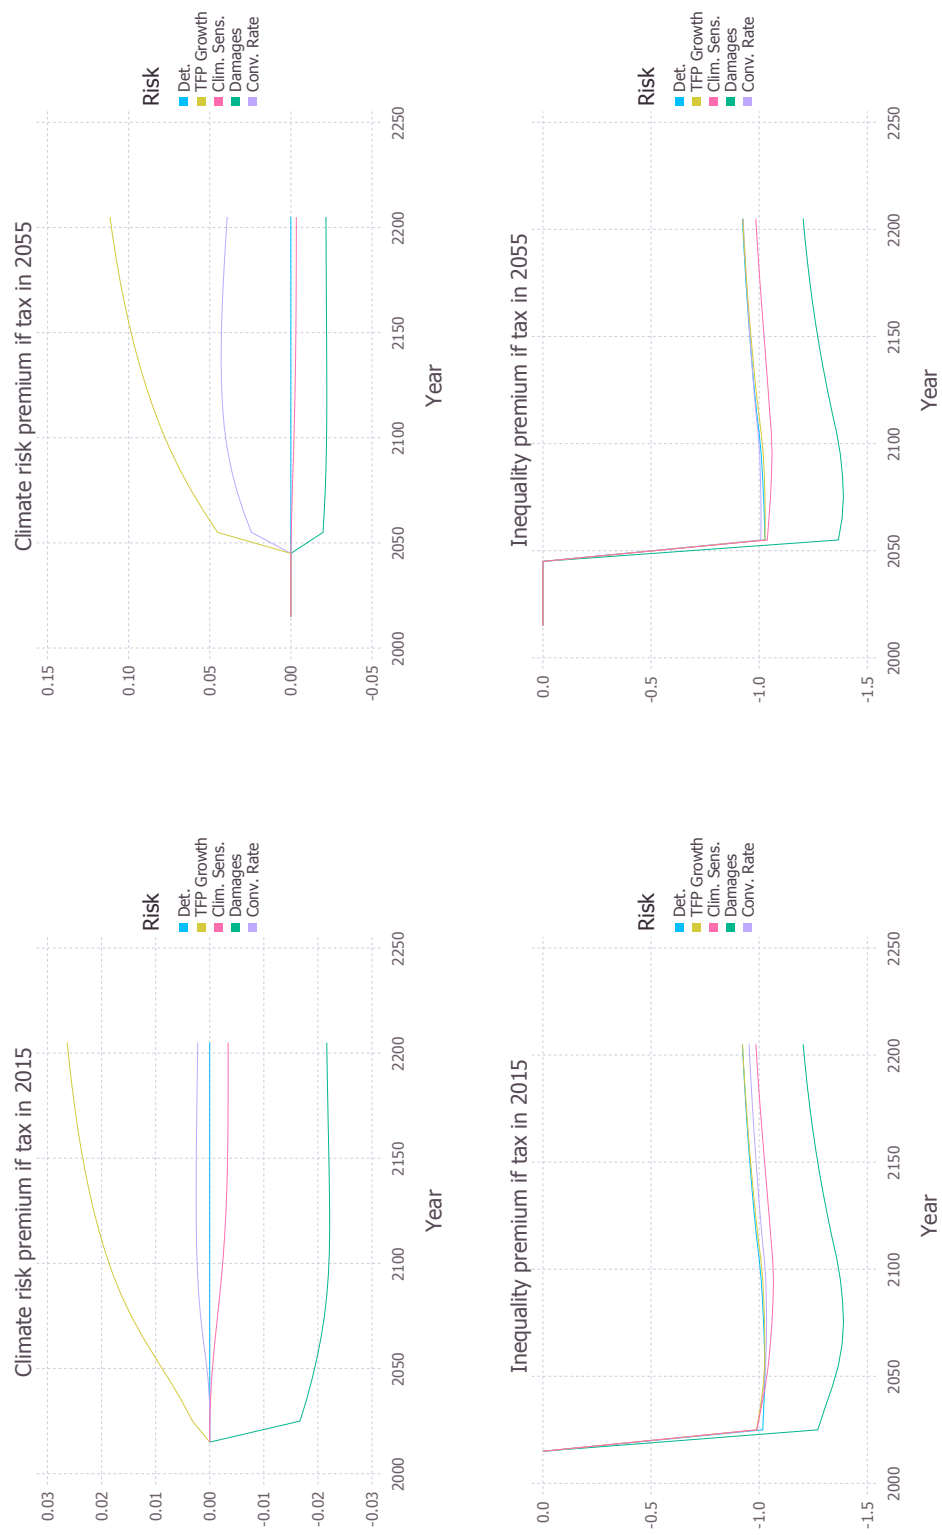


Figure 9: Taxes across models, deterministic




Figure 10: Learning, NICER $e = 1$

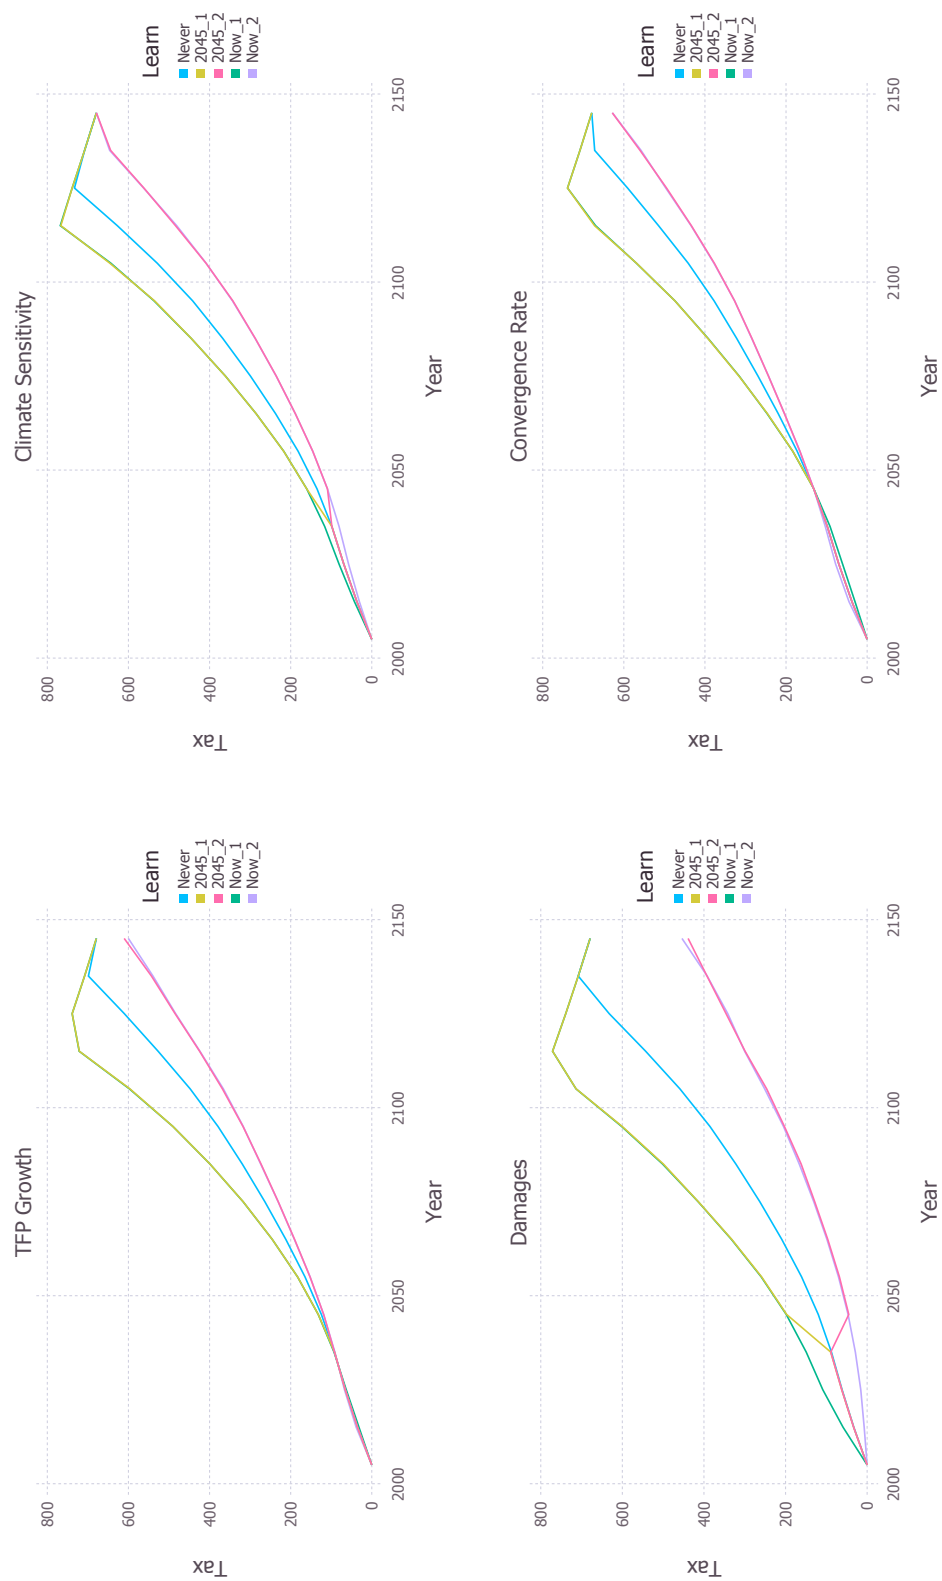


Figure 11: Learning, NICER $e = 0$

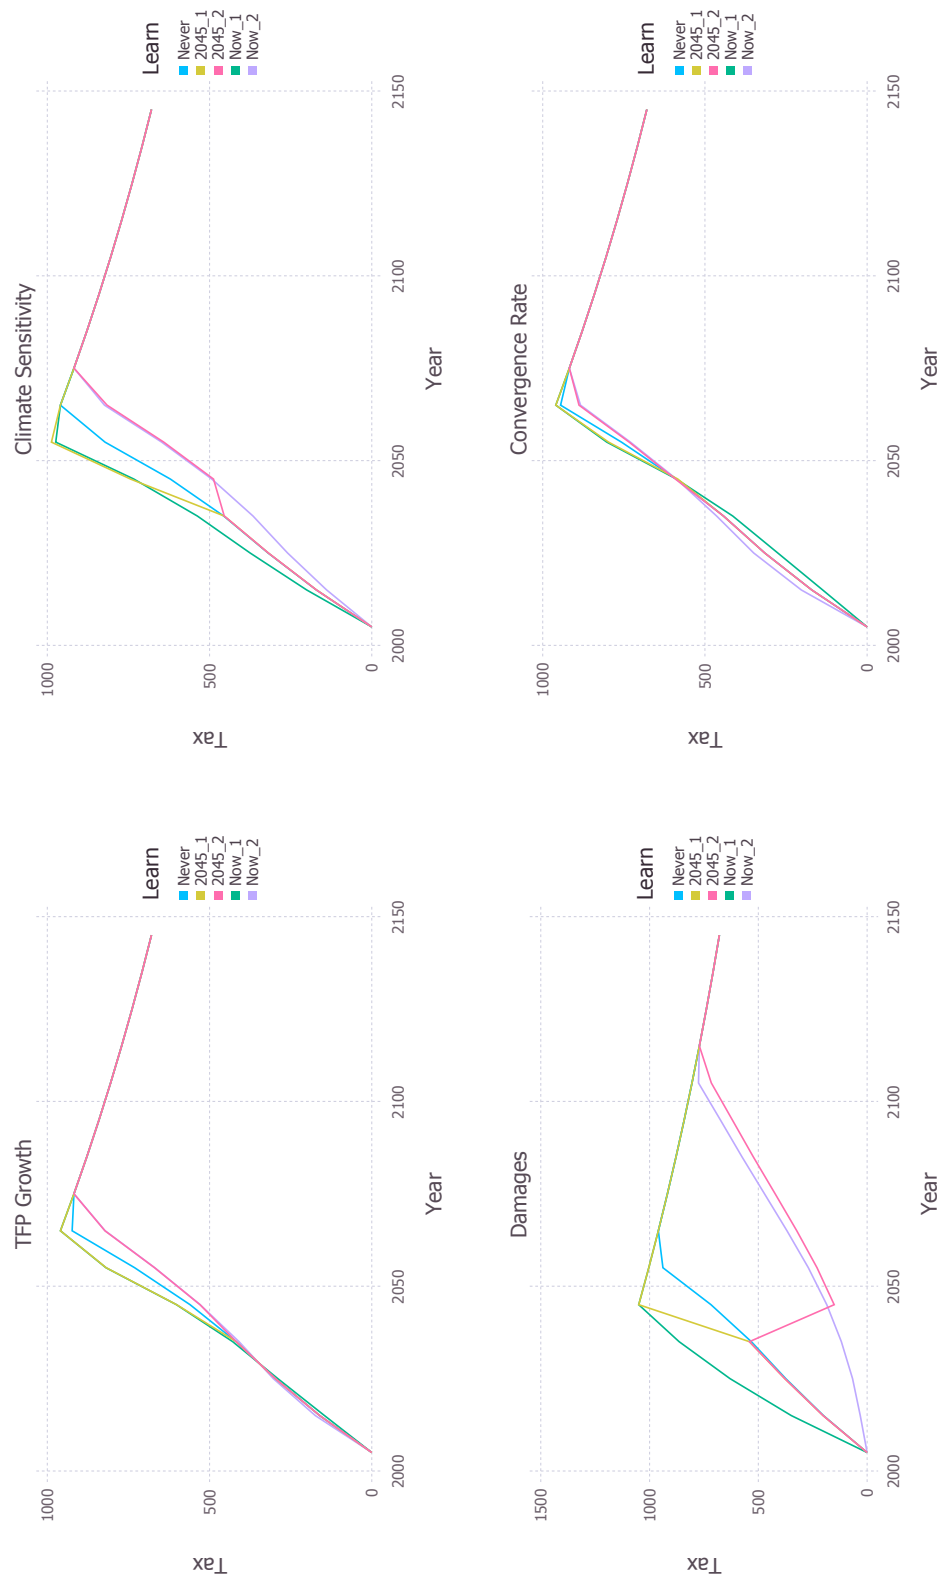


Figure 12: Delay: deterministic, NICE $e = 1$




Figure 13: Delay: deterministic, NICE $e = 0$





Figure 14: Delay: deterministic, NICE $e = 1$ decomposition (%)



The graph shows the Total percentage change in consumption when we delay action (measured in terms of expected equally distributed equivalent) and the disaggregation into its main drivers: change in mitigation costs, change in damages, change in the expected value of inequality within regions, change in expected inequality across regions and change in aggregate consumption risk.

Figure 15: Delay: deterministic, NICE $e = 0$ decomposition (%)




Figure 16: Delay: TFP risk, NICE $e = 0$ decomposition (%)




Figure 17: Delay: Damage risk, NICE $e = 0$ decomposition (%)




Figure 18: Delay: NICE $e = 1$ decomposition, Africa




Figure 19: Delay: NICE $e = 0$ decomposition, Africa




Figure 20: Delay: NICE $e = 1$ decomposition, India




Figure 21: Delay: NICE $e = 0$ decomposition, India




Figure 22: Delay: NICE $e = 1$ decomposition, USA




Figure 23: Delay: NICE $e = 0$ decomposition, USA






Figure 24: Delay: regional WTAs



Regions: 1) Aggregate; 2) USA; 3) OECD Europe; 4) Japan; 5) Russia; 6) Non-Russia Eurasia;
 7) China; 8) India; 9) Middle East; 10) Africa; 11) Latin America; 12) Other High Income;
 13) Other non-OECD Asia

Figure 25: Temperatures under Delay and Learning



Each curve shows the path of temperature under the optimal taxes in cases of delay and learning under TFP growth risk for the realizations in deciles 1, 5, and 10 of the distribution

B The NICER model

As in Nordhaus' RICE model, we assume that the world is populated by 12 regions, each of them characterized by a regional production function $Y_{rt} = F(K_{rt}, L_{rt})$, which depends on regional population L_{rt} and regional capital stock K_{rt} , and by a regional budget constraint $C_{rt} = Q_{rt} - K_{rt+1}$, where C_{rt} denotes total regional consumption, and Q_{rt} regional income net of climate damages and mitigation costs: $Q_{rt} = \frac{1-\Lambda_{rt}}{1+D_{rt}} Y_{rt}$. The damage function is quadratic $D_{rt} = \psi_1 T_t + \psi_2 T_t^2$, and both the mitigation costs, the relationship between production and emissions, and the climate module follow exactly the RICE model.

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. Regions maximize their own inter-temporal utility function U_r :

$$U_r = \sum_t R^t \frac{L_{rt}}{5} \sum_i \frac{\mathbb{E}[c_{irst}^{1-\eta}]}{1-\eta}$$

where $c_{irst} = 5c_{rt}(q_{ir} + D_{rst}(q_{ir} - d_{ir}))$, as explained in the main text. The maximization will result in a stream of optimal mitigation levels $\{\mu_{rst}^*(\tau)\}_{\forall s,t}$, optimal capital levels $\{K_{rst}^*(\tau)\}_{\forall t,s}$, and optimal consumption levels $\{c_{irst}^*\}_{\forall t,s}$, given a path for the carbon tax $\tau = \{\tau_t\}_{\forall t}$. Following again Nordhaus, it turns out that the optimal regional mitigation level μ_{rt} is state-independent¹², which implies that also the mitigation costs Λ_{rt} are state-independent. 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 2017), and almost constant over time. That is true also in our case; therefore, we assume that at the optimum $c_{rsj} \simeq (1-S) \frac{Q_{rsj}}{L_{rj}}$, where S denotes the constant saving rate.

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

$$W = \sum_r U_r = \sum_t R^t \sum_r \frac{L_{rt}}{5} \sum_i \frac{\mathbb{E}[c_{irst}^{1-\eta}]}{1-\eta}$$

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

¹²From the RICE manual, we assume that $\mu_{rt} = \left(\frac{\tau_t}{b_t}\right)^{\frac{1}{\theta-1}}$, where b_t is the backstop price at t and θ a parameter governing the mitigation costs Λ_{rt} .

t , yields the following condition:

$$R^t \sum_r L_{rt} \sum_i \mathbb{E}_s c_{irst}^{-\eta} \frac{-\Lambda'_{rst}}{1 - \Lambda_{rst}} \frac{Q_{rst}}{L_{rt}} \gamma_{irst} + \\ + \sum_{j \geq t} R^j \sum_r L_{rj} \sum_i \mathbb{E}_s c_{irsj}^{-\eta} \left[\frac{-D'_{rsj}}{1 + D_{rsj}} \frac{Q_{rsj}}{L_{rj}} \gamma_{irsj} + \frac{Q_{rsj}}{L_{rj}} D'_{rsj} (q_{ir} - d_{ir}) \right] = 0$$

with $\Lambda'_{rst} = \frac{\partial \Lambda_{rst}}{\partial \tau_t}$, $\gamma_{irst} = q_{ir} + D_{rst}(q_{ir} - d_{ir})$, and $D'_{rsj} = \frac{\partial D_{rsj}}{\partial \tau_t}$. Given the definition of quintile i consumption, $c_{irsj} = 5(1 - S) \frac{Q_{rsj}}{L_{rj}} \gamma_{irsj}$, the previous condition becomes:

$$-R^t L_t \sum_r \frac{L_{rt}}{L_t} \frac{1}{5} \sum_i \mathbb{E}_s c_{irst}^{1-\eta} \lambda_{rt} + \sum_{j \geq t} R^j L_j \sum_r \frac{L_{rj}}{L_j} \frac{1}{5} \sum_i \mathbb{E}_s c_{irsj}^{1-\eta} \delta_{rsj} \left(1 - \frac{(q_{ir} - d_{ir})(1 + D_{rsj})}{\gamma_{irsj}} \right) = 0$$

where $\lambda_{rt} \equiv \frac{\Lambda'_{rst}}{1 - \Lambda_{rst}}$ and $\delta_{rsj} \equiv \frac{-D'_{rsj}}{1 + D_{rsj}}$. By denoting $\tilde{\delta}_{irsj} \equiv \delta_{rsj} \left(1 - \frac{(q_{ir} - d_{ir})(1 + D_{rsj})}{\gamma_{irsj}} \right) = \delta_{rsj} \frac{d_{ir}}{\gamma_{irsj}}$, and after some re-adjustments, we recover expression (2).

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 :

$$W^A = \sum_t R^t L_t \frac{\mathbb{E}_s [c_{st}^{1-\eta}]}{1 - \eta}$$

where $c_{st} = \sum_r \frac{L_{rt}}{L_t} \frac{1}{5} \sum_i c_{irst}$ is the average per-capita consumption in state s . By assuming that the optimal mitigation costs are state-independent and optimal consumption is a fixed amount of output, the maximization with respect to the tax τ_t yields the following expression:

$$R^t \mathbb{E}_s c_{st}^{-\eta} \sum_r L_{rt} \frac{-\Lambda'_{rst}}{1 - \Lambda_{rst}} \frac{Q_{rst}}{L_{rt}} + \sum_{j \geq t} R^j \mathbb{E}_s c_{js}^{-\eta} \frac{-D'_{rst}}{1 + D_{rst}} \frac{Q_{rst}}{L_{rt}} = 0$$

After some re-adjustments, we obtain the following condition:

$$\lambda_t = - \sum_{j \geq t} R^{j-t} \frac{L_j}{L_t} \frac{\hat{\mathbb{E}}_s c_{sj}^{-\eta} \delta_{rsj} c_{rsj}}{\hat{\mathbb{E}}_s c_{st}^{-\eta} \lambda_{rt} c_{rst}}$$

where the variables λ_{rt} and δ_{rsj} are defined as before, and $\lambda_t = \sum_r \frac{L_{rt}}{L_t} \lambda_{rt}$. By defining $\xi_j = \left(\frac{\hat{\mathbb{E}}_s c_{sj}^{-\eta} \delta_{rsj} c_{rst}}{\hat{\mathbb{E}}_s c_{st}^{-\eta} \lambda_{rt} c_{rst}} \right) \left(\frac{\mathbb{E}_s [c_{st}^{1-\eta}]}{\mathbb{E}_s [c_{sj}^{1-\eta} \delta_{sj}]} \right)$, we can recover expression (1).

C Calibrations

C.1 Calibration of Damages

It is useful to recall the damage function that maps temperature increases into damages,

$$D(T) = \psi_1 T + \psi_2 T^2$$

We calibrate the damage function to capture the key features of the data evidenced by Tol (2012), who argues that there is a non-negligible probability that a warming of 2.5°C will lead to negative damages (i.e. higher output) as opposed to positive damages and lower output (see Table 2 and related discussion in his paper). While previous papers (see Dietz et al. 2016 for a review) have attempted to calibrate damage uncertainty via a distribution on the quadratic parameter ψ_2 , we note that for a reasonable calibration based on the Tol data, the damage function becomes non-convex in order to generate negative damages at 2.5°C, which means that there is no incentive to mitigate since higher temperatures are always output-increasing. To solve this problem in a parsimonious manner, we instead focus our calibration on the linear parameter ψ_1 , leaving $\psi_2 > 0$ and thus preserving the convexity of damages.

Let X be a random variable describing total damages as a percentage of global output at 2.5°C of warming. Tol's data can be summarized as saying that $X \sim \mathcal{N}(0.94, 1.28)$. Given the damage function specification and the fact that after-damage output is computed as $Q = \frac{1-\Lambda}{1+D} Y$, % output lost can be expressed as

$$\frac{D(T)}{1 + D(T)} = \frac{\psi_1 T + \psi_2 T^2}{1 + \psi_1 T + \psi_2 T^2}$$

We first extend this distribution into each region as follows. Let X_r be a random variable describing total damages in region r as a percentage of regional output at 2.5°C of warming. Then, $X_r \sim \mathcal{N}(\mu_r, \sigma_r)$ where μ_r is computed using the parameters of the regional damage

function in the risk-less case at 2.5°C, and σ_r satisfies

$$\frac{1.28}{0.94} = \frac{\sigma_r}{\mu_r}$$

so that all regions' damage uncertainties feature the same coefficient of variation.

Therefore, in any region, for a given realization of X_r , x_r , and setting $T = 2.5$, we can write

$$x_r = \frac{\psi_{1r}(2.5) + \psi_{2r}(2.5)^2}{1 + \psi_{1r}(2.5) + \psi_{2r}(2.5)^2}$$

so that holding the ψ_{2r} parameter fixed, ψ_{1r} can be expressed as

$$\psi_{1r} = \frac{\psi_{2r}(2.5)^2(1 - x_r) - x_r}{(2.5)(x_r - 1)}$$

which implies a distribution for the parameter ψ_{1r} as a function of the underlying risky damage variable X_r .

In scenarios featuring damage uncertainty, we simply draw from the distribution of X_r for each region, and then compute what the equivalent draw of the parameter ψ_{1r} must be.

C.2 Calibration of the other random variables

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); Middle-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.

D The approximated cost of delay and its decomposition

Let $\Delta_t k = \frac{EED E_t^d - EED E_t^{nd}}{EED E_t^{nd}}$ be the proportional change in expected equally distributed equivalent consumption when delaying, where EED E is defined as the uniform certain consumption level that is socially equivalent to the initial uncertain, unequal distribution of consumption across the population. Formally, for each period t :

$$L_t u(EED E_t) = \sum_r L_{rt} \frac{1}{5} \sum_i \mathbb{E}_s u(c_{irst})$$

Let us consider a public project that, by changing current consumption by the proportional amount $\alpha_D(k)$ induces a stream of proportional returns $\Delta_t k$, for each time t , and leaves total welfare unchanged. In our case, this project yields mainly costs in the future, as future consumption decreases due to delay. Thus, $\alpha_D(k)$ represents the minimum amount to be paid to the current generation so as to let it accept the project, where k represents the size of the project:

$$L_0 u(EED E_0^{nd}(1 + \alpha_D(k))) + \sum_t R^t L_t u(EED E_t^{nd}(1 + k\Delta_t)) = \sum_t R^t L_t u(EED E_t^{nd})$$

Let us assume that we are interested in a marginal project, i.e. a project whose size k goes to 0 and which does not change the structure of the economy. Then, the cost α_D can be approximated as $\alpha_D(0) \simeq \alpha'_D(0)k$. By fully differentiating the previous condition and evaluating it at $k = 0$, we find that

$$\alpha_D \simeq - \sum_{t=0}^T R^t \frac{L_t}{L_0} \left(\frac{EED E_t^{nd}}{EED E_0^{nd}} \right)^{1-\eta} \Delta_t k$$

By definition, $EDE_t = c_t(1 - \phi_t)$, where $c_t = \sum_r \frac{L_{rt}}{L_t} \mathbb{E}_s \frac{(1-\Lambda_{rt})Y_{rst}}{1+D_{rst}}$ is the average expected consumption at time t , and ϕ_t an index representing the value of risk and inequality (in particular, how much the generation is willing to pay to completely eliminate both risk and inequality). Then, the proportional change in EEDE can be decomposed as:

$$\frac{EEDE^d - EEDE^{nd}}{EEDE^{nd}} = \frac{c^d - c^{nd}}{c^{nd}} + \frac{c^d}{c^{nd}} \frac{\phi^{nd} - \phi^d}{1 - \phi^{nd}}$$

Moreover, given the definition of c_t , the change in average expected consumption can be further decomposed in a change in mitigation costs and a change in damages as follows:

$$\frac{c^d - c^{nd}}{c^{nd}} = \mathbb{E}_s \sum_r \frac{L_r}{L} \frac{c_{rs}^{nd}}{c^{nd}} \left(\frac{(1 - \Lambda_r^d)Y_{rs}^d - (1 - \Lambda_r^{nd})Y_{rs}^{nd}}{(1 - \Lambda_r^{nd})Y_{rs}^{nd}} + \frac{c_{rs}^d}{c_{rs}^{nd}} \frac{D_{rs}^{nd} - D_{rs}^d}{1 + D_{rs}^{nd}} \right)$$

The index ϕ_t will reflect both the degree of regional and sub-regional inequality, and the size of the risk. Let us consider, first of all, the equally distributed consumption inside each region in state s , EDE_{rst} , which is defined as: $u(EDE_{rst}) = \frac{1}{5} \sum_{i=1}^5 u(c_{irst})$. By using a regional inequality premium I_{rst} , the regional equally distributed equivalent in state s can be rewritten as: $EDE_{rst} = c_{rst} - I_{rst}$. The index I_{rst} measures the absolute value of within region inequality in a given state of nature. Given the definition of EDE_{rst} , the aggregate equally distributed equivalent in state s will be such that:

$$\sum_r \frac{L_{rt}}{L_t} u(c_{rst} - I_{rst}) = u(EDE_{st}) = u \left(c_{st} - \sum_r \frac{L_{rt}}{L_t} I_{rst} - \psi_{st} \right)$$

where ψ_{st} denotes the inequality premium that must be paid in state s to eliminate inequality in regional equally distributed equivalent consumptions EDE_{rst} . In other words, ψ_{st} reflects the degree of inequality across regions in a given state of nature. Given the previous definitions, the EEDE is equal to:

$$u(EEDE_t) = \mathbb{E}_s u(EDE_{st}) = u \left(c_t - \sum_r \frac{L_{rt}}{L_t} \mathbb{E}_s I_{rst} - \mathbb{E}_s \psi_{st} - \lambda_t \right)$$

where λ_t denotes the risk premium that must be paid to eliminate the volatility in aggregate equally distributed equivalents. Therefore, λ_t represents the value of aggregate risk. As a consequence:

$$\phi_t = \sum_r \frac{L_{rt}}{L_t} \mathbb{E}_s \frac{I_{rst}}{c_t} + \mathbb{E}_s \psi_{st} c_t + \lambda_t$$

The total premium ϕ_t depends on the value of aggregate risk, on the expected degree of inequality across regions, and on the average expected degree of inequality within regions. Likewise the change in this total index due to delay:

$$\frac{\phi^{nd} - \phi^d}{1 - \phi^{nd}} = \frac{\mathbb{E}_{r,s} I_{rs}^{nd} - \mathbb{E}_{r,s} I_{rs}^d}{1 - \phi^{nd}} + \frac{\mathbb{E}_s \psi_s^{nd} - \mathbb{E}_s \psi_s^d}{1 - \phi^{nd}} + \frac{\lambda^{nd} - \lambda^d}{1 - \phi^{nd}}$$

A similar decomposition can be performed for the regional cost of delay. For the sake of simplicity, let us consider only the deterministic case. In the approximation, the regional cost of delay α_D^r will be equal to:

$$\alpha_D^r \simeq - \sum_t R^t \frac{L_{rt}}{L_{r0}} \left(\frac{EDE_{rt}^{nd}}{EDE_{r0}^{nd}} \right)^{1-\eta} \frac{EDE_{rt}^d - EDE_{rt}^{nd}}{EDE_{rt}^d}$$

where the variation in regional equally distributed equivalent consumption can be decomposed as:

$$\frac{EDE_{rt}^d - EDE_{rt}^{nd}}{EDE_{rt}^{nd}} = \frac{(1 - \Lambda_r^d)Y_r^d - (1 - \Lambda_r^{nd})Y_r^{nd}}{(1 - \Lambda_r^{nd})Y_r^{nd}} + \frac{c_r^d}{c_r^{nd}} \frac{D_r^{nd} - D_r^d}{1 + D_r^{nd}} + \frac{c_r^d}{c_r^{nd}} \frac{\phi_r^{nd} - \phi_r^d}{1 - \phi_r^{nd}}$$

Finally, we can show that the aggregate cost of delay does not coincide exactly with the weighted sum of regional costs of delay. Indeed, given the definitions (8) and (10), a few computations yield:

$$\alpha_D = \left(\sum_{r=1}^{12} \nu_r (1 + \alpha_r)^{1-\eta} \right)^{\frac{1}{1-\eta}} - 1$$

where $\nu_r = \frac{U_{r0}^d}{W_0^d}$, and $U_{r0} = L_{r0} 0.2 \sum_{i=1}^5 \frac{c_{ir0}^{1-\eta}}{1-\eta}$. As a consequence:

$$\alpha_D < \sum_{r=1}^{12} \nu_r \alpha_D^r$$