On Current and Future Carbon Prices in a Risky World

Stan Olijslagers* Frederick van der Ploeg† Sweder van Wijnbergen‡

Revised 31 January 2022

Abstract

We analyse the optimal paths of abatement and carbon prices under a variety of economic, temperature and damage risks. Carbon prices grow in line with economic growth, but with convex damages and temperature-dependent risks of climatic tipping points grow more quickly and with gradual resolution of uncertainty grow more slowly. With temperature-dependent economic damage tipping points carbon prices are higher, but when the tipping point occurs, the price jumps downward. With a temperature cap the efficient carbon price rises at the risk-adjusted interest rate. Allowing for damages as well as a cap leads to a higher carbon price which grows more slowly. But as temperature and cumulative emissions approach their caps, the carbon price is ramped up ever more. Policy makers should expect a rising path of carbon prices.

JEL codes: H23, Q44, Q51, Q54
Keywords: CO2 prices, growth uncertainty, tipping points, damages, gradual resolution of damage uncertainty, temperature caps

*Faculty of Economics and Business, University of Amsterdam.
†Department of Economics, University of Oxford. Also affiliated with University of Amsterdam, CEPR and Tinbergen Institute.
‡Faculty of Economics and Business, University of Amsterdam. Also affiliated with CEPR and Tinbergen Institute.

1 Introduction

Rising temperatures and the ensuing threat to our planet and the economy constitute the biggest market failure we know (Stern et al., 2006). One way to correct this market failure is to price carbon uniformly throughout the global economy at a price equal to the social cost of carbon (SCC) which is the expected present discounted value of all future damages resulting from emitting one ton of carbon today. In this sense, the carbon price or SCC is an asset price just like a share or house price. Another way is to ensure that any production inefficiencies in the production of renewable energies are properly internalized. Production of renewable energies is subject to Swanson's law with any doubling of installed solar panels or windmills leading a reduction of 20 or 40 percent in the unit costs of a panel or windmill. The mechanism underlying Swanson's law is technical progress through learning by doing. This naturally leads to the policy recommendation that renewable energies should receive a subsidy equal to the social benefit of learning or SBL which is the expected present discounted value of all future renewable cost reductions resulting from using one unit of renawable energy today.

Our aim is threefold. First, to analyse the initial level of an optimal carbon price internalizing that externality. Second, to investigate what the entire time path of carbon prices should look like and how various shocks to the economy, damages and temperature, preference structures and parameter choices influence the shape of that time path and its initial starting point. The time path of (expected) carbon prices arguably matters as much as the initial carbon price since much of the adjustment and mitigation efforts will have to take place through new investments and these depend on the trade-off between current costs and future prices.Third, to gain insight into the level and time path of the optimal subsidy for renewable energy and how this interacts with the optimal carbon price.

Pricing carbon reduces demand for carbon-intensive goods, encourages green innovation, carbon capture and sequestration, and locks up fossil fuel in the crust of the earth. This Pigouvian solution charges emissions at a price that implements the optimal policy and in this way internalizes the global warming externality (Pigou, 1920). This price can be implemented as a carbon tax with the revenue rebated in lump-sum manner to the private sector or one could adopt a Coasian approach where property rights to emit or the right to a clean planet are allocated (Coase, 1960), with subsequent trade allowed. There is a burgeoning literature on how high that tax should be today, but some argue that future prices should decline from high initial levels (e.g. Daniel, Litterman, and Wagner (2019)) while most others argue for the exact opposite. Our contribution is to shed light on these differences and offer better understanding of the determinants of the shape of the entire time path of optimal carbon prices.

The learning-by-doing externality should thus be dealt with using a separate instrument: early and direct subsidies of green energy. Given that most integrated assessment models of climate and the economy used for optimal policy analysis are based on the model of Nordhaus (2017) and are concerned with the optimal proportion of energy that is carbon-free, they cannot distinguish between a carbon price and a renewable energy subsidy. This is why the learning-by-doing subsidy and the carbon price are often lumped together with carbon prices as in Daniel et al. (2019)), which then leads to an unwarranted early spike in carbon prices. This may discourage investment in clean technology by raising input costs while not raising future prices commensurately and should thus not be taken literally. We argue that appropriately targeting the second externality by a separate renewable energy subsidy separate from the carbon price unambiguously results in a rising expected time path for the latter.

In climate economics the Pigouvian price is referred to as the social cost of carbon or SCC. Imposing the Pigouvian tax as the private price of carbon actually leads to decentralized implementation of the optimal policy. But the SCC is in fact a more general concept than a Pigouvian tax as it can be evaluated along other paths than the optimal path. For example, the SCC evaluated along a business-as-usual path where global warming externalities are not internalized by private actors in the economy will be higher than along the optimal path if damages are convex enough (Olijslagers, 2021a).

Policy makers must evaluate the SCC under huge uncertainties regarding the wealth of future generations and future global warming damages resulting from emissions today. This involves difficult trade-offs between consumption today and the risks of damages from global warming to consumption and the risks of shocks to temperature in the near and distant future. For that reason we focus extensively on the nature of the stochastic processes driving these uncertainties, on whether we know their distribution or not and on the interaction with the preference structures determining society's current attitude with respect to future uncertainty.

Our benchmark is the case where damages to aggregate production are linear in temperature. Given that recent insights in atmospheric science suggest that temperature is approximately linear in cumulative emissions (Allen et al., 2009; Dietz & Venmans, 2019; Matthews, Gillett, Stott, & Zickfeld, 2009; van der Ploeg, 2018), the function relating the percentage loss in aggregate production to cumulative emissions is then also approximately linear in cumulative emissions.[1] Since damages are proportional to aggregate production, we can show that for this benchmark the optimal carbon price grows at the same rate of growth as the economy. We then consider step by step four generalizations of our benchmark and show how they impact the qualitative pattern of the time path of optimal carbon prices.

First, we show that if damages are a convex function of temperature as has been argued by Weitzman (2012) and Dietz and Stern (2015), the optimal carbon price will start at a higher level than in the benchmark case and will also grow faster than the economy.

Second, we confirm an earlier result by Daniel et al. (2019) that if there is gradual resolution of uncertainty in the damage ratio, there is a component of the optimal

[1]This is related to Golosov, Hassler, Krusell, and Tsyvinski (2014), who take a different perspective but also end up with a linear relation between damages and the concentration of atmospheric carbon. Their damage function is a convex function of temperature and their temperature relationship is a concave function of the stock of atmospheric carbon. They then notice that their exponential damage function is roughly linear in the stock of atmospheric carbon.

carbon price which falls over time.[2] But we show that when there is sufficient growth of the economy, this component is outweighed by the growing component of the carbon price resulting from growing damages. The key insight is thus that gradual resolution of uncertainty slows down the rate of growth of the optimal carbon price but under plausible assumptions does not reverse it. Gerlagh and Liski (2018) also find that learning and resolution of uncertainty slows the rise in the optimal carbon price.[3]

Third, we show that climatic and economic tipping points whose arrival rates increase in temperature boost the carbon price. Once a climate tipping point occurs, it will suddenly increase the sensitivity of temperature to cumulative emissions which in turn should prompt policy makers to boost carbon prices and abatement significantly right now. Immediately after the tip has occurred, climate policy is ramped up resulting in an instantaneous further upward jump in the carbon price and abatement. An economic tipping point that becomes more likely with global warming and abruptly leads to a percentage destruction of production also leads to a higher path of carbon prices and abatement ex ante. But immediately after the tip the optimal level of the carbon price and abatement jump down. Different types of tipping points thus have radically different implications. However immediately after such a downward jump, optimal carbon prices will start rising again.

Fourth, although economists usually take a conventional welfare-maximizing approach, the International Governmental Panel on Climate Change (IPCC) and most countries have adopted the more pragmatic approach of agreeing that policy makers aim to stay below a temperature ceiling. They will do their utmost best to keep global mean temperature well below 2 degrees Celsius and aim for 1.5 degrees Celsius. A temperature cap which bites implies that the optimal carbon price should grow at a rate equal to the risk-adjusted interest rate (cf. Gollier (2020)).[4] Once allowance is made for the risk premium, this Hotelling path for the carbon price is typically faster than the rate of growth of the economy (even when the safe return is below the economic growth rate). Hence, the initial carbon price and abatement will be lower upfront but higher in the future. We find that taking into account risk and uncertainty, climate policy is stepped up hugely as temperature gets closer to its cap. The reason is that policy makers must prevent temperature overshooting the cap. If policy makers adopt a tighter cap, they need to boost the carbon price and abatement. We also show that if policy makers take account of a temperature cap

[2]Daniel et al. (2019) employ the workhorse recursive dynamic asset pricing model consisting of a discrete-time decision tree with a finite horizon extended to allow for Epstein-Zin preferences and generate optimal carbon dioxide price paths based on probabilistic assumptions about climate damages. They argue that it is optimal to have a high price today that is expected to decline over time as the ‘insurance’ value of mitigation declines and technological change makes emission cuts cheaper.

[3]For learning and optimal climate policy, see also Kelly and Kolstad (1999), Kelly and Tan (2015) and van Wijnbergen and Willems (2015).

[4]Gollier (2020) shows in his analysis of the optimal carbon price needed to ensure that a temperature cap is not violated that this rate equals the safe rate plus the beta (the regression coefficient if rate of change in marginal abatement costs is regressed on rate of growth in aggregate consumption) times the aggregate risk premium.

and internalize damages from global warming to aggregate production, the optimal carbon price will grow faster than if only damages are internalized but slower than if only a temperature cap is imposed.

Overall, our results suggest that in face of a wide range of risks and uncertainties and under a wide range of assumptions about society's preferences with respect to various aspects of uncertainty, and in particular whether we know their distribution or not, *policy makers should aim for gradually rising carbon prices.*

Our framework of analysis is a simple endowment economy where the endowment is subject to normal economic shocks (modelled by a geometric Brownian motion) and by rare macroeconomic disasters as in Barro (2006, 2009) and Barro and Jin (2011). Temperature is driven by cumulative emissions, and the fraction of damages lost due to global warming is a power function of temperature and is subject to stochastic shocks with a distribution that is skewed and has mean reversion as in van den Bremer and van der Ploeg (2021). Our short-cut approach to modelling gradual resolution of damage uncertainty is slow release of information. We distinguish aversion to risk from aversion to intertemporal fluctuations, so we use recursive preferences (Duffie & Epstein, 1992; Epstein & Zin, 1989, 1991). This allows for a preference for early resolution of uncertainty when the coefficient of relative risk aversion exceeds the inverse of the elasticity of intertemporal substitution, as empirical evidence strongly suggests Epstein and Zin (1991).

Our paper is closely related to recent contributions by Lemoine (2021) and van den Bremer and van der Ploeg (2021) who also study the effect of damage ratio uncertainty and uncertainty about the economic growth rate in an endowment economy and offer analytical insights into the deterministic, precautionary, damage scaling and growth insurance determinants of the optimal SCC. Our paper differs in that we distinguish relative risk aversion from the inverse of the elasticity of intertemporal substitution and thereby allow for preferences for early resolution of uncertainty. Also we have more general forms of uncertainty and allow for skewness and declining volatility of the shocks to the damage ratio (cf. Daniel et al. (2019)), the risk of rare macroeconomic disasters, and both economic and climatic tipping risk whose frequency increases with temperature. We also allow for learning-by-doing effects in mitigation and thus for the consequent need for renewable energy subsidies. Furthermore, another contribution of our study is that we analyse the effects of temperature caps under uncertainty (both with and without damages to the economy) on the time path of the optimal carbon price under uncertainty. Temperature caps are also analysed in Gollier (2020) but his two period framework precludes the analysis of intertemporal changes in the risk premium which we show to be of importance.

Our paper is also related to an extensive literature on optimal discounting under uncertainty (Gollier, 2002a, 2002b, 2008, 2011, 2012; Olijslagers & van Wijnbergen, 2019; Weitzman, 1998, 2007, 2009, 2011) and optimal climate policy under uncertainty (Crost & Traeger, 2013, 2014; Jensen & Traeger, 2014; Traeger, 2021; van den Bremer & van der Ploeg, 2021). It also relates to a growing literature on optimal climate policy in the presence of climatic and economic tipping points (Cai, Lenton, & Lontzek, 2016; Cai & Lontzek, 2019; Lemoine & Traeger, 2014, 2016; van der Ploeg & de Zeeuw, 2018).

Like much of this literature, we present a simple general equilibrium asset pricing model to answer many of the questions regarding uncertainty and tipping points in this literature. Our focus is different however, in that we explicitly aim to understand the qualitative nature of the *time path* of the path of optimal carbon prices and abatement. We also study temperature caps with time-varying risk premia in a continuous-time, infinite-horizon integrated assessment model of the economy and the climate. In the absence of damages from global warming to the economy, we show that the expected growth in the marginal abatement cost and the price of carbon equals the risk-free rate plus an insurance premium. Compared to Gollier (2020), we additionally consider the implementation of a temperature cap while at the same time internalizing the damages to aggregate production caused by climate change. This gives an expected growth of the carbon price that is in between the growth rate of the economy and the risk-adjusted interest rate.

2 Integrated assessment model for optimal climate policy evaluation under risk

To make the trade-off between sacrifices in current consumption against less consumption due to global warming in the future, we use preferences which recursively defines a value function giving the expected welfare from time t onwards, i.e. V_t (Duffie & Epstein, 1992; Epstein & Zin, 1989, 1991). This formulation distinguishes the coefficient of relative risk aversion γ, from the inverse of the elasticity of intertemporal substitution (we denote the latter by ϵ). Policy makers prefer early (late) resolution of uncertainty if γ exceeds (is less than) $1/\epsilon$. Econometric evidence on financial markets strongly suggests this separation and that the coefficient of relative risk aversion γ exceeds $1/\epsilon$ by a substantial margin (van Binsbergen, Fernández-Villaverde, Koijen, & Rubio-Ramírez, 2012; Vissing-Jørgensen & Attanasio, 2003). Hence, the risk-adjusted interest rate incorporates a so-called 'timing premium' (Epstein, Farhi, & Strzalecki, 2014). If $\gamma = 1/\epsilon$ as is the case with the power utility function, policy makers are indifferent about the timing of the resolution of uncertainty and there is no timing premium in interest rates. Mathematically, this is represented as follows. All agents have identical preferences and endowments, so all the agents can be replaced by one representative agent. If the coefficient of relative risk aversion equals γ and the elasticity of intertemporal substitution equals ϵ, preferences of this agent follow recursively from:

$$V_t = \max_{u_t} E_t\Big[\int_t^\infty f(C_s, V_s)ds\Big] \quad \text{with}$$
$$f(C,V) = \frac{\beta}{1-1/\epsilon}\frac{C^{1-1/\epsilon} - \Big((1-\gamma)V\Big)^{1/\zeta}}{\Big((1-\gamma)V\Big)^{1/\zeta-1}}, \tag{2.1}$$

where $\zeta = (1-\gamma)/(1-1/\epsilon)$ and $\beta > 0$ denotes the utility discount rate or rate of time impatience. If $\gamma = 1/\epsilon$, equation (2.1) boils down to the expected utility approach with no preference for early (or late) resolution of uncertainty.

The endowment of the economy Y_t follows a geometric Brownian motion with drift μ and volatility σ_Y and includes additional terms to allow for disaster shocks with constant mean arrival rate λ_1. The size of the shocks is a random variable with time-invariant distribution. The endowment thus follows the stochastic process:

$$dY_t = \mu Y_t dt + \sigma_Y Y_t dW_t^Y - J_1 Y_t dN_{1,t}, \tag{2.2}$$

where W_t^Y is a standard Brownian motion, $N_{1,t}$ is a Poisson process with arrival rate λ_1, and J_1 is a random variable and corresponds to the share of endowment destroyed if a disaster hits the economy. We assume that $x = 1 - J_1$ follows a power distribution with density $f(x) = \alpha_1 x^{\alpha_1 - 1}$, so $E[x^n] = \alpha_1/(\alpha_1 + n)$ and $0 \leq E[J_1] = 1/(\alpha_1 + 1) \leq 1$. For all moments to exist, we assume that $\gamma < \alpha_1$. This process for the evolution of the economy thus incorporates both normal macroeconomic uncertainty (captured by the geometric Brownian motion) and the risk of rare macroeconomic disasters as in Barro (2006, 2009).

Consumption equals:

$$C_t = \frac{1 - A_t}{1 + D_t} Y_t, \tag{2.3}$$

where A_t denotes the fraction of output used for abatement and D_t is the damage ratio associated with global warming. Note that D_t is always zero or positive by construction, so we also have $0 < C_t \leq Y_t$. The business-as-usual flow of emissions E_t is equal to the product of endowment and the carbon intensity ψ_t:

$$E_t = \psi_t Y_t. \tag{2.4}$$

This carbon intensity ψ_t declines at the following rate: $\delta_t^\psi = \delta_0^\psi e^{-\alpha_\psi t} + \delta_\infty^\psi (1 - e^{-\alpha_\psi t})$. δ_t^ψ is calibrated such that the emissions flow E_t is (in expectation) initially increasing over time, but in the long run the business-as-usual flow of emissions will go towards zero since fossil fuels are exhaustible. Actual emission flows are $(1 - u_t)E_t$ where u_t denotes the abatement rate. Without carbon capture and sequestration (CCS), the upper bound of the abatement rate equals 1 in which case all emissions are fully abated and the economy effectively only uses renewable energy. Carbon capture can be represented by having u_t exceed one which we however rule out in this paper [5].

The cost function for abatement is:

$$A_t = c_0 e^{-c_1 X_t} u_t^{c_2}, \tag{2.5}$$

where X_t is the accumulated stock of knowledge in using renewable energy and c_1 is the parameter that controls how fast the costs decline over time as the stock of knowledge about green technology increases. Furthermore we assume that $c_2 > 1$, so abatement costs are a convex function of the abatement rate. We consider two different processes for the stock of knowledge. In the standard case, the stock of

[5] For a simple reason: at least as of now there do not seem to exist technologies having for carbon capture at a scale to make a material difference

knowledge grows exogenously over time. The abatement cost function then boils down to the cost function of Nordhaus (2017):

$$A_t = c_0 e^{-c_1 t} u_t^{c_2}. \tag{2.6}$$

Technological progress in this case is exogenous. In the alternative, more general case (2.5) we allow for learning by doing by letting the growth of the stock of knowledge be a function of the cumulative amount of emissions that have been abated. Additionally, we consider abatement uncertainty. The stock of knowledge then follows:

$$dX_t = u_t E_t dt + \sigma_X dW_t^X, \tag{2.7}$$

where W_t^X is a standard Brownian motion.

Temperature is a linear function of *cumulative* carbon emissions, so the *change* in temperature depends on the flow of emissions, which gives us for the dynamics of temperature T_t:

$$dT_t = \chi(1 - u_t)E_t dt, \tag{2.8}$$

where χ denotes the transient climate response to cumulative emissions (TCRCE). The damage ratio is a function of temperature and shocks that take some time to have their full impact and follow a skewed distribution to reflect 'tail' risk. The damage ratio is given by:

$$D_t = T_t^{1+\theta_T} \omega_t^{1+\theta_\omega} \quad \text{with } d\omega_t \quad = \upsilon(\overline{\omega} - \omega_t)dt + \sigma_t^\omega dW_t^\omega, \tag{2.9}$$

where ω_t follows a Vasicek (or Ornstein-Uhlenbeck) process with short-run volatility σ_t^ω, mean reversion υ and long-run mean, $\overline{\omega}$ and W_t^ω is a standard Brownian motion (cf. van den Bremer and van der Ploeg (2021)). Here θ_T controls the convexity with respect to temperature and θ_ω controls the skew of the shock hitting the damage ratio. In our benchmark case we assume linear damages in temperature corresponding to $\theta_T = 0$, but we extensively analyse the case of convex damages where $\theta_T > 0$. A novel feature of our analysis is that we allow for a declining time path of volatility, which we capture by the specification:

$$\sigma_t^\omega = \max\left[(1 - t/\overline{t}^\omega)\sigma_0^\omega, 0\right], \tag{2.10}$$

so volatility starts with σ_0^ω and falls linearly to zero after $\overline{t}^\omega$ years. This allows for gradual resolution of damage uncertainty. Volatility is constant if we set $\overline{t}^\omega \rightarrow \infty$. When a temperature cap is implemented, we analyse the consequences of imposing a Paris-like restriction $T_t \leq T^{cap}$. This is in our setup equivalent to the restriction that either only renewable energy must be used once temperature is at its cap, i.e. $u_t = 1$ if $T_t = T^{cap}$ or any emission must be offset by equal carbon capture once that temperature cap is reached.

Finally, we allow for the possibility of *economic* and *climatic* tipping points. We assume that the probability of a tipping point occurs increases in global mean temperature. The hazard rate of the *economic* tipping point equals $\lambda_2 T_t$ where λ_2 indicates the rate at which the hazard rate increases with temperature. We assume that when the system tips, a share J_2 of endowment is destroyed. J_2 is a random variable which also follows a power distribution, but with parameter α_2. The main difference between the *economic* tipping point and the disaster process, is that the tipping point can only tip once, while the Barro-style disasters recur over time. We also consider a *climatic* tipping point for which the sensitivity of temperature with respect to cumulative emissions suddenly increases after a tip. More specifically, we assume that before the tip the transient climate response to cumulative emissions is equal to χ_0 and after the tip it jumps to χ_1. The hazard rate for the *climatic* tipping point equals $\lambda_3 T_t$. We show that the two different specifications have very different implications for the optimal carbon price and abatement rate.

2.1 Implementation of first best in a decentralized economy

We can solve the problem of maximizing expected welfare subject to equations (2.2) to (2.10) using the method of dynamic programming (see Appendix A). The resulting social optimum gives rise to the optimal SCC and can be implemented in a decentralized market economy when the carbon tax is set to the SCC and the revenue is rebated to the private sector as lump sums (see Appendix B). The numerical implementation is discussed in Appendix C.

Define the value function as function of the four state variables and time: $V_t^{i,j} = Z^{i,j}\Big(Y_t, T_t, \omega_t, X_t, t\Big)$, where $i \in \{0,1\}$ and $j \in \{0,1\}$ indicate whether respectively the economic and the climate tipping point has already occurred. The social cost of carbon (SCC) corresponds to the expected present discounted value of all present and future damages to the economy resulting from emitting one ton of carbon today. It equals the welfare loss of emitting one unit of carbon scaled by the instantaneous marginal utility of consumption:

$$SCC_t = -\chi \frac{\partial Z^{i,j}/\partial T_t}{f_C(C_t, V_t)}. \tag{2.11}$$

We consider two cases for abatement costs. In our benchmark case abatement costs decline exogenously over time. In the learning-by-doing case abatement costs are endogenous and increase in the stock of accumulated past abatements (a proxy for accumulated knowledge about the use of green energy). The social benefit of learning corresponds to all the present and future marginal benefits in terms of lower mitigation costs resulting from using one unit of mitigation more today:

$$SBL_t = \frac{\partial Z^{i,j}/\partial X_t}{f_C(C_t, V_t)}. \tag{2.12}$$

In the benchmark case without learning by doing, the SBL is simply equal to zero.

When choosing optimal abatement policy, policy makers must recognize that abatement serves two purposes in our set-up: 1) it reduces emissions and thus global warming, which leads to less climate damages in the future and 2) due to learning by doing, abatement reduces future abatement costs. But abatement is costly. Policy makers must sacrifice current consumption to make room for abatement if they want to curb global warming and increase (expected) future consumption. Optimal abatement u_t thus follows from the condition that the marginal abatement cost (MAC) must equal the social cost of carbon plus the social benefit of learning:

$$MAC_t = SCC_t + SBL_t \quad \text{where} \quad MAC_t = -\frac{\partial C_t/\partial u_t}{E_t}. \tag{2.13}$$

The marginal abatement cost is the cost of abating one more unit of carbon emissions. It increases in the abatement rate u_t since abatement costs are a convex function of the abatement rate. The economy increases abatement until the marginal abatement costs equal the benefits of abatement. If there is no learning by doing, the only benefit of abatement is the reduction of climate change damages. In that case the marginal abatement cost is equal to the SCC, which is the expected present discounted value of all current and future damages caused by emitting one more ton of carbon today. The learning-by-doing externality gives an additional incentive to reduce emissions. The marginal abatement cost thus equals the sum of the social cost of carbon and the social benefit of learning. We denote the optimal abatement policy that solves the dynamic programming problem by u_t^*.

When the government implements a carbon tax which is set it to $\tau_t = SCC_t$ and a renewable energy subsidy which is set to $s_t = SBL_t$, and the net revenue of these policy instruments are rebated as lump sums, the social optimum can be replicated in the decentralized market economy (see Appendix B). Competitive energy producing firms will then choose the energy mix such that the amount of fossil fuel use equals $F_t = (1 - u_t^*)E_t$ and the amount of renewable energy use equals $R_t = u_t^* E_t$, where E_t is the total amount of energy use in the economy (which we have previously referred to as business-as-usual emissions).

We adapt the simple but widely used energy model of Nordhaus (2017) and extend it to allow for uncertainty and tipping points in the economy, the climate sensitivity, and damages from global warming. One drawback of this is that in our setting, taxing carbon is equivalent to subsidizing renewable energy since total energy use is not endogenously chosen by the energy producers and since fossil and green energy are perfect substitutes. Optimal policy could thus in such a framework also be replicated by setting a carbon tax equal to $\tau_t = SCC_t + SBL_t$. However, it is important to stress that this is not the case in more general models. When fossil fuel and renewable energy use can be optimally chosen separately, replication of the command optimum can only be done by setting $\tau_t = SCC_t$ and $s_t = SBL_t$ (e.g. Rezai and van Der Ploeg (2017)). Taxing carbon then has different implications than subsidizing green energy. In a more general setting with directed technical change, it can be shown that when green and dirty inputs are sufficiently substitutable, a temporary green energy subsidy is optimal to fight climate change by kickstarting the economy in directions of green technical progress (e.g. Acemoglu, Aghion, Bursztyn, and Hemous (2012)).[6] Although taxing carbon and subsidizing green energy are equivalent in our

[6] Bovenberg and Smulders (1995, 1996) offer early contributions on climate policy and endogenous growth. It has also been argued that subsidizing green energy technology is not effective to fight climate change, since it leads to higher energy use in total instead of substantially less fossil fuel use (Hassler, Krusell, Olovsson, & Reiter, 2020).

simple framework, we do interpret the social cost of carbon as the optimal carbon tax and the social benefit of learning as the optimal renewable energy subsidy, to stress that the two are in general not equivalent.

We assume that negative emissions are not possible (or at least not at a competitive price) and thus impose an upper bound on the abatement rate of 1. Hence, when it would be optimal to abate more than 100% of the emissions, the optimality condition (2.13) cannot be satisfied anymore. In this case the marginal abatement costs are smaller than the sum of the social cost of carbon and social benefit of learning.

2.2 Effects of a temperature cap on carbon pricing

Optimal policy in presence of a temperature cap still satisfies the first-order condition, but the SCC now must account for the temperature cap. A temperature cap in our model is equivalent to the restriction that $u_t = 1$ when $T_t = T^{cap}$ since only then the cap will not be violated. We show that in the case of a pure temperature cap (i.e. no effect of climate change on damages to aggregate production), intertemporal optimization implies that the expected growth rate of the SCC and of the MAC must equal the risk-free interest rate plus a risk premium (for a proof, see Appendix D). Let π_t be the stochastic discount factor. In this case, we thus have that expected growth in the MAC and in SCC (in the absence of learning by doing) equals:

$$E_t\Big[\frac{dMAC_t}{MAC_t}\Big] = r_t + E_t\Big[-\frac{d[\pi_t, MAC_t]}{\pi_t MAC_t}\Big], \tag{2.14}$$

where $d[\pi_t, MAC_t]$ is the quadratic variation between the stochastic discount factor and the marginal abatement cost. In expectation, the growth rate of marginal abatement costs is therefore equal to the risk-free rate r_t plus a risk premium related to the correlation between π_t and MAC_t. If changes in consumption and marginal abatement costs are positively correlated, then the stochastic discount factor and marginal abatement costs are negatively correlated and the risk premium is positive. This leads to a faster rate of growth of the SCC and thus the optimal carbon price.[7]

This result echoes the result derived by Gollier (2020) for a two-period model. It follows from the assumption that temperature increases linearly in cumulative emissions. In that case, we get an equivalent of the Hotelling rule with the carbon price growing at a rate equal to the risk-adjusted interested rate.[8] This price path achieves intertemporal efficiency and ensures that temperature does not exceed the cap. In other words, the risk-adjusted discounted MAC is the same for each period.

Marginal abatement costs (without learning by doing) are given by:

$$MAC_t = -\frac{\partial C_t/\partial u_t}{E_t} = \frac{Y_t}{E_t}\frac{\partial A_t}{\partial u_t} = \frac{Y_t}{E_t} c_0 e^{-c_1 t} c_2 u_t^{c_2 - 1} = \frac{c_0 e^{-c_1 t} c_2 u_t^{c_2-1}}{\psi_t}. \tag{2.15}$$

[7] We discuss the risk premium in more detail in the numerical part of the paper in section 4.6.

[8] Normally, the Hotelling rule states that the growth rate of the efficient carbon price equals the (risk-adjusted) interest rate plus the rate of decay of atmospheric carbon. Since in our model temperature is driven by cumulative emissions rather than the stock of atmospheric carbon, the decay rate does not appear in our Hotelling rule.

The assumption that both total abatement costs and emissions are proportional to endowment implies that abatement costs per unit of emissions do not scale with endowment. Uncertainty within marginal abatement costs is driven by uncertainty around the optimal abatement rate u_t^*. We show that the risk premium has a clear negative time trend, which is not visible in a 2-period model like Gollier (2020). We will analyse the risk premium in more detail in the results section.

3 Calibration and benchmark results

We discuss our benchmark calibration and then present and discuss the corresponding optimal time path for respectively the carbon price, the learning-by-doing subsidy, abatement, and temperature.

3.1 Calibration

In our benchmark calibration, we choose for the coefficient of relative risk aversion $\gamma = 7$, for the elasticity of intertemporal substitution $\epsilon = 1.5$ and for the rate of impatience $\beta = 2\%$ per year. These are values that are typically used in the asset pricing literature with Epstein-Zin preferences (e.g. Table 1, Cai and Lontzek (2019)) based on extensive empirical evidence. The details of our calibration are reported in Table 1.

The initial endowment is set to world consumption (using purchasing power parities) of 80 trillion 2015 US dollars. We suppose this endowment is subject to normal shocks captured by geometric Brownian motion with a drift of 2% per year and an annual volatility of 3%. In addition, we have macroeconomic disaster shocks along the lines of Barro (2006, 2009). Here the mean size of the disaster shocks is 8.7% and the mean arrival rate of these shocks is 0.035 per year corresponding to a mean arrival time of 29 years. This calibration yields a real risk-free interest rate of 0.75% and a risk premium of 2.65% if we abstract from the adverse effects of climate change on the economy. Since in the past century climate change has arguably had no effect on interest rates, we can compare these numbers to historical averages.

Dimson, Marsh, and Staunton (2011) calculate that the global real risk-free rate has been on average 1% and the risk premium 4% over the period 1900-2010. We are currently in a low interest environment and in the long run it is questionable whether interest rate will return to their old average levels, which makes 0.75% a reasonable real risk-free interest rate. Our risk premium is lower than the historical average, but our main purpose is not to solve the equity premium puzzle. Furthermore, a risk premium of 2.65% is more realistic compared to most other climate-economy models in which the risk premium is often small or non-existent.[9] These numbers are also close to Gollier (2020) who calibrates the risk-free rate at 1% and the risk premium at 2.5%.

[9] Our model does not distinguish between asset price volatility and output growth volatility. We have calibrated the volatility in our model to output growth volatility. Calibrating the volatility to asset price volatility would lead to a higher risk premium.

Table 1: *Calibration details*

Preferences	Coefficient of relative risk aversion: $\gamma = 7$ Elasticity of intertemporal substitution: $\epsilon = 1.5$ Rate of impatience: $\beta = 2\%$
Economy	Initial endowment: $Y_0 = 80$ trillion US dollars *Geometric Brownian motion* Drift in endowment $\mu = 2\%/year$ Volatiltiy of shocks to endowment $\sigma_Y = 3\%/\sqrt{year}$ *Macroeconomic disasters* Arrival rate of disasters $\lambda_1 = 0.035/year$ Mean size of disasters $E[J_1] = 8.7\%$ Shape parameter of power distribution $\alpha_1 = 10.5$
BAU emissions	Initial flow of global emissions in BAU scenario $E_0 = 10GtC/year$ Initial growth rate of carbon intensity $\delta_0^{\psi} = -0.5\%/year$ Long run growth rate of carbon intensity $\delta_{\infty}^{\psi} = -6.5\%/year$ Carbon intensity parameter $\alpha_{\psi} = 0.0025$
Abatement costs (Benchmark case)	Initial cost of full decarbonization $c_0 = 7.41\%$ of GDP Rate of technological progress $c_1 = 1.9\%/year$ Convexity parameter cost function $c_2 = 2.6$ Maximum abatement $u \leq 1$
Abatement costs (Learning by doing case)	Initial level of knowledge stock $X_0 = 0$ Initial cost of full decarbonization $c_0 = 7.41\%$ of GDP Rate of technological progress $c_1 = 0.375\%$/unit of knowledge Convexity parameter cost function $c_2 = 2.6$ Abatement cost volatility $\sigma_X = 5$ Maximum abatement $u \leq 1$
Temperature	Initial temperature $T_0 = 1^oC$ Transient climate response to cumulative emissions $TCRCE = \chi_0 = 1.8^oC/TtC$ Temperature cap $T^{cap} = 2^oC$ or $T^{cap} = \infty$
Damage ratio	Convexity parameter $\theta_T = 0$ (linear) or $\theta_T = 0.56$ (convex) Skew parameter of shocks $\theta_{\omega} = 2.7$ Mean reversion of shocks $\upsilon = 0.2/year$ Initial and mean steady-state value of shocks $\omega_0 = \overline{\omega} = 0.21$ Constant volatility variant $\sigma_0^{\omega} = 0.05$, $\overline{t}^{\omega} = \infty$ Declining volatility variant $\sigma_0^{\omega} = 0.05$, $\overline{t}^{\omega} = 100$ years
Economic tipping point	Arrival rate of tipping point $\lambda_2 = 0.01T_t$ Mean tipping damage level $E[J_2] = 2.5\%$ Shape parameter of power distribution $\alpha_2 = 39$
Climatic tipping point	Arrival rate of tipping point $\lambda_3 = 0.006T_t$ $TCRCE$ after tipping $\chi_1 = 2.5^oC/TtC$

Figure 1: *Expected costs of full abatement ($u_t = 1$) in the benchmark and in the learning-by-doing case for two different damage specifications (linear and convex)*

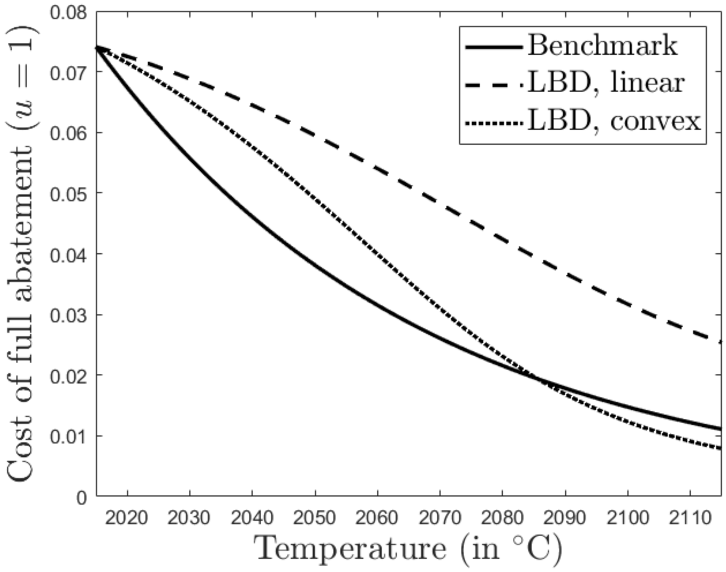

Parameters for the carbon intensity are chosen to match in expectation the baseline emissions scenario in Nordhaus (2017). The parameters c_0, c_1 and c_2 of the abatement cost function in the benchmark case are taken from the DICE calibration (Nordhaus, 2017). For the learning by doing calibration, we take the same value for c_0 (cost of full abatement in initial period) and for c_2 (convexity of abatement costs in abatement rate u_t). The parameter c_1 now represents the decline in abatement costs when one additional Gt of carbon emissions is abated and is set to $c_1 = 0.375\%$ (cf. Rezai and van Der Ploeg (2017)). With learning by doing in renewable energy production, future abatement costs depend on cumulative past abatement efforts and thus also depend on the damage calibration. Figure 1 compares abatement costs of the benchmark case with the learning-by-doing case, both when damages from global warming are linear and when they are convex in temperature.

We set the transient climate response to cumulative emissions (TCRCE) to $1.8^oC/TtC$ (Matthews et al., 2009). The parameters of the uncertain damage shock and of the convexity parameter θ_T are taken from van den Bremer and van der Ploeg (2021). For the variant with gradual resolution of damage uncertainty, we assume that the volatility of the damage shock is linearly declining to 0 over a period of 100 years as in equation (2.10).

Finally, we assume that initially an economic tipping point is expected to tip once every 100 years (i.e. an arrival rate of 0.01) but that the arrival rate increases linearly with the temperature. Thus when temperature increases to two (four) degrees Celsius, this becomes a tip once every 50 (25) years. The size of the damages caused by the tipping disaster is assumed to be on average 2.5%. For the climate tipping point, it takes initially on average 167 years for the climate system to tip. With 2 degrees Celsius warming the average time reduces to 83 years. When the system tips,

Figure 2: *Benchmark with linear damages, no learning by doing, no gradual resolution of uncertainty, no tipping points, and no temperature cap*

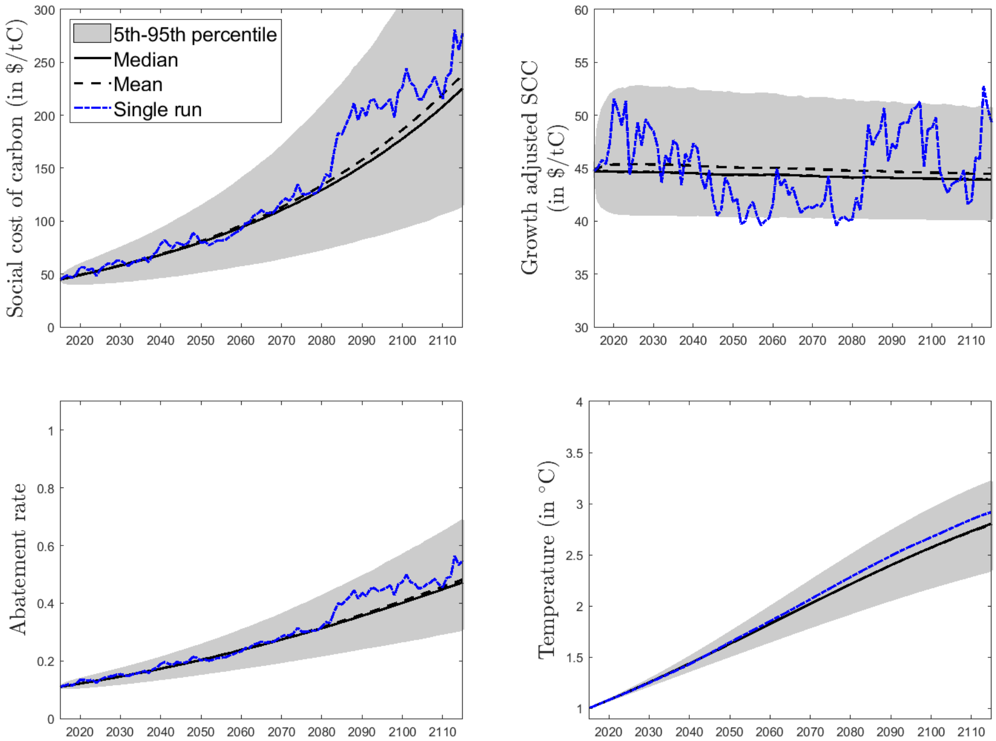

the TCRCE jumps from $1.8^oC/TtC$ to $2.5^oC/TtC$. Overall, the main message of the two tipping point calibrations is that the probability of tipping in both cases is quite small initally but goes up as temperatures rise, and we will show that the impact on optimal carbon prices and the abatement rate is considerable.

3.2 Benchmark optimal carbon prices

With this calibration, the benchmark SCC (with linear damages, no learning by doing and no temperature cap) is shown in Figure 2. The SCC corresponds to the optimal carbon price. The most striking feature of the top two panels is that the ex-ante *mean* and *median* paths of the optimal carbon price start at almost 45\$/tC and then grow almost in tandem with the growth of the economy.

In fact, there is a modest decline in carbon price corrected for the growth of the economy as can be seen from the top right panel. The median carbon price path lies below the mean carbon price path, and the 5% and 95% bounds become wider for carbon prices that are further in the future as one should expect given that a function of GBM processes is a GBM process itself (cf Shreve (2004)). As a result of the technological progress in abatement technology, there is a gradual rise in abatement efforts over time. Due to the rise in business-as-usual emissions, temperature rises in expectation to almost 3 degrees Celsius in the next century but by rather less than in the absence of abatement efforts. The plots also indicate a sample run in blue. This suggests that for individual sample paths of the optimal carbon price there may be substantial volatility, which does not show up in the ex-ante time path for the mean (or median) optimal carbon price. There is also substantial uncertainty around optimal abatement policy and future temperature levels.

4 Five generalizations of the benchmark

We now discuss five generalizations of the benchmark. For expositional reasons, we discuss these generalizations one by one. In practice, all these generalizations are relevant at the same time. We discuss first the effects of convex damages, then present the effects of learning by doing and a combination of convex damages and learning by doing. After that we discuss the implications of gradual resolution of damage uncertainty and then show the differential impacts of climatic and economic tipping points. We then analyse the effects of temperature caps on the time path of carbon prices both with and without damages to economic production. Finally, we present a combined optimal policy simulation exercise as an alternative for the benchmark case.

4.1 Convex damages

Figure 3 presents the effects of convex damages captured by the proportion of output lost due to global warming being a convex rather than a linear function of temperature. Following van den Bremer and van der Ploeg (2021), we let this function be proportional to temperature to the power of 1.56. This is slightly less convex than the damage function of Nordhaus (2017) but serves to illustrate the effects of convex damages. The most striking effect of convex damages is that the carbon price starts at a higher level, 87$/tC instead of 44$/tC, and then grows in expectation at a faster pace than in the benchmark. We can see this most strikingly by comparing the top right panel of Figure 2 with the left panel of Figure 3. This shows that with convex damages, the path of optimal carbon prices corrected for growth of the economy rises whilst with linear damages, this path declined mildly. Hence, the abatement efforts are much stronger. The average mitigation rate now rises in a century to 83% instead of 47% in the benchmark. We thus confirm the earlier Monte-Carlo results of Dietz and Stern (2015) in our fully stochastic framework: climate policies get intensified if damages are convex.[10]

4.2 Learning by doing in abating emissions

Including learning by doing into the analysis gives an additional reason for abatement. The MAC is now equal to the SCC plus the SBL. The SCC adjusted for economic growth is very similar to the base situation, so changing the abatement cost structure

[10]Crost and Traeger (2013) show that Monte-Carlo simulations do not properly take account of uncertainty on the optimal carbon price and can lead to misleading results.

Figure 3: *Convex damages*

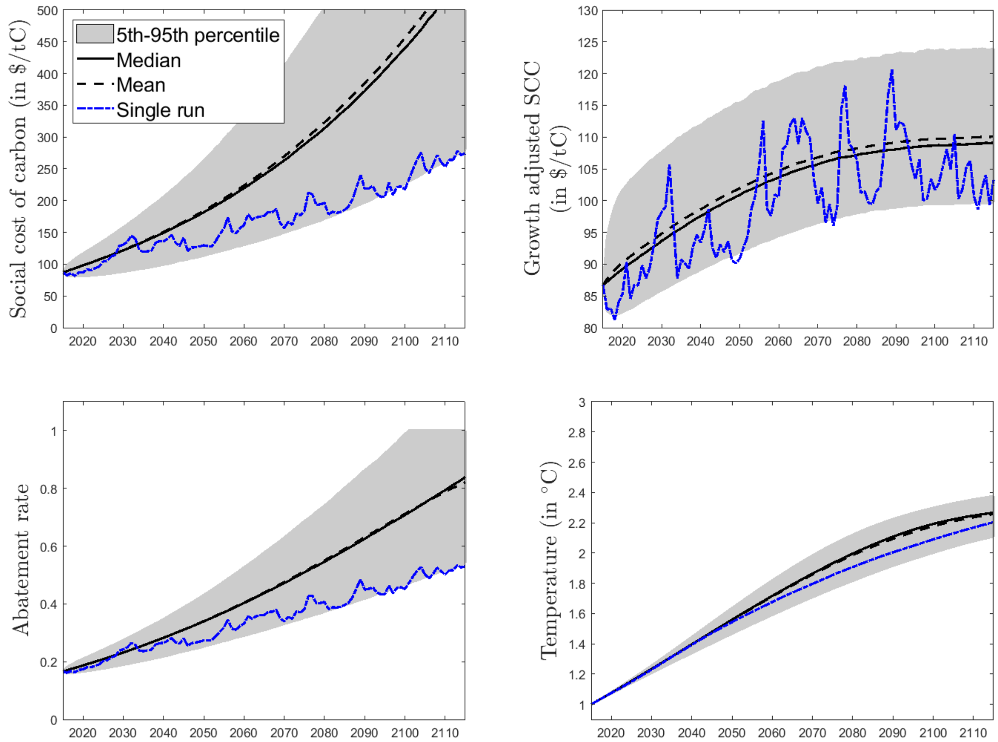

has no significant effect on optimal carbon prices. The only difference is that growth adjusted carbon prices have a slight downward trend over time (i.e. grow slightly slower than GDP). But optimal carbon prices still grow almost in tandem with the economy (see top left panel of Figure 4).

The SBL shown in the top right panel of Figure 4 has a very different shape. It grows faster than the economy in the first 50 years: the panel displays the growth-adjusted SBL, hence an upward-sloping time path of this SBL implies that the SBL grows at a higher rate than the economy. But in the second half of this century (after 2050) abatement costs have been reduced substantially because of learning by doing to such an extent that even lower abatement costs do not give much additional benefit anymore, the SBL falls below its 2020 value. Another notable difference is due to the stochastic nature of the learning-by-doing externality: there now is much more uncertainty in the optimal abatement rate. This is reflected in much wider confidence intervals for the abatement rate u_t^*.

4.3 Convex damages and learning by doing in abatement

Figure 5 shows that combining convex damages and learning by doing leads to a much stronger incentive for abating energy. The optimal carbon price is again quite similar to the optimal carbon price without learning by doing. The SBL now has a clearer hump shape compared the linear case. It starts much higher and declines towards

Figure 4: *Learning by doing in abating emissions*

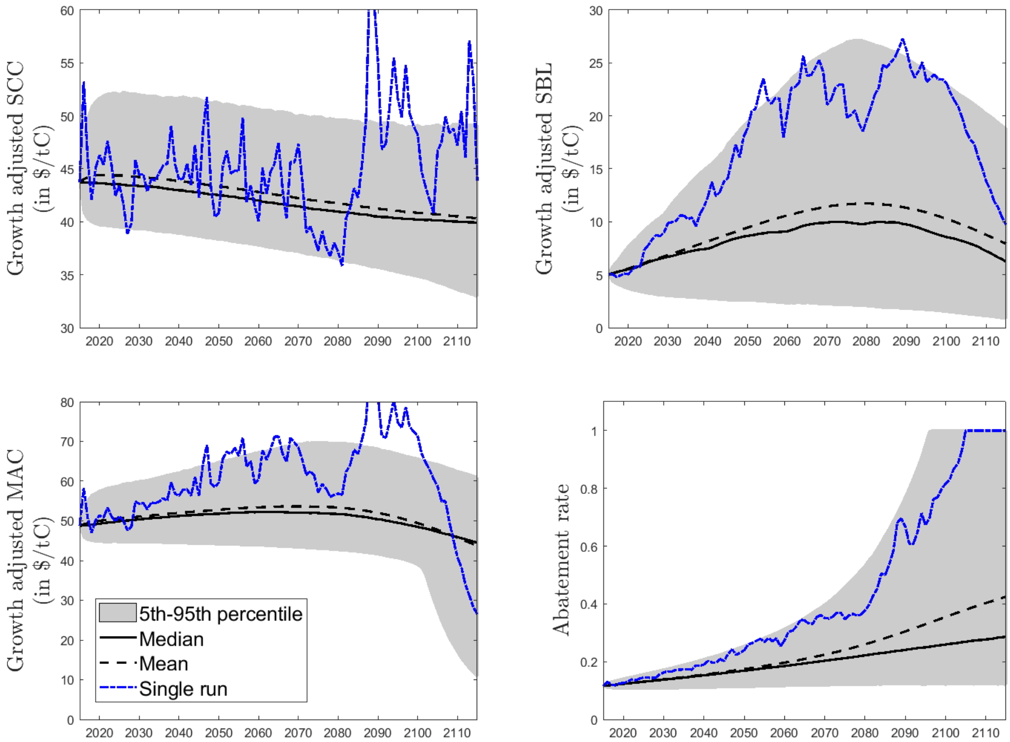

zero faster. Since damages are more severe, more abatement is optimal and lowering abatement costs by investing in knowledge is even more beneficial, which explains the higher level of the SBL. In this scenario it is optimal to fully decarbonize the economy around the end of the century in the median scenario, much earlier than in the previous scenarios. The main takeaway from the learning-by-doing simulations is that optimal abatement of emissions is understated if learning-by-doing externalities are not properly internalized.

4.4 Gradual resolution of damage uncertainty

Our third generalization is to allow for gradual resolution of damage uncertainty. More precisely, we let the annual volatility of the damage ratio fall to zero linearly in a century. This is a shortcut to capturing slow resolution of uncertainty without delving into the intricacies of learning [11]. The left panel of Figure 6 indicates that the expected optimal path of carbon prices now grows more slowly than the economy itself: corrected for growth of the economy the SCC now falls over time, much more strongly than the modest decline shown in the benchmark (see top right panel of Figure 2). We find that the optimal carbon price does not only grow much more slowly than the economy, but also starts at only 33$/tC instead of 44$/tC. The fact

[11]See Gerlagh and Liski (2018) for similar results in a formal model of learning and uncertainty resolution over time.

Figure 5: *Convex damages and learning by doing in abating emissions*

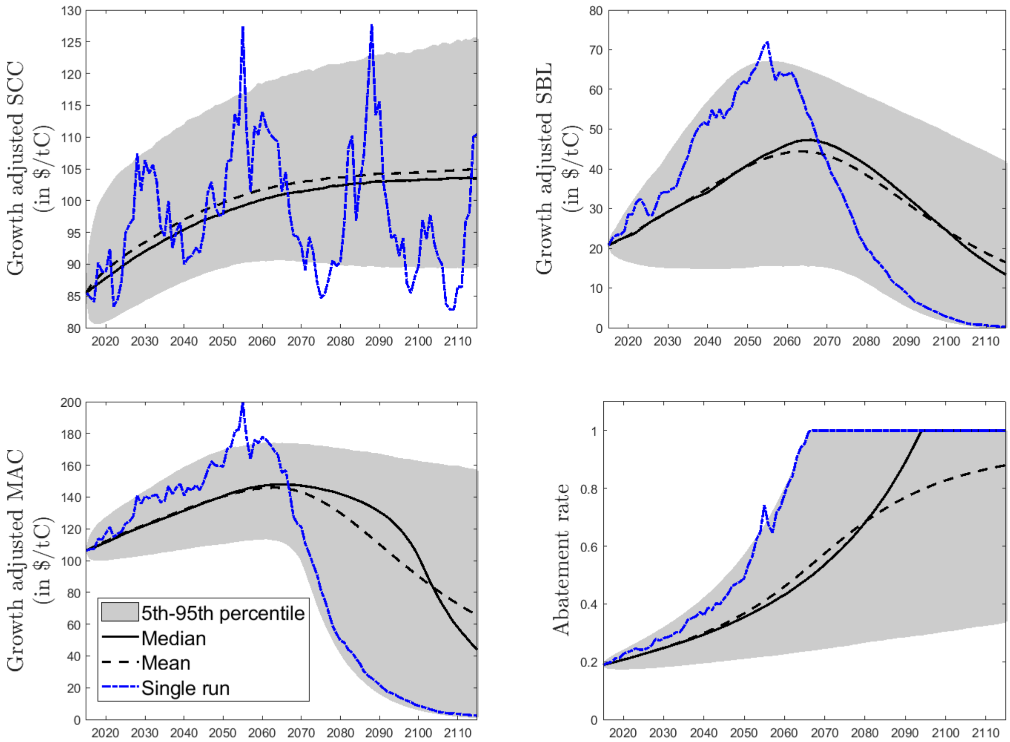

that there is declining uncertainty about the damage ratio means that policy makers can pursue a less vigorous climate policy than in the benchmark. Declining volatility in the future already has an impact on the optimal carbon price today. This implies that the mitigation rate rises in a century to only 34% compared to 47% in the benchmark. Note that if there is no or very little growth in the economy, the optimal carbon price would decline over time as found by Daniel et al. (2019) for a 7-period model for integrated assessment of economy and the climate. The general point is that gradual resolution of damage uncertainty slows down the rate of growth of the optimal carbon price and actually shifts the entire time path down also. The initial carbon price is also lower than in the benchmark case.

4.5 Climatic and economic tipping points

Our fourth generalization is to allow for climatic and economic tipping points. There is a growing literature on the effects of various stochastic tipping points on optimal climate policy (Cai & Lontzek, 2019; Lemoine & Traeger, 2014, 2016; van der Ploeg & de Zeeuw, 2018). Most of these studies are quite challenging from a numerical point of view. Here we simply present the effects (relative to our benchmark) of two illustrative tipping points.

We first present a single climatic tipping point for which we assume that there is a risk of a regime shift in which the transient climate response to cumulative emissions

Figure 6: *Gradual resolution of damage uncertainty.*

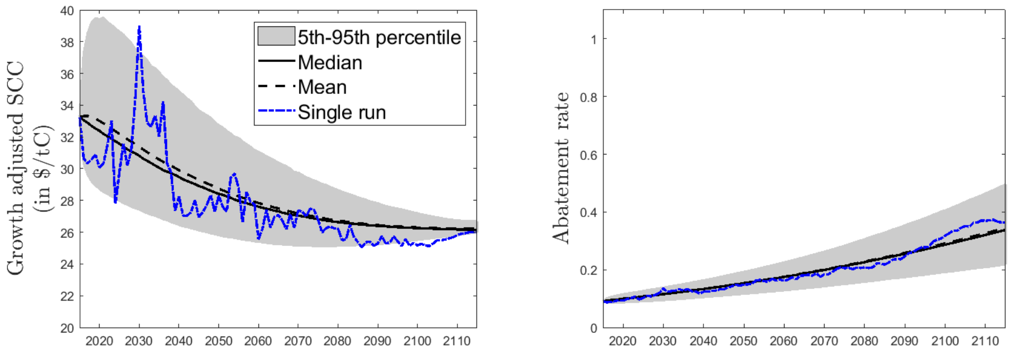

Figure 7: *Risk of a climate tipping point*

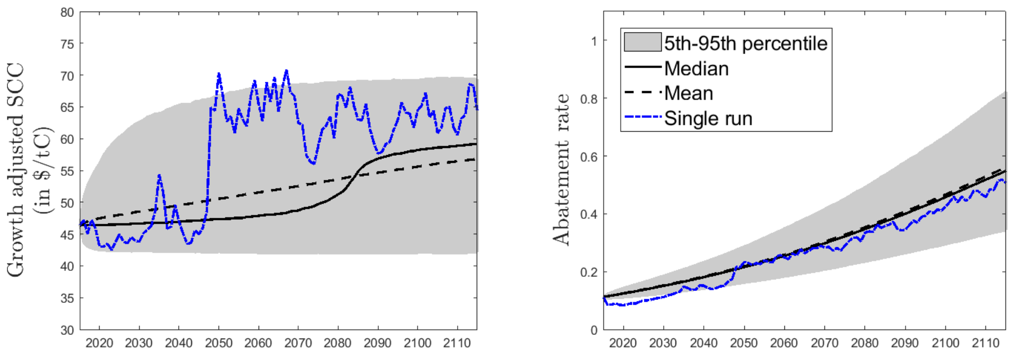

at an unknown future moment suddenly jumps up from $1.8^oC/TtC$ to $2.5^oC/TtC$. Moreover, we assume that the arrival rate is higher at higher temperatures: the initial hazard of this tip at the initial temperature of 1 ^oC is 0.006, which implies an expected arrival time of 167 years, but for every increase in temperature by 1 ^oC we let the hazard rate rise by a further 0.006. This means that at 3 ^oC the hazard is 0.018 and the mean arrival time for the catastrophe is only 56 years. Global warming makes the tipping point thus more imminent. Although these small risks are likely to occur in the very distant future, they have consequences on optimal climate policy now already, as can be seen by comparing Figure 7 with Figure 2. The mean optimal carbon price now starts somewhat higher at 48$/tC than in the benchmark and then rises over time. Hence, the mitigation rate ends up higher after a century, at 55% instead of 47%. The blue lines indicate a sample path with the tipping point occurring in 2047. At that time, the carbon price jumps up substantially because of the bigger climate challenge resulting from the increased sensitivity of temperature to cumulative emissions.

Figure 8 shows the optimal policy simulations for a different type of tipping point, namely one that leads to a higher effect of global warming on damages instead of increased temperature sensitivity. We assume that the size of the economy drops on average by 2.5% once this tipping point occurs. The initial hazard of this tip at initial temperature is 0.01, which implies an expected arrival time of 100 years. For each increase in temperature by 1 ^oC, the hazard rate is assumed to rise by 0.01. Hence, at 3 ^oC the hazard is 0.03 and the mean arrival time for the tipping point goes down to 33 years. This economic tipping point is thus expected to occur more rapidly than the climate tipping point. The most striking feature is that for this tipping point, the initial carbon price is much higher than in the benchmark, i.e., 85\$/tC instead of 44\$/tC, but that the mean and median paths of the optimal carbon price corrected for growth of the economy fall strongly over time. The blue line indicates a sample run where the tipping point occurs at the end of the century. At that time, the carbon price drops down instantaneously and, as a result, the mitigation rate drops down at that time too. The intuition behind this drop is obvious: initially, a large fraction of the carbon price is reflecting the urgency of preventing the tipping point. A higher carbon price leads to more mitigation efforts and therefore a lower probability of the tipping. But when despite these additional abatement efforts, the system tips eventually, there are no further tipping points to prevent. Moreover, after the tip has occurred the economy is smaller because of the sudden increase in damages. The social costs of carbon are proportional to output, which is another factor behind the drop in the SCC after the damage catastrophe occurs. Another way of saying this is that the stochastic discount factor is higher in this case, which implies both a higher initial SCC and a slower growth rate. The benefit of carbon reduction after the tip is the same as the benefit in the benchmark model without the tipping point for the same level of output.

This is an important point: a tipping point in the climate system that speeds up warming or leads to a slower decay of carbon emissions has very different implications than a tipping point that directly damages economic production. In the case of a climatic tipping point abatement efforts can be higher before the tip to prevent tipping, but when the system tips eventually abatement efforts jump up even further since one unit of emissions now leads to more global warming. The expected growth adjusted carbon price is therefore growing faster than economic growth. In the case of an economic tipping point, abatement efforts before tipping are also higher than in the absence of a tipping point to prevent tipping, but once the damage tipping point has happened, the economy is actually smaller in the future and the carbon price jumps down since damages are still proportional to the economy.

We can also combine both types of tipping points in a single simulation. Figure 9 shows a sample path in which the climate tipping point tips very early and in which the economic tipping point tips around 2055. The initial carbon price is equal to 88\$/tC. The left panel indicates that the declining effect of the economic tipping point dominates the increasing effect of the climate tipping point. However, the growth-adjusted carbon price or SCC is now much flatter compared to the left panel of Figure 8. Abatement efforts are higher when both tipping points are present; the optimal abatement rate is 58% after a century.

Of course, in practice, the full impact of a tipping point (e.g., the effects of the melting of the Greenland Ice Sheet) may take a very long time to materialize (Cai & Lontzek, 2019; van der Ploeg & de Zeeuw, 2018). We have abstracted from this, but protracted effects of tipping points are clearly important in terms of the result-

Figure 8: *Risk of an economic tipping point*

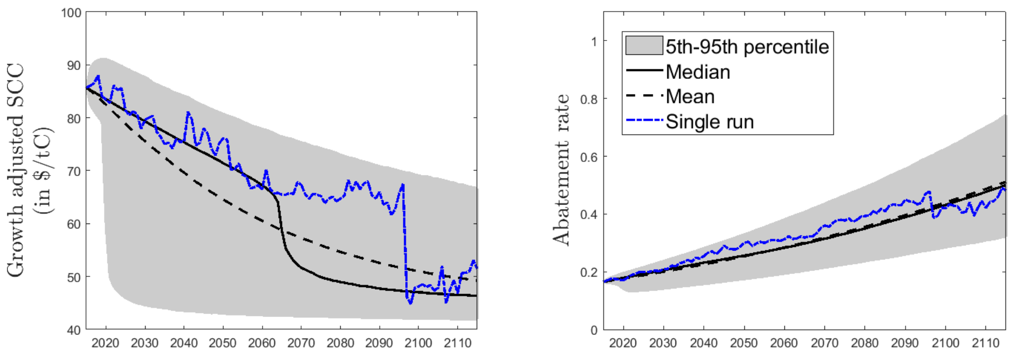

Figure 9: *Risk of two tipping points affecting the climate system and the size of the economy*

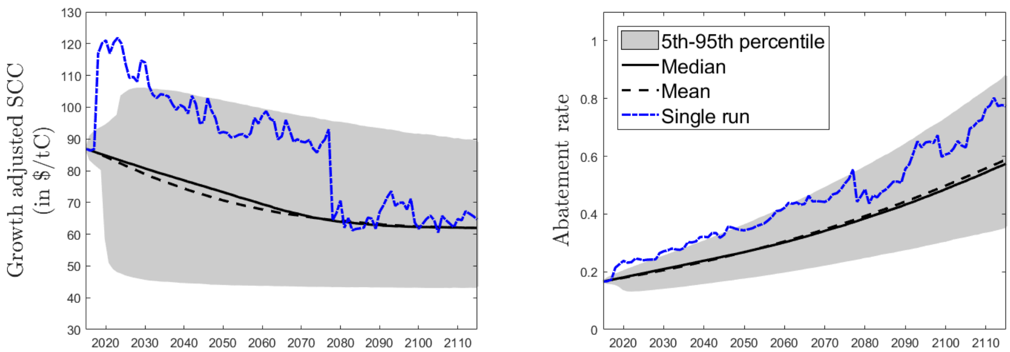

Figure 10: *Effects of a 2 degrees Celsius temperature cap without damages*

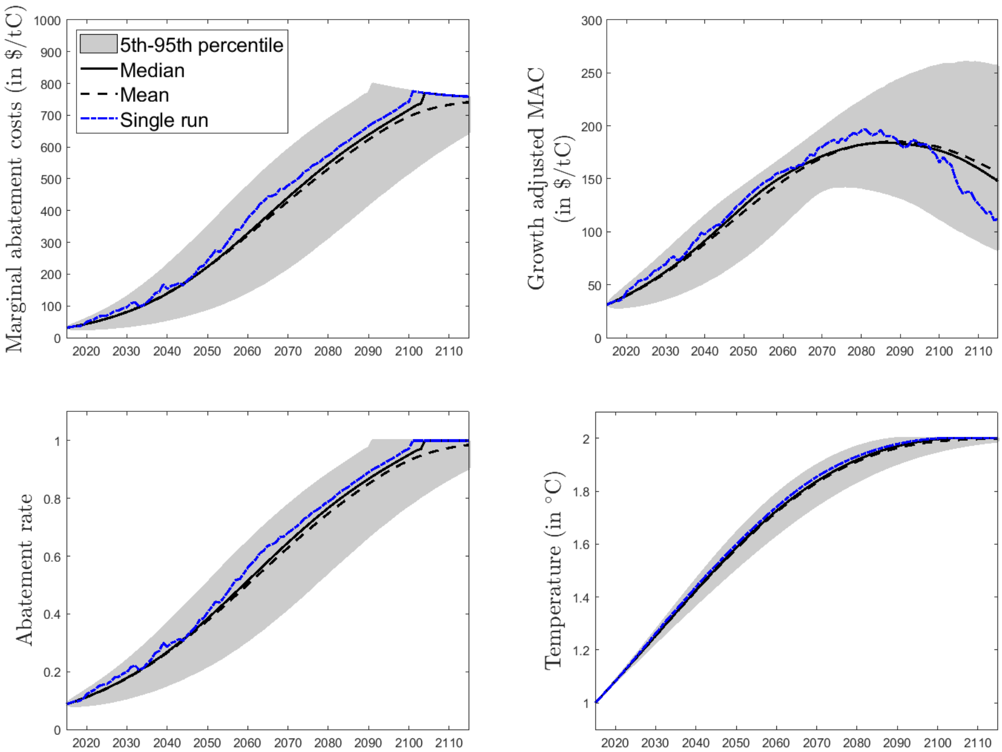

ing time path of optimal policy, which will then change more gradually. It is also important to allow for cascading tipping points where the onset of one tip might increase the likelihood of another tipping point occurring, by more than implied by the temperature-dependence of the hazard rate (Cai et al., 2016; Lemoine & Traeger, 2016). In particular, the downward jump after the damage tip occurs will be smaller in that case since there is the remaining incentive to delay future tipping points.

4.6 Pricing carbon to enforce temperature caps

Although most economists have adopted a welfare-maximizing approach where policymakers internalize the global warming externalities, many governments (as well as central banks and the Network of Greening the Financial System) have followed the IPCC and have decided that the best way to deal with global warming is to enforce a ceiling on global mean temperature.

Given that temperature increases with cumulative emissions, the optimal carbon price must then grow at a rate that is equal to the interest rate plus a risk premium. In Figure 10 we show the optimal climate policies when a cap on global mean temperature of 2 ^{o}C is implemented and where we abstract from damages to global warming. The top left panel indicates a rise in the median carbon price. The carbon price initially also rises when adjusted for the growth rate of the economy, but eventually the growth rate of the carbon price drops below the growth rate of the economy. The

Figure 11: *Expected growth rate of marginal abatement costs under a temperature cap of 2 degrees Celsius without global warming damages to production*

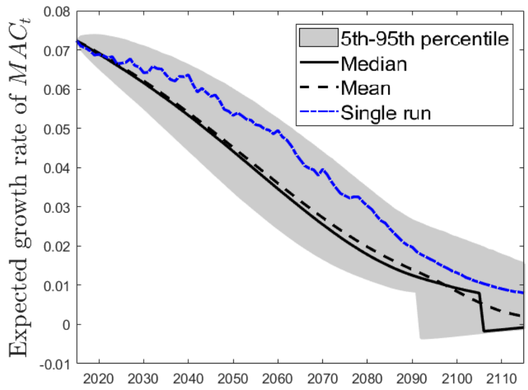

initial carbon price is a bit lower than in the benchmark, but the carbon price grows much faster than the growth rate of the economy. This steep growth in carbon prices ensures a rapid rise in the abatement rate and quick decarbonization of the economy (bottom left panel). Hence, temperature is much lower in a century: 2 ^oC instead of almost 3 ^oC (bottom right panel).

We numerically confirm our theoretical result that the expected growth rate of the carbon price and the MAC indeed equals the risk-free interest rate plus the risk premium. As discussed in section 2.2, we have that:

$$E_t\Big[\frac{dMAC_t}{MAC_t}\Big] = r_t + E_t\Big[-\frac{d[\pi_t, MAC_t]}{\pi_t MAC_t}\Big]. \tag{4.1}$$

What matters for the risk premium is the correlation between marginal abatement costs MAC_t and the stochastic discount factor π_t (via consumption). Marginal abatement costs are given by: $MAC_t = \frac{c_0 e^{-c_1 t} c_2 u_t^{c2}}{\psi_t}$. Uncertainty in MAC_t comes directly from uncertainty in u_t. When endowment and thus consumption is high (in 'good' states of the world), emissions are also high. The emissions control rate must be high in these states as well, in order to comply with the temperature cap. There is thus clearly a positive correlation between MAC_t and consumption C_t. This results in a positive risk premium and the expected growth rate of the carbon price is therefore higher than the rate of interest.

Figure 11 shows however that the expected growth rate of MAC_t is *declining* over time. This is because the risk premium is declining over time. The intuition behind this result is as follows. In the beginning, the abatement rate u_t^* is close to 0. Actual emissions are equal to $(1-u_t^*)\psi_t Y_t$. When there is a negative shock to endowment (and thus to consumption), for example because of a disaster shock, emissions are very responsive. The drop in emissions will then lead to a drop in abatement, since there are less emissions and the temperature cap is the same. Over time, abatement becomes cheaper. So further in the future, emissions are less responsive to a change in output (since u_t^* is larger). The abatement policy reaction to a change in endowment

Figure 12: *Effects of a 2 degrees Celsius temperature cap with convex damages*

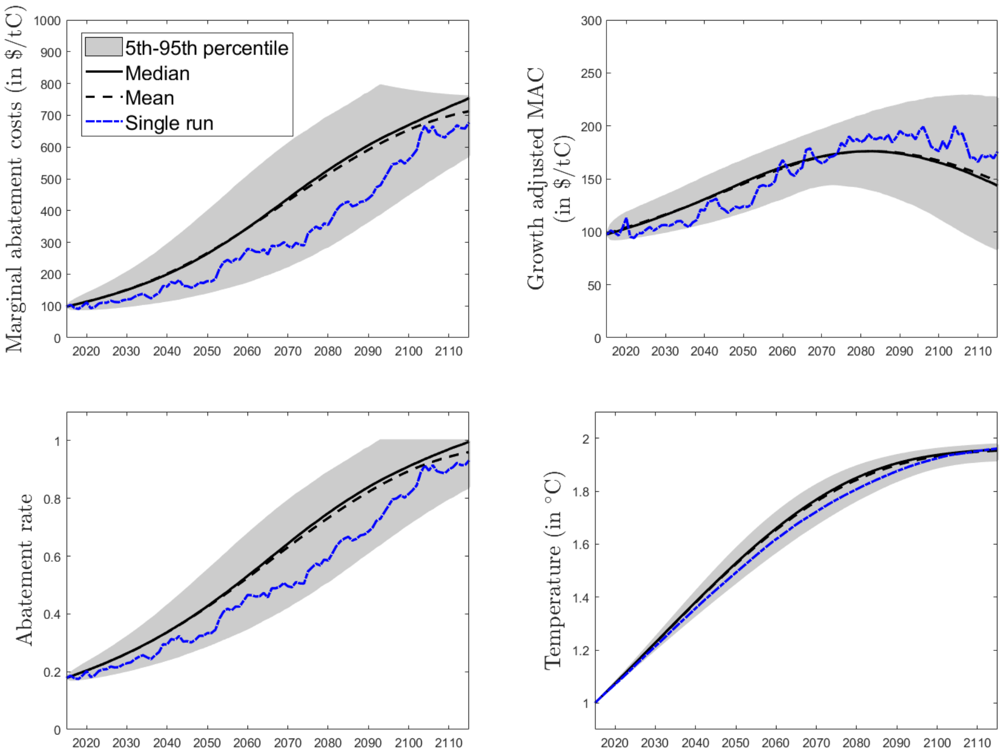

becomes smaller over time. It is also intuitive that there is not much room to change policy when the temperature level is close to the temperature cap. This results in a declining risk premium, and thus the rate of growth of the optimal carbon price tapers off as full mitigation is reached (see top left panel of Figure 10).

In practice, this implies that the optimal path of the carbon price with a temperature cap starts low initially. It then grows fast in the beginning, but the growth rate is declining over time. Around the point where u_t^* is close to 1, the risk premium is close to zero and growth rate is almost equal to the interest rate. The result that the risk premium is declining when the economy gets closer to carbon-neutral cannot be found in a 2-period model like Gollier (2020), since there is only a single uncertain period in that case.

Figure 12 plots the optimal climate policies under a 2 ^oC cap when there are also convex damages from global warming to the aggregate economy. We then find that the growth rate of the optimal path of carbon prices is somewhere in between the risk-adjusted rate of interest and the rate of economic growth (cf. van der Ploeg (2018)). Postponing abatement can be more cost-efficient due to discounting and technological progress in abatement technology, but that also leads to more warming and therefore more damages. The initial price, with both a temperature cap and damages, is therefore much higher (100\$/tC compared to 35\$/tC without damages) and the growth rate lower.

Finally, figure 13 shows the optimal policy function at the initial date (2015) for the

Figure 13: *Optimal policy functions for the abatement rate as function of temperature at the initial date (2015)*

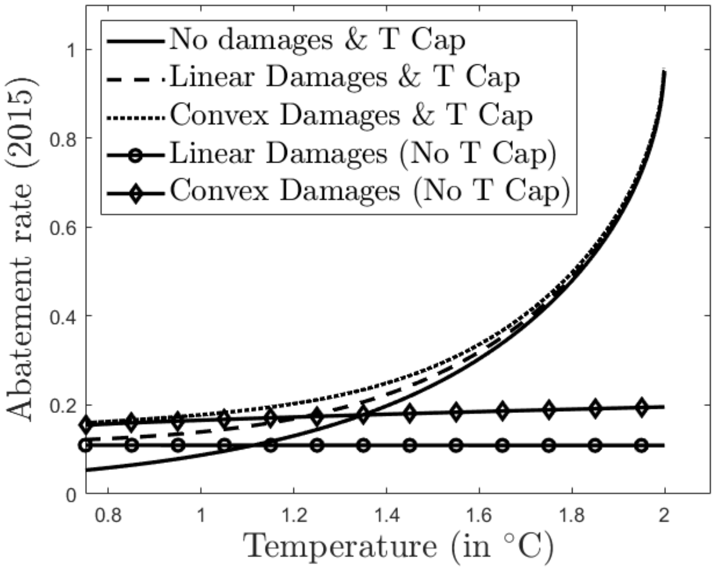

abatement rate in state space, so as a function of global mean temperature. The solid line shows that the abatement rate increases more and more rapidly in the direction of 100% mitigation as the temperature of cap of 2 degrees Celsius is approached. This very nonlinear feature is necessary to ensure that temperature stays below its cap. One can see that the corresponding optimal policy function for the benchmark case of linear damages is flat. The optimal policy function for the case of convex damages is, of course, much higher and slopes gently upwards as the convexity of damages kicks in. Although the policy function with convex damages starts higher than the one with a temperature cap, it rapidly is overtaken as temperature increases. If we combine linear damages and a temperature cap, the policy function starts slightly higher compared to the case with linear damages only. A similar result emerges for the case with convex damages.

4.7 Alternative to benchmark: combination of effects

Most cases discussed in section 4 imply that the growth rate of the carbon price should be at least as high as the rate of economic growth. But in particular economic damage tipping points may at times reverse or slow down that trend. So what do we conclude from all these conflicting trends? It is important to note that all the extensions considered separately are in fact all relevant simultaneously. Therefore we conclude with a combination run featuring convex damages, economic and climate tipping points, and an economists' approach to deriving optimal policy: minimizing damages/maximizing welfare. We leave out active learning by doing, not because it is irrelevant but because it should be stimulated by a dedicated subsidy to new technology, not by manipulating carbon prices; and we have shown that once such

Figure 14: *Convex damages and risk of two tipping points affecting the climate system and the size of the economy*

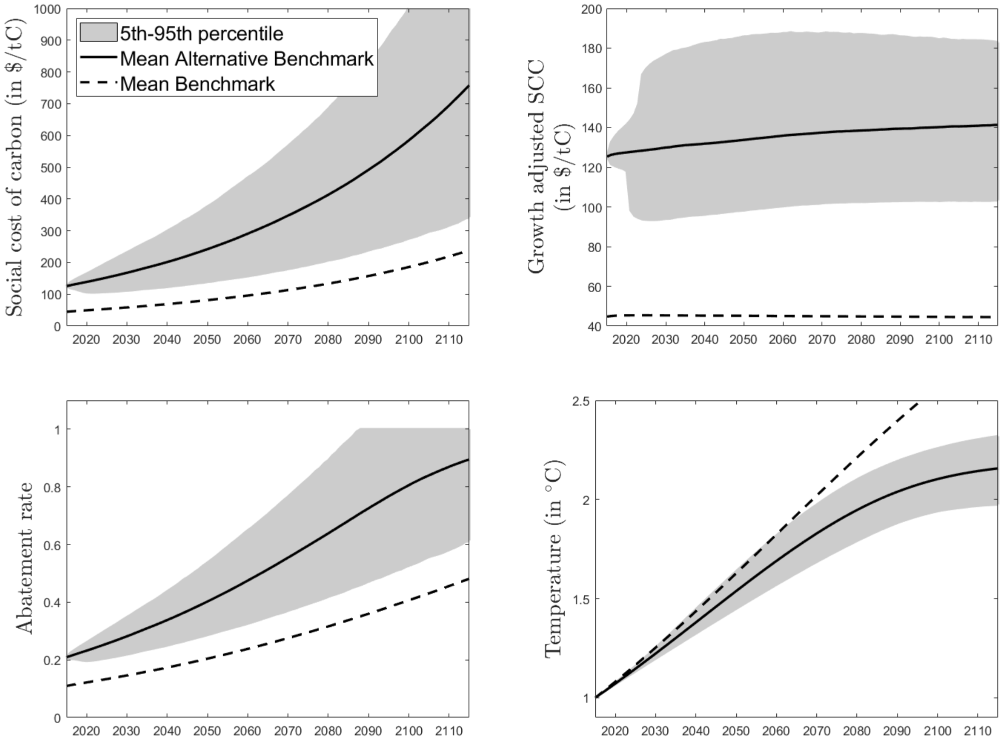

a subsidy is introduced the time path of carbon prices is mostly unaffected[12]. To illustrate this we present an optimal policy simulation run which we consider to be more relevant than the benchmark case in Fig 14.

Clearly, allowing for the combined effects does not only lead to a higher SCC and price of carbon but also to a growth of carbon prices that is somewhat faster than that of the rate of growth of GDP (see top two panels of Fig 14. As a result, the abatement efforts are higher and the transition to a fully carbon-free economy occurs more quickly (third panel) than under the benchmark. The global mean temperature is lower and reaches a plateau of a little more than 2 degrees Celsius relative to pre-industial by the end of this century (fourth panel), approximately equal to the target agreed upon in Paris.

[12]Others have argued that it is optimal to have an upfront spike in carbon prices followed by a decline in carbon prices if there are learning-by-doing effects in renewable energies (e.g. Daniel et al. (2019)). As we have seen, the carbon price that is put forward in Daniel et al. (2019) is a combination of a gradual rise in carbon prices and a spike in renewable energy subsidies. Once the carbon price and renewable energy subsidies are introduced separately, the spike upfront in carbon prices disappears

5 Conclusion

We have shown that convex damages, tipping points and temperature caps all argue in favour of a rising path of carbon prices. Only if there is gradual resolution of uncertainty will there be a declining component in the optimal carbon price, but this effect is dominated by rising components if damages and the economy are growing at empirically plausible rates. Furthermore, convex damages and especially temperature caps require that the carbon prices grow at a faster rate than the economy. Our policy recommendation is therefore that decision makers should start with a significant carbon price and at the same time commit to a rising path of carbon prices. This rising path of carbon prices can, if required by learning-by-doing externalities, be supplemented with renewable energy subsidies separate from taxing emissions through the carbon price. These two policies give the best guarantee for redirecting investments from carbon-intensive to green technologies.

More generally, if in addition to the normal growth uncertainty, risk of macroeconomic disasters and uncertainty about the damage ratio highlighted in our model, account is taken of climatic forms of uncertainty such as uncertainty in the carbon stock and temperature dynamics (van den Bremer & van der Ploeg, 2021) or about tipping of the Greenland or Antarctic Ice Sheet or reversal of the Gulf Stream (Cai & Lontzek, 2019), the optimal response is also a rising path of carbon prices. If integrated assessment models are extended to allow for long-run risk in economic growth with temperature-induced tail risks, the temperature risk premium increases with temperature (Bansal, Kiku, & Ochoa, 2016; Bansal & Yaron, 2004) and it is even more difficult to get a declining carbon price.[13]

An interesting extension is to investigate whether a growing and credible path of carbon prices has the added advantage that businesses get clear incentives to invest in the long-term projects necessary to make the transition from carbon-intensive to carbon-free production. We did not allow for irreversible investments and did not consider hold-up problems in investments resulting from policy transition risk and uncertainty about the path of future carbon prices. However, real option theory might be used to investigate whether by credibly committing to a growing carbon price path corporations are more likely to make the investments that are needed to transition to the carbon-free economy.

References

Acemoglu, D., Aghion, P., Bursztyn, L., & Hemous, D. (2012). The environment and

[13] Olijslagers and van Wijnbergen (2019) show that ambiguity aversion has a major positive impact on the optimal carbon price: the direct effect on the aversion-adjusted valuation of future income flows substantially exceeds the effect ambiguity aversion also has in the opposite direction because it also increases the appropriate discount rate (from the worst-case assumption that optimality requires one to take when faced with the multiple-priors framework (Gilboa & Schmeidler, 1989)). With learning so that after a tipping point it becomes known that the climate sensitivity has increased or carbon sinks have been weakened, the optimal response is to have a rising path of carbon prices before and a rising but higher path after the tipping point (Lemoine & Traeger, 2014).

directed technical change. *American economic review*, *102*(1), 131–66.
Allen, M. R., Frame, D. J., Huntingford, C., Jones, C. D., Lowe, J. A., Meinshausen, M., & Meinshausen, N. (2009). Warming caused by cumulative carbon emissions towards the trillionth tonne. *Nature*, *458*(7242), 1163–1166.
Bansal, R., Kiku, D., & Ochoa, M. (2016). *Price of long-run temperature shifts in capital markets* (Tech. Rep.). National Bureau of Economic Research.
Bansal, R., & Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The journal of Finance*, *59*(4), 1481–1509.
Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, *121*(3), 823–866.
Barro, R. J. (2009). Rare disasters, asset prices, and welfare costs. *American Economic Review*, *99*(1), 243–64.
Barro, R. J., & Jin, T. (2011). On the size distribution of macroeconomic disasters. *Econometrica*, *79*(5), 1567–1589.
Bovenberg, A. L., & Smulders, S. (1995). Environmental quality and pollution-augmenting technological change in a two-sector endogenous growth model. *Journal of public Economics*, *57*(3), 369–391.
Bovenberg, A. L., & Smulders, S. A. (1996). Transitional impacts of environmental policy in an endogenous growth model. *International Economic Review*, 861–893.
Cai, Y., Lenton, T. M., & Lontzek, T. S. (2016). Risk of multiple interacting tipping points should encourage rapid co 2 emission reduction. *Nature Climate Change*, *6*(5), 520–525.
Cai, Y., & Lontzek, T. S. (2019). The social cost of carbon with economic and climate risks. *Journal of Political Economy*, *127*(6), 2684–2734.
Carr, P., & Mayo, A. (2007). On the numerical evaluation of option prices in jump diffusion processes. *The European Journal of Finance*, *13*(4), 353-372.
Coase, R. H. (1960). The problem of social cost. In *Classic papers in natural resource economics* (pp. 87–137). Springer.
Crost, B., & Traeger, C. P. (2013). Optimal climate policy: uncertainty versus monte carlo. *Economics Letters*, *120*(3), 552–558.
Crost, B., & Traeger, C. P. (2014). Optimal co 2 mitigation under damage risk valuation. *Nature Climate Change*, *4*(7), 631–636.
Daniel, K. D., Litterman, R. B., & Wagner, G. (2019). Declining co2 price paths. *Proceedings of the National Academy of Sciences*, *116*(42), 20886–20891.
Dietz, S., & Stern, N. (2015). Endogenous growth, convexity of damage and climate risk: how nordhaus' framework supports deep cuts in carbon emissions. *The Economic Journal*, *125*(583), 574–620.
Dietz, S., & Venmans, F. (2019). Cumulative carbon emissions and economic policy: in search of general principles. *Journal of Environmental Economics and Management*, *96*, 108–129.
Dimson, E., Marsh, P., & Staunton, M. (2011). Equity premia around the world. *Available at SSRN 1940165*.
Duffie, D., & Epstein, L. G. (1992). Stochastic differential utility. *Econometrica: Journal of the Econometric Society*, 353–394.

Epstein, L. G., Farhi, E., & Strzalecki, T. (2014). How much would you pay to resolve long-run risk? *American Economic Review*, *104*(9), 2680–97.
Epstein, L. G., & Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica: Journal of the Econometric Society*, 937–969.
Epstein, L. G., & Zin, S. E. (1991). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis. *Journal of political Economy*, *99*(2), 263–286.
Gerlagh, R., & Liski, M. (2018). Carbon prices for the next hundred years. *The Economic Journal*, *128*(609), 728–757.
Gilboa, I., & Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, *18*(2), 141 - 153.
Gollier, C. (2002a). Discounting an uncertain future. *Journal of public economics*, *85*(2), 149–166.
Gollier, C. (2002b). Time horizon and the discount rate. *Journal of Economic Theory*, *107*(2), 463-473.
Gollier, C. (2008). Discounting with fat-tailed economic growth. *Journal of Risk and Uncertainty*, *37*(2), 171–186.
Gollier, C. (2011). On the underestimation of the precautionary effect in discounting. *The Geneva Risk and Insurance Review*, *36*(2), 95–111.
Gollier, C. (2012). *Pricing the planet's future.* Princeton University Press.
Gollier, C. (2020). The cost-efficiency carbon pricing puzzle.
Golosov, M., Hassler, J., Krusell, P., & Tsyvinski, A. (2014). Optimal taxes on fossil fuel in general equilibrium. *Econometrica*, *82*(1), 41–88.
Hassler, J., Krusell, P., Olovsson, C., & Reiter, M. (2020). On the effectiveness of climate policies. *IIES WP*.
Jensen, S., & Traeger, C. P. (2014). Optimal climate change mitigation under long-term growth uncertainty: Stochastic integrated assessment and analytic findings. *European Economic Review*, *69*, 104–125.
Kelly, D. L., & Kolstad, C. D. (1999). Bayesian learning, growth, and pollution. *Journal of economic dynamics and control*, *23*(4), 491–518.
Kelly, D. L., & Tan, Z. (2015). Learning and climate feedbacks: Optimal climate insurance and fat tails. *Journal of Environmental Economics and Management*, *72*, 98–122.
Lemoine, D. (2021). The climate risk premium: how uncertainty affects the social cost of carbon. *Journal of the Association of Environmental and Resource Economists*, *8*(1), 27–57.
Lemoine, D., & Traeger, C. (2014). Watch your step: optimal policy in a tipping climate. *American Economic Journal: Economic Policy*, *6*(1), 137–66.
Lemoine, D., & Traeger, C. P. (2016). Ambiguous tipping points. *Journal of Economic Behavior & Organization*, *132*, 5–18.
Matthews, H. D., Gillett, N. P., Stott, P. A., & Zickfeld, K. (2009). The proportionality of global warming to cumulative carbon emissions. *Nature*, *459*(7248), 829–832.
Nordhaus, W. D. (2017). Revisiting the social cost of carbon. *Proceedings of the*

National Academy of Sciences, 201609244.
Olijslagers, S. (2021a). The social cost of carbon: Optimal policy versus business-as-usual. *University of Amsterdam.*
Olijslagers, S. (2021b). Solution methods for dsge models in continuous time: application to a climate-economy model. *University of Amsterdam.*
Olijslagers, S., & van Wijnbergen, S. (2019). Discounting the future: on climate change, ambiguity aversion and epstein-zin preferences.
Pigou, A. C. (1920). *The economics of welfare.* London: Macmillan.
Rezai, A., & van Der Ploeg, F. (2017). Abandoning fossil fuel: how fast and how much. *The Manchester School*, *85*, e16–e44.
Shreve, S. E. (2004). *Stochastic calculus for finance ii: Continuous-time models* (Vol. 11). Springer Science & Business Media.
Stern, N., Peters, S., Bakhshi, V., Bowen, A., Cameron, C., Catovsky, S., ... Edmonson, N. (2006). *Stern review: The economics of climate change* (Vol. 30). HM treasury London.
Traeger, C. P. (2021). Ace-analytic climate economy.
van Binsbergen, J. H., Fernández-Villaverde, J., Koijen, R. S., & Rubio-Ramírez, J. (2012). The term structure of interest rates in a dsge model with recursive preferences. *Journal of Monetary Economics*, *59*(7), 634–648.
van den Bremer, T. S., & van der Ploeg, F. (2021, September). The risk-adjusted carbon price. *American Economic Review*, *111*(9), 2782-2810.
van der Ploeg, F. (2018). The safe carbon budget. *Climatic change*, *147*(1), 47–59.
van der Ploeg, F., & de Zeeuw, A. (2018). Climate tipping and economic growth: precautionary capital and the price of carbon. *Journal of the European Economic Association*, *16*(5), 1577–1617.
van Wijnbergen, S., & Willems, T. (2015). Optimal learning on climate change: why climate skeptics should reduce emissions. *Journal of Environmental Economics and Management*, *70*, 17–33.
Vissing-Jørgensen, A., & Attanasio, O. P. (2003). Stock-market participation, intertemporal substitution, and risk-aversion. *American Economic Review*, *93*(2), 383–391.
Weitzman, M. L. (1998). Why the far-distant future should be discounted at its lowest possible rate. *Journal of environmental economics and management*, *36*(3), 201–208.
Weitzman, M. L. (2007). A review of the stern review on the economics of climate change. *Journal of economic literature*, *45*(3), 703–724.
Weitzman, M. L. (2009). Additive damages, fat-tailed climate dynamics, and uncertain discounting. *Economics: The Open-Access, Open-Assessment E-Journal*, *3*.
Weitzman, M. L. (2011). Fat-tailed uncertainty in the economics of catastrophic climate change. *Review of Environmental Economics and Policy*, *5*(2), 275–292.
Weitzman, M. L. (2012). Ghg targets as insurance against catastrophic climate damages. *Journal of Public Economic Theory*, *14*(2), 221–244.

Appendix A Solving for optimal climate policy

Since we include two tipping points each of which can only tip once, we must solve four sub-problems. Define by $V_t^{1,1}$ the value function for the problem where both tipping points have already taken place. $V_t^{1,0}$ is the value function for the problem where the economic (or more precisely the endowment) tipping point has tipped but the climate tipping point has not tipped yet. $V_t^{0,1}$ is defined similarly for the climate tipping point. Lastly, $V_t^{0,0}$ is the value function before any of the two tipping points have taken place. Each of the four sub-problems satisfies its own Hamilton-Jacobi-Bellman (HJB) equation. The HJB-equation for $V_t^{i,j}$, $i \in \{0,1\}$, $j \in \{0,1\}$ equals:

$$
\begin{aligned}
0 = \max_{u_t} \Bigg\{ & f(C_t, V_t^{i,j}) + Z_Y^{i,j} \mu Y_t + \frac{1}{2} Z_{YY}^{i,j} \sigma_Y^2 Y_t^2 + Z_t^{i,j} + Z_T^{i,j} \chi_j (1-u_t) \psi_t Y_t \\
& + Z_\omega^{i,j} \upsilon(\overline{\omega} - \omega_t) + Z_X^{i,j} u_t E_t + \frac{1}{2} Z_{\omega\omega}^{i,j} (\sigma_t^\omega)^2 + \frac{1}{2} Z_{XX}^{i,j} \sigma_X^2 \\
& + \lambda_1 E\Big[Z^{i,j}\Big((1-J_1)Y_t, T_t, \omega_t, X_t, t\Big) - Z^{i,j}\Big(Y_t, T_t, \omega_t, X_t, t\Big)\Big] \\
& + \mathbb{I}_{i=0} \lambda_2 T_t E\Big[Z^{i+1,j}\Big((1-J_2)Y_t, T_t, \omega_t, X_t, t\Big) - Z^{i,j}\Big(Y_t, T_t, \omega_t, X_t, t\Big)\Big] \\
& + \mathbb{I}_{j=0} \lambda_3 T_t E\Big[Z^{i,j+1}\Big(Y_t, T_t, \omega_t, X_t, t\Big) - Z^{i,j}\Big(Y_t, T_t, \omega_t, X_t, t\Big)\Big] \Bigg\}
\end{aligned}
\tag{A.1}
$$

subject to $u_t = 1$ if $T_t = T^{cap}$, where the value function $V_t^{i,j} = Z^{i,j}\Big(Y_t, T_t, \omega_t, X_t, t\Big)$ depends on the four state variables and time and its partial derivatives are denoted by subscripts. $\mathbb{I}$ is an indicator function. Note that the state-variable X_t and all related terms drop out when abatement costs are exogenous.

Define $g_t^{i,j} = h^{i,j}\Big(Y_t, T_t, \omega_t, X_t, t\Big)$ such that $V_t^{i,j} = g_t^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma}$. The derivatives of V are:

$$
\begin{aligned}
Z_t^{i,j} &= h_t^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma}, \\
Z_Y^{i,j} &= h_Y^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma} + g_t^{i,j} Y_t^{-\gamma}, \\
Z_{YY}^{i,j} &= h_{YY}^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma} + 2 h_Y^{i,j} Y_t^{-\gamma} - \gamma g_t^{i,j} Y_t^{-\gamma-1}, \\
Z_k^{i,j} &= h_k^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma}, \quad k \in \{\omega, X, T, t\}, \\
Z_{kk}^{i,j} &= h_{kk}^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma}, \quad k \in \{\omega, X, T\}.
\end{aligned}
\tag{A.2}
$$

Furthermore, we can calculate:

$$
\begin{aligned}
f(C_t, V_t^{i,j}) &= \frac{\beta}{1-1/\epsilon} \frac{\left(Y_t \frac{C_t}{Y_t}\right)^{1-1/\epsilon} - \left(g_t^{i,j} Y_t^{1-\gamma}\right)^{\frac{1}{\zeta}}}{\left(g_t^{i,j} Y_t^{1-\gamma}\right)^{\frac{1}{\zeta}-1}} \\
&= \frac{\beta}{1-1/\epsilon}\Big((g_t^{i,j})^{1-\frac{1}{\zeta}}\Big(\frac{C_t}{Y_t}\Big)^{1-1/\epsilon} Y_t^{1-\gamma} - g_t^{i,j} Y_t^{1-\gamma}\Big) \\
&= \beta\zeta\Big((g_t^{i,j})^{-\frac{1}{\zeta}}\Big(\frac{C_t}{Y_t}\Big)^{1-1/\epsilon} - 1\Big) g_t^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma}.
\end{aligned}
\tag{A.3}
$$

Substituting everything into the HJB equation gives:

$$
\begin{aligned}
0 = \max_{u_t} \Bigg\{ & \beta\zeta\Big((g_t^{i,j})^{-\frac{1}{\zeta}}\Big(\frac{C_t}{Y_t}\Big)^{1-1/\epsilon} - 1\Big) g_t^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma} + \Big(h_Y^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma} + g_t^{i,j} Y_t^{-\gamma}\Big)\mu Y_t \\
& + \frac{1}{2}\Big(h_{YY}^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma} + 2h_Y^{i,j} Y_t^{-\gamma} - \gamma g_t^{i,j} Y_t^{-\gamma-1}\Big)\sigma_Y^2 Y_t^2 \\
& + h_t^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma} + h_T^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma} \chi_j (1-u_t)\psi_t Y_t + h_\omega^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma} \upsilon(\overline{\omega} - \omega_t) \\
& + h_X^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma} \mu_X + \frac{1}{2} h_{\omega\omega}^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma} (\sigma_t^\omega)^2 + \frac{1}{2} h_{XX}^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma} \sigma_X^2 \\
& + \lambda_1 E\Big[h^{i,j}\Big((1-J_1)Y_t, T_t, \omega_t, X_t, t\Big) \frac{\Big((1-J_1)Y_t\Big)^{1-\gamma}}{1-\gamma} - g_t^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma}\Big] \\
& + \mathbb{I}_{i=0} \lambda_2 T_t E\Big[h^{i+1,j}\Big((1-J_2)Y_t, T_t, \omega_t, X_t, t\Big) \frac{\Big((1-J_2)Y_t\Big)^{1-\gamma}}{1-\gamma} - g_t^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma}\Big] \\
& + \mathbb{I}_{j=0} \lambda_3 T_t \Big(g_t^{i,j+1} \frac{Y_t^{1-\gamma}}{1-\gamma} - g_t^{i,j} \frac{Y_t^{1-\gamma}}{1-\gamma}\Big)\Bigg\}.
\end{aligned}
\tag{A.4}
$$

Dividing by $\frac{Y_t^{1-\gamma}}{1-\gamma}$ and rearranging yields:

$$\begin{aligned}
0 = \min_{u_t} \Bigg\{ & \beta\zeta\Big((g_t^{i,j})^{-\frac{1}{\zeta}}\Big(\frac{C_t}{Y_t}\Big)^{1-1/\epsilon} - 1\Big)g_t^{i,j} + \Big((1-\gamma)(\mu - \frac{1}{2}\gamma\sigma_Y^2)\Big)g_t^{i,j} \\
& + h_Y^{i,j}\Big(\mu + (1-\gamma)\sigma_Y^2\Big)Y_t + \frac{1}{2}h_{YY}^{i,j}\sigma_Y^2 Y_t^2 \\
& + h_t^{i,j} + h_T^{i,j}\chi_j(1-u_t)\psi_t Y_t + h_\omega^{i,j}\upsilon(\overline{\omega} - \omega_t) \\
& + h_X^{i,j}\mu_X + \frac{1}{2}h_{\omega\omega}^{i,j}(\sigma_t^\omega)^2 + \frac{1}{2}h_{XX}^{i,j}\sigma_X^2 \\
& + \lambda_1 E\Big[h^{i,j}\Big((1-J_1)Y_t, T_t, \omega_t, X_t, t\Big)(1-J_1)^{1-\gamma}\Big] - g_t^{i,j} \\
& + \mathbb{I}_{i=0}\lambda_2 T_t\Big(E\Big[h^{i+1,j}\Big((1-J_2)Y_t, T_t, \omega_t, X_t, t\Big)(1-J_2)^{1-\gamma}\Big] - g_t^{i,j}\Big) \\
& + \mathbb{I}_{j=0}\lambda_3 T_t\Big(g_t^{i,j+1} - g_t^{i,j}\Big)\Bigg\},
\end{aligned} \tag{A.5}$$

subject to $u_t = 1$ if $T_t = T^{cap}$.

We define the SCC as the welfare loss of emitting one unit of carbon divided by the instantaneous marginal utility of consumption:

$$SCC_t = -\chi\frac{\partial Z_t^{i,j}/\partial T_t}{f_C(C_t, V_t)}. \tag{A.6}$$

The SBL corresponds to all the present and future marginal benefits in terms of lower mitigation costs resulting from using one unit of mitigation more today:

$$SBL_t = \frac{\partial Z_t^{i,j}/\partial X_t}{f_C(C_t, V_t)}. \tag{A.7}$$

The optimality of the abatement rate implies that u_t is chosen such that the MAC is equal to the sum of the SCC and the SBL. Abatement on the one hand leads to lower emissions and on the other hand lowers the costs for future abatement, which implies that $SCC_t + SBL_t = MAC_t$, where:

$$MAC_t = -\frac{\partial C_t/\partial u_t}{E_t} = \frac{Y_t/(1+D_t)}{E_t}\frac{\partial A_t}{\partial u_t} = \frac{1}{\psi_t(1+D_t)}c_0 e^{-c_1 X_t} c_2 u_t^{c_2-1}. \tag{A.8}$$

The relation $SCC_t + SBL_t = MAC_t$ holds if the restriction $u \leq 1$ is not binding. If $u = 1$, then the sum of the SCC and the SBL will be larger than MAC, but it is not possible to abate more (in the absence of negative emissions). The single control variable u_t thus tackles both externalities.

The main insight is that in more disaggregated models of energy use two separate policy instruments should be included. In that case carbon emissions should be priced at the SCC whilst mitigation should be subsidized at the SBL. We also refer to the SCC as the optimal carbon price and to the SBL as the optimal mitigation subsidy, while we note that this relation only holds as long as there is an interior solution to optimal abatement.

We also report the growth-adjusted quantities of the SCC, SBL and MAC to analyse the determinants of these variables other than economic growth. We define the growth-adjusted social cost of carbon by $SCC_t \frac{Y_0}{Y_t}$. The growth-adjusted SCC is scaled with Y_0 to make the initial SCC equal to the actual initial SCC. The growth-adjusted SBL and MAC are defined in the same way.

Appendix B Decentralized market economy

In the decentralized market economy, we need to consider energy producers, households, and the government separately. We assume that the households own the energy producers. We denote the consumer price for fossil fuel by p_t. Since fossil fuel and renewable energy are perfect substitutes, the consumer price for renewable energy is also equal to p_t. We let τ_t and s_t denote the specific tax on fossil fuel and the subsidy on renewable energy, respectively. Fossil fuel use is denoted by F_t and renewable energy use by R_t, so that the mitigation rate is defined by $u_t = \frac{R_t}{F_t+R_t}$. Total energy use is equal to $E_t = \psi_t Y_t$. Profits of and lump-sum rebates to energy producers are denoted by Π_t and S_t, respectively. Profits of energy firms, the household budget constraint and the government budget constraint are given by:

$$\begin{aligned}\Pi_t &= p_t F_t + p_t R_t - \tau_t F_t + s_t R_t - A(u_t, X_t)\frac{Y_t}{1+D_t},\\ C_t &= \frac{Y_t}{1+D_t} + \Pi_t - \tau_t F_t - p_t F_t - p_t R_t,\\ S_t &= \tau_t F_t - s_t R_t.\end{aligned} \tag{B.1}$$

Provided that it is not optimal to fully decarbonize the economy, the first-order optimality conditions for fossil fuel and renewable energy use are:

$$\begin{aligned}p_t &= \tau_t - A_u(u_t, X_t)u_t(1-u_t)\frac{Y_t}{F_t(1+D_t)},\\ p_t &= -s_t + A_u(u_t, X_t)u_t(1-u_t)\frac{Y_t}{R_t(1+D_t)}.\end{aligned} \tag{B.2}$$

Now use that $F_t = (1-u_t)E_t$ and $R_t = u_t E_t$ to obtain:

$$\begin{aligned}p_t &= \tau_t - A_u(u_t, X_t)u_t\frac{Y_t}{E_t(1+D_t)},\\ p_t &= -s_t + A_u(u_t, X_t)(1-u_t)\frac{Y_t}{E_t(1+D_t)}.\end{aligned} \tag{B.3}$$

Combining these two equations gives:

$$\tau_t + s_t = \frac{Y_t}{E_t(1+D_t)}A_u(u_t, X_t) = \frac{1}{\psi_t(1+D_t)}A_u(u_t, X_t). \tag{B.4}$$

Note that $MAC_t = \frac{1}{\psi_t(1+D_t)} A_u(u_t, X_t)$. Imposing a carbon tax and a renewable energy subsidy implies that optimal policy is chosen such that the marginal abatement cost equals the sum of the carbon tax and the renewable energy subsidy. We can therefore replicate optimal policy of the command optimum by setting $\tau_t = SCC_t$ and $s_t = SBL_t$ and rebating any net revenue in lump-sum to the private sector.

Appendix C Numerical implementation

The HJB-equation is a set of partial integro-differential equations. We solve this system of partial differential equations using a finite-difference method. More details on the finite difference method are given in Olijslagers (2021b). One difference is that we have to integrate within the finite-difference scheme to calculate the expectations within the HJB-equation. Recall that the HJB-equation was given by:

$$
\begin{aligned}
0 = \min_{u_t} \Bigg\{ & \beta\zeta\Big((g_t^{i,j})^{-\frac{1}{\zeta}}\Big(\frac{C_t}{Y_t}\Big)^{1-1/\epsilon} - 1\Big)g_t^{i,j} + \Big((1-\gamma)(\mu - \frac{1}{2}\gamma\sigma_Y^2)\Big)g_t^{i,j} \\
& + h_Y^{i,j}\Big(\mu + (1-\gamma)\sigma_Y^2\Big)Y_t + \frac{1}{2}h_{YY}^{i,j}\sigma_Y^2 Y_t^2 \\
& + h_t^{i,j} + h_T^{i,j}\chi_j(1-u_t)\psi_t Y_t + h_\omega^{i,j}\upsilon(\overline{\omega} - \omega_t) \\
& + h_X^{i,j}\mu_X + \frac{1}{2}h_{\omega\omega}^{i,j}(\sigma_t^\omega)^2 + \frac{1}{2}h_{XX}^{i,j}\sigma_X^2 \\
& + \lambda_1 E\Big[h^{i,j}\Big((1-J_1)Y_t, T_t, \omega_t, X_t, t\Big)(1-J_1)^{1-\gamma}\Big] - g_t^{i,j} \\
& + \mathbb{I}_{i=0}\lambda_2 T_t E\Big[h^{i+1,j}\Big((1-J_2)Y_t, T_t, \omega_t, X_t, t\Big)(1-J_2)^{1-\gamma}\Big] - g_t^{i,j} \\
& + \mathbb{I}_{j=0}\lambda_3 T_t\Big(g_t^{i,j+1} - g_t^{i,j}\Big)\Bigg\},
\end{aligned} \tag{C.1}
$$

subject to $u_t = 1$ if $T_t = T^{cap}$. For the numerical implementation, it is useful to make the change of variables $\tilde{Y}_t = \log(Y_t)$ (see for example(Carr & Mayo, 2007)). The

HJB-equation becomes:

$$
\begin{aligned}
0 = \min_{u_t} \Bigg\{ & \beta\zeta\Big((g_t^{i,j})^{-\frac{1}{\zeta}}\Big(\frac{C_t}{Y_t}\Big)^{1-1/\epsilon} - 1\Big)g_t^{i,j} + \Big((1-\gamma)(\mu - \frac{1}{2}\gamma\sigma_Y^2)\Big)g_t^{i,j} \\
& + h_{\tilde{Y}}^{i,j}\Big(\mu - \gamma\sigma_Y^2\Big) + \frac{1}{2}h_{\tilde{Y}\tilde{Y}}^{i,j}\sigma_Y^2 \\
& + h_t^{i,j} + h_T^{i,j}\chi_j(1-u_t)\psi_t e^{\tilde{Y}_t} + h_\omega^{i,j}\upsilon(\overline{\omega} - \omega_t) \\
& + h_X^{i,j}\mu_X + \frac{1}{2}h_{\omega\omega}^{i,j}(\sigma_t^\omega)^2 + \frac{1}{2}h_{XX}^{i,j}\sigma_X^2 \\
& + \lambda_1 E\Big[h^{i,j}\Big(\tilde{Y}_t + \log(1-J_1), T_t, \omega_t, X_t, t\Big)(1-J_1)^{1-\gamma}\Big] - g_t^{i,j} \\
& + \mathbb{I}_{i=0}\lambda_2 T_t\Big(E\Big[h^{i+1,j}\Big(\tilde{Y}_t + \log(1-J_2), T_t, \omega_t, X_t, t\Big)(1-J_2)^{1-\gamma}\Big] - g_t^{i,j}\Big) \\
& + \mathbb{I}_{j=0}\lambda_3 T_t\Big(g_t^{i,j+1} - g_t^{i,j}\Big)\Bigg\}.
\end{aligned}
\tag{C.2}
$$

Define $\tilde{J}_1 = \log(1 - J_1)$. We then obtain that:

$$
\begin{aligned}
& E\Big[h^{i,j}\Big(\tilde{Y}_t + \log(1-J_1), T_t, \omega_t, X_t, t\Big)(1-J_1)^{1-\gamma}\Big] \\
& = \int_0^1 h^{i,j}\Big(\tilde{Y}_t + \log(1-J_1), T_t, \omega_t, X_t, t\Big)(1-J_1)^{1-\gamma}\alpha_1(1-J_1)^{\alpha_1-1}dJ_1 \\
& = \alpha_1 \int_0^1 h^{i,j}\Big(\tilde{Y}_t + \log(1-J_1), T_t, \omega_t, X_t, t\Big)(1-J_1)^{\alpha_1-\gamma}dJ_1 \\
& = \alpha_1 \int_{-\infty}^0 h^{i,j}\Big(\tilde{Y}_t + \tilde{J}_1, T_t, \omega_t, X_t, t\Big)e^{\tilde{J}_1(\alpha_1+1-\gamma)}d\tilde{J}_1.
\end{aligned}
\tag{C.3}
$$

$\tilde{J}_2$ is defined similarly. We use an equally spaced grid for the finite-difference method. To calculate the expectation, we use fourth-order Newton-Cotes integration. The integral is approximated at the points of the finite difference grid with step size δ_Y in the Y dimension:

$$
\begin{aligned}
& \alpha_1 \int_{-\infty}^0 h\Big(\tilde{Y}_t + \tilde{J}_1, T_t, \omega_t, X_t, t\Big)e^{\tilde{J}_1(\alpha_1+1-\gamma)}d\tilde{J}_1 \\
& \approx \alpha_1 \sum_{i=1}^N w_i h\Big(\tilde{Y}_t - i\delta_Y, T_t, \omega_t, X_t, t\Big)e^{-i\delta_Y(\alpha_1+1-\gamma)},
\end{aligned}
\tag{C.4}
$$

where w_i are the Newton-Cotes weights.

We can solve the model analytically when there are no climate damages. We use this as our initial guess at time $t_{max} = 500$, and from there solve the system backwards with time step $\delta_t = 1$. The four-dimensional grid is equally spaced with upper boundaries $[Y^{max} \;\; T^{max} \;\; \omega^{max} \;\; X^{max}]' = [\log(1.5) \;\; 5 \;\; 0.6 \;\; 1750]'$ if there is learning by doing and $[Y^{max} \;\; T^{max} \;\; \omega^{max}]' = [\log(1.5) \;\; 5 \;\; 0.6]'$ without learning by doing. The lower boundaries are equal to $[Y^{min} \;\; T^{min} \;\; \omega^{min} \;\; X^{min}]' = [\log(0.01) \;\; 0.75 \;\; 0 \;\; -25]'$

with learning by doing and $[Y^{min} \;\; T^{min} \;\; \omega^{min}]' = [\log(0.01) \;\; 0.75 \;\; 0]'$ without learning by doing. When a temperature cap is implemented, T^{max} is chosen equal to T^{cap}. Optimal policy is calculated every period by solving for u_t such that the sum of the SCC and the SBL is equal to the marginal abatement cost. If this requires $u_t > 1$, we set $u_t = 1$. The restriction of the temperature cap is implemented by imposing $u_t = 1$ on the boundary $T_t = T^{max}$. The restriction at the boundary also affects optimal policy at all interior grid points of temperature since it will affect the derivative of the value function with respect to temperature. A temperature cap will thus lead to a higher SCC and to a higher emissions control rate u_t.

To speed up the computations, we use a sparse grid combination method. This method solves multiple sub-problems on smaller full grid and combines the solutions of the sub-problems. The sparse-grid combination method that we use in this paper is slightly different from the implementation in Olijslagers (2021b). The difference is that we allow for asymmetry in the grids. Define the 'level' of the grid for dimension i by $L_i, i \in \{T, \omega, X\}$. The number of grid points on the edge of the sparse grid in dimension i is equal to $2^{L_i} + 1$. The level therefore controls the amount of grid points and the accuracy in dimension i. When the value function is non-linear in a specific dimension it is possible to have more grid points in that dimension. This is for example useful when we solve the problem with a temperature cap, since in this case the value function becomes quite non-linear in the temperature dimension.

Let $\mathcal{L} = \left\{ l : \frac{l_T - 1}{L_T - 1} + \frac{l_\omega - 1}{L_\omega - 1} + \frac{l_X - 1}{L_X - 1} \leq 1 \right\}$ be the set of all admissible sub-grids where $l = (l_T, l_\omega, l_X)$. The weight of sub-grid l is equal to:

$$w_l = \sum_{i_T=0}^{1} \sum_{i_\omega=0}^{1} \sum_{i_X=0}^{1} (-1)^{i_T + i_\omega + i_X} \mathbb{I}_{(l_T + i_T, l_\omega + i_\omega, l_X + i_X) \in \mathcal{L}}. \tag{C.5}$$

We solve for g on all subgrids that have a non-zero weight w_l. Note that all grids have different grid points. To find the approximation g_l on sub-grid l in a specific point, we use linear interpolation. We then combine the solutions on all sub-grids by summing over the product of the weight and the solutions: $g = \sum_{l \in \mathcal{L}} w_l g_l$.

Figure 15 shows an example of the sparse-grid combination method in two dimensions. In the example $L_1 = 3$ and $L_2 = 4$, so the sparse grid will be denser in the second dimension. First, the set L is constructed, which in this example consists of the following grids: $(1, 1), (1, 2), (1, 3), (1, 4), (2, 1), (2, 2), (3, 1)$. Of all grids within this set, the grids $(1, 4), (2, 2), (3, 1)$ all have weight $+1$ and the grids $(1, 2), (2, 1)$ have weight -1. The other two grids have weight zero and therefore these do not have to be evaluated.

Figure 15: *The sparse grid combination method.*

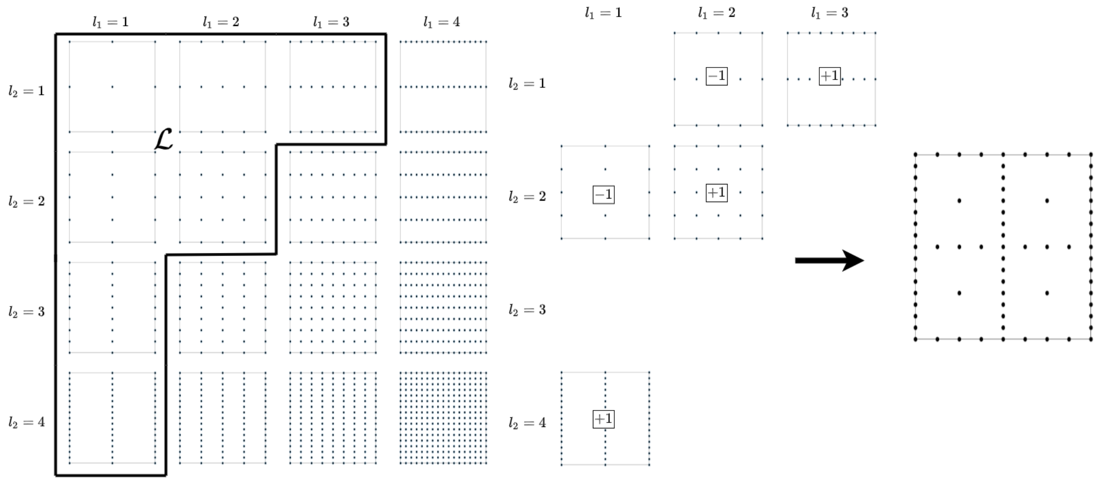

Appendix D Derivation of the growth rate of marginal abatement costs with a temperature cap and no damages

If climate damages are not taken account of and a temperature cap is in place instead, it does not matter for the time at which the temperature cap is reached whether a unit of emissions is abated today or in some period in the future before that time, at least as long as the relationship between temperature and cumulative emissions is linear. Therefore, along the optimal path, a marginal increase of abatement today combined with a marginal decrease of abatement in the future should not lead to a change in welfare. The cost of a marginal increase of abatement today equals MAC_0, while the benefit of a marginal decrease of abatement in time t equals MAC_t. Optimal behaviour therefore implies that $\pi_0 MAC_0 = E_0[\pi_t MAC_t]$ where $\pi_t = \exp\Big(\int_0^t f_V(C_s, V_s)ds\Big) f_C(C_t, V_t)$ is the stochastic discount factor (Duffie & Epstein, 1992). We therefore must have that the product $\pi_t MAC_t$ is a martingale. Now calculate:

$$\frac{d\pi_t MAC_t}{\pi_{t-} MAC_{t-}} = \frac{d\pi_t}{\pi_{t-}} + \frac{dMAC_t}{MAC_{t-}} + \frac{d[\pi_t, MAC_t]}{\pi_{t-} MAC_{t-}}. \tag{D.1}$$

Applying the martingale property and rearranging gives:

$$E_t\Big[\frac{dMAC_t}{MAC_{t-}}\Big] = E_t\Big[-\frac{d\pi_t}{\pi_{t-}}\Big] + E_t\Big[-\frac{d[\pi_t, MAC_t]}{\pi_{t-} MAC_{t-}}\Big], \tag{D.2}$$

where $[\pi_t, MAC_t]$ denotes the quadratic covariation for the processes π_t and MAC_t. Note that the first term $E_t\Big[-\frac{d\pi_t}{\pi_t}\Big]$ is exactly equal to the real risk-free interest rate, while the second term is a risk premium related to the correlation between the stochastic discount factor and the marginal abatement costs. Equation (D.2) implies that the optimal carbon price must grow at a rate equal to the sum of the real risk-free interest rate plus an interest premium to be determined, similar to Gollier (2020). In the following we derive the risk-free rate and the risk premium.

D.1 Derivation of stochastic discount factor, interest rate, and risk premium

We can work out the stochastic discount factor π_t and the marginal abatement cost function MAC_t. The model without climate damages can be written as follows. The log-endowment follows from:

$$d\tilde{Y}_t = (\mu - \frac{1}{2}\sigma_Y^2)dt + \sigma_Y dW_t^Y + \log(1-J_1)dN_{1,t}. \tag{D.3}$$

Consumption is equal to endowment minus abatement expenditure: $C_t = (1 - A_t)e^{\tilde{Y}_t}$, where the abatement cost function $A_t = c_0 e^{-c_1 t} u_t^{c_2}$. Define the consumption-endowment ratio $\xi_t = 1 - A_t$, which depends on the two state variables and time. The state variable T_t (temperature) follows from:

$$dT_t = \chi(1-u_t)\psi_t e^{\tilde{Y}_t} dt. \tag{D.4}$$

The temperature cap adds the restriction $u_t = 1$ if $T_t = T^{cap}$. The HJB-equation corresponding to the value function V_t for this problem is thus given by:

$$\begin{aligned}
0 = \min_{u_t} \Bigg\{ & \beta\zeta\Big(g_t^{-1/\zeta}\Big(\frac{C_t}{Y_t}\Big)^{1-1/\epsilon} - 1\Big)g_t \\
& + (1-\gamma)\Big(\mu - \frac{1}{2}\gamma\sigma_Y^2\Big)g_t \\
& + h_{\tilde{Y}}\Big(\mu - \gamma\sigma_Y^2\Big) + \frac{1}{2}h_{\tilde{Y}\tilde{Y}}\sigma_Y^2 \\
& + h_t + h_T\chi(1-u_t)\psi_t e^{\tilde{Y}_t} \\
& + \lambda_1\Big(E\Big[h\Big(\tilde{Y}_t + \log(1-J_1), T_t, t\Big)(1-J_1)^{1-\gamma}\Big] - g_t\Big)\Bigg\}
\end{aligned} \tag{D.5}$$

subject to $u_t = 1$ if $T_t = T^{cap}$. The derivatives of instantaneous utility $f(C_t, V_t)$ can be calculated as:

$$\begin{aligned}
f_C(C_t, V_t) &= \frac{\beta C_t^{-1/\epsilon}}{\Big((1-\gamma)V_t\Big)^{1/\zeta - 1}}, \\
f_V(C_t, V_t) &= \beta\zeta\Big((1-1/\zeta)C_t^{1-1/\epsilon}\Big((1-\gamma)V_t\Big)^{-1/\zeta} - 1\Big).
\end{aligned} \tag{D.6}$$

Now substitute in $V_t = g_t \frac{Y_t^{1-\gamma}}{1-\gamma}$ and $\xi_t = \frac{C_t}{Y_t}$ to obtain:

$$\begin{aligned}
f_C(C_t, V_t) &= \beta\xi_t^{-1/\epsilon} g_t^{1-1/\zeta} Y_t^{-\gamma}, \\
f_V(C_t, V_t) &= \beta\zeta\Big((1-1/\zeta)\xi_t^{1-1/\epsilon} g_t^{-1/\zeta} - 1\Big).
\end{aligned} \tag{D.7}$$

Substituting this and $Y_t^{-\gamma} = e^{-\gamma \tilde{Y}_t}$ into the stochastic discount factor gives:

$$\pi_t = \exp\Big(\int_0^t \beta\zeta\Big((1-1/\zeta)\xi_s^{1-1/\epsilon} g_s^{-1/\zeta} - 1\Big)ds\Big)\beta\xi_t^{-1/\epsilon} g_t^{1-1/\zeta} e^{-\gamma\tilde{Y}_t}. \tag{D.8}$$

Define $f_t = \xi_t^{-1/\epsilon} g_t^{1-1/\zeta} = \nu(\tilde{Y}_t, T_t, t)$. Write π_t as a differential equation:

$$\begin{aligned} \frac{d\pi_t}{\pi_t} &= \beta\zeta\Big((1-1/\zeta)\xi_t^{1-1/\epsilon} g_t^{-1/\zeta} - 1\Big)dt \\ &\quad + \frac{df_t}{f_t} + \frac{de^{-\gamma\tilde{Y}_t}}{e^{-\gamma\tilde{Y}_t}} + \frac{d[f_t, e^{-\gamma\tilde{Y}_t}]}{f_t e^{-\gamma\tilde{Y}_t}} \\ &\quad + \frac{\nu\Big(\tilde{Y_{t-}} + \log(1-J_1), T_t, t\Big)e^{-\gamma\tilde{Y_{t-}}}(1-J_1)^{-\gamma} - f_{t-}e^{-\gamma\tilde{Y_{t-}}}}{f_{t-}e^{-\gamma\tilde{Y_{t-}}}} dN_{1,t}. \end{aligned} \tag{D.9}$$

Applying Ito's lemma to $\tilde{Y}_t$ (without jump terms) gives:

$$\frac{de^{-\gamma\tilde{Y_{t-}}}}{e^{-\gamma\tilde{Y_{t-}}}} = -\gamma\Big(\mu - \frac{1}{2}(\gamma+1)\sigma_Y^2\Big)dt - \gamma\sigma_Y dW_t^Y. \tag{D.10}$$

Similarly, we apply Ito's lemma to f_t to get:

$$\frac{df_t}{f_t} = \Big(\frac{\nu_t}{f_t} + \frac{\nu_T}{f_t}\chi(1-u_t)\psi_t e^{\tilde{Y}_t} + \frac{\nu_{\tilde{Y}}}{f_t}(\mu - \frac{1}{2}\sigma_Y^2) + \frac{1}{2}\frac{\nu_{\tilde{Y}\tilde{Y}}}{f_t}\sigma_Y^2\Big)dt + \frac{\nu_{\tilde{Y}}}{f_t}\sigma_Y dW_t^Y. \tag{D.11}$$

Define $\mu_{f,t} = \frac{\nu_t}{f_t} + \frac{\nu_T}{f_t}\chi(1-u_t)\psi_t e^{\tilde{Y}_t} + \frac{\nu_{\tilde{Y}}}{f_t}(\mu - \frac{1}{2}\sigma_Y^2) + \frac{1}{2}\frac{\nu_{\tilde{Y}\tilde{Y}}}{f_t}\sigma_Y^2$. Putting everything together yields:

$$\begin{aligned} \frac{d\pi_t}{\pi_{t-}} &= \Big\{\beta\zeta\Big((1-1/\zeta)\xi_t^{1-1/\epsilon} g_t^{-1/\zeta} - 1\Big) - \gamma\Big(\mu - \frac{1}{2}(\gamma+1)\sigma_Y^2\Big) + \mu_{f,t} \\ &\quad - \gamma\sigma_Y^2\frac{\nu_{\tilde{Y}}}{f_t}\Big\}dt - \gamma\sigma_Y dW_t^Y + \frac{\nu_{\tilde{Y}}}{f_t}\sigma_Y dW_t^Y \\ &\quad + \frac{\nu\Big(\tilde{Y_{t-}} + \log(1-J_1), T_t, t\Big)(1-J_1)^{-\gamma} - f_{t-}}{f_{t-}} dN_{1,t}. \end{aligned} \tag{D.12}$$

The interest rate can be calculated as:

$$\begin{aligned} r_t dt &= E_t\Big[-\frac{d\pi_t}{\pi_t}\Big] \\ &= \Big\{-\beta\zeta\Big((1-1/\zeta)\xi_t^{1-1/\epsilon} g_t^{-1/\zeta} - 1\Big) + \gamma\Big(\mu - \frac{1}{2}(\gamma+1)\sigma_Y^2\Big) - \mu_{f,t} \\ &\quad - \lambda_1\frac{E\Big[\nu\Big(\tilde{Y_{t-}} + \log(1-J_1), T_t, t\Big)(1-J_1)^{-\gamma}\Big] - f_{t-}}{f_{t-}}\Big\}dt. \end{aligned} \tag{D.13}$$

Marginal abatement costs are given by:

$$MAC_t = -\frac{\partial C_t/\partial u_t}{\psi_t Y_t} = \psi_t^{-1}\frac{\partial A_t}{\partial u_t} = \psi_t^{-1}c_0 e^{-c_1 t}c_2 u_t^{c_2-1} = \eta(\tilde{Y}_t, T_t, t). \tag{D.14}$$

We then get:

$$
\begin{aligned}
&E_t\Big[\frac{d[\pi_t, MAC_t]}{\pi_{t-}MAC_{t-}}\Big] = -\gamma\sigma_Y^2\frac{\eta_{\tilde{Y}}}{MAC_{t-}}dt + \sigma_Y^2\frac{\nu_{\tilde{Y}}\eta_{\tilde{Y}}}{f_t MAC_{t-}}dt \\
&- \lambda_1\frac{E\Big[\nu\Big(\tilde{Y_{t-}} + \log(1-J_1), T_t, t\Big)(1-J_1)^{-\gamma}\Big] - f_{t-}}{f_{t-}}dt \\
&+ \lambda_1\frac{E\Big[\nu\Big(\tilde{Y_{t-}} + \log(1-J_1), T_t, t\Big)\eta\Big(\tilde{Y_{t-}} + \log(1-J_1), T_t, t\Big)(1-J_1)^{-\gamma}\Big]}{MAC_{t-}f_{t-}}dt \\
&- \lambda_1\frac{E[\eta\Big(\tilde{Y_{t-}} + \log(1-J_1), T_t, t\Big)]}{MAC_{t-}}dt.
\end{aligned}
\tag{D.15}
$$