Optimal emissions under exogenous and endogenous learning

Coppens Léo[*], Dietz Simon[†], Venmans Frank[‡]

31st January 2022

Abstract

Some abatement technologies become cheaper over time due to spillover effects from other sectors. In a climate cost-benefit analysis, this process is "exogenous", in the sense that it only depends on time. Other abatement technologies become cheaper only if they are deployed on a large scale. This "endogenous" process of learning by doing decreases abatement costs as a function of cumulative abatement. We aim to shed light on the differentiated impact of endogenous and exogenous learning on the optimal mitigation path. This is particularly important in a time when many models and scenarios are ignoring the dynamic characteristic of learning by doing. We develop a cost-benefit integrated assessment model which includes both types of learning dynamics as well as inertia. Theoretically, endogenous learning leads to a supplementary term in the optimality condition: the "learning gains", whereas exogenous learning only creates an incentive to postpone climate action. We show analytically and numerically that including endogenous and exogenous learning steepens the abatement path. In a cost-benefit analysis, both types of learning leads to lower peak warming. Moreover, endogenous learning leads to negative emissions in the long run. Besides, the common practice of modelling endogenous learning as an exogenous process underestimates optimal abatement by 9% in 2050.

Keywords: climate change mitigation, induced innovation, integrated assessment models, learning by doing, optimal abatement path

JEL references: C61, O30, Q54, Q55, Q58

[*]Department of Finance, Warocqué School of Business and Economics, University of Mons, Place Warocqué 17, 7000 Mons, Belgium, leo.coppens@umons.ac.be

[†]Department of Geography and Environment, ESRC Centre for Climate Change Economics and Policy, Grantham Research Institute on Climate Change and the Environment, London School of Economics and Political Science, UK, S.Dietz@lse.ac.uk

[‡]ESRC Centre for Climate Change Economics and Policy, Grantham Research Institute on Climate Change and the Environment, London School of Economics and Political Science, UK, f.venmans1@lse.ac.uk

1 Introduction

Defining optimal emissions trajectories is crucial to inform climate policy. Many models have been developed to do so, and numerous scenarios have been produced. Most notoriously, Intergovernmental Panel on Climate Change (IPCC) reports gather many trajectories coming from various modelling tools.

In the transition to the climate objectives, technological change (TC) plays an important role. Gillingham et al. (2008) made a review on how TC is included in models: it is either modelled exogenously or endogenously, and in the latter case, this can be included through R&D investments, learning by doing, and price-induced learning. Regarding learning-by-doing, Wright (1936) was certainly the first to quantify technological learning: he found that unit labour costs in the aircraft industry decreased with cumulative output. Our analysis will use this way of including TC.

The learning dynamics and how they are modelled may have significant consequences on the optimal path. As concluded by Grubb et al. (2020), dynamic characteristics (including TC) "can have a radical impact on the optimal approach to tackling climate change".

However, among the main widely used models and in the IPCC reports, many do not include endogenous learning dynamics in their optimisation. Learning is often only exogenously included. It is also the case in many bottom-up energy models. Therefore, better understanding the effect of learning dynamics on the optimal path will help one's analysis of the scenarios produced by those models, which disregard the incentive created by learning by doing for early abatement. In our results section, we show that approximating endogenous learning by an exogenous process as it is done in many models underestimates the optimal carbon price and abatement (by 9% in 2050).

To do so, we consider a cost-benefit model where marginal abatement costs depend on time directly (exogenous learning) or on cumulative abatement A (endogenous learning), such that we can write the decrease of marginal abatement costs over time as the sum of both processes.

$$\frac{dMAC(t,A)}{dt} = \underbrace{\frac{\partial MAC}{\partial t}}_{exogenous\ learning} + \underbrace{\frac{\partial MAC}{\partial A}\frac{dA}{dt}}_{endogenous\ learning} \tag{1}$$

We calibrate our MAC to the climate scenarios database of IPCC and NGFS (Network for Greening the Financial System) (Rogelj et al., 2018; Huppmann et al., 2019; NGFS, 2021). This framework allows us to investigate the consequences of learning on the optimal abatement path and to focus on the differences among the impact of endogenous, exogenous learning and the absence of learning. This is highly relevant for 3 main reasons: learning dynamics matter when looking at the optimal path; many models do not include endogenous learning; the question of the impact of endogenous (versus exogenous) learning does not seem to be sorted out in the literature.

Theoretically, endogenous learning leads to a supplementary term in the optimality condition: the "learning gains", whereas exogenous learning only creates an incentive to postpone climate action. However, compared to a model without learning, including exogenous or endogenous learning steepens the abatement path and leads to less emissions reduction in the short term, but thanks to more efforts in the medium/long term, peak warming is lower. In order to isolate the exogenous learning effect, we compare the exogenous learning model with a model where there is no learning, but identical marginal abatement costs (by fitting a polynomial). We show that the exogenous learning dynamic has no effect on the optimal trajectory. Following a similar approach, we isolate the endogenous learning effect and show its significant impact on the optimal trajectory (the endogenous dynamic leads to lower emissions throughout this century).

The structure of the paper is as follow. Section 2 briefly presents the related literature. Section 3 introduces the model and shows analytical results. Section 4 explains how we isolate the exogenous and endogenous effect on the optimal path. Section 5 presents how we calibrated our model to the IPCC and NGFS database. Section 6 provides numerical results. We analyse the central cases related to the different types of learning dynamics and provide a sensitivity analysis on the learning parameters as well as on the discount rate. Then, we isolate the effects of endogenous and exogenous learning dynamics. Section 7 concludes and discusses the main findings.

2 Literature

Recently, Grubb et al. (2021) criticised optimising IAMs because many models do not consider path dependency (inertia and learning dynamics). We actually base our paper on one of the models they criticised (Dietz and Venmans, 2019) and add path dependency via endogenous learning and inertia while still developing a surprisingly simple model to assess optimal climate policy.

The most closely related paper to our work is certainly the one from Goulder and Mathai (2000), in the sense that we both look at the impact of induced technological change on the optimal path and that they derive analytical and numerical results to support their analysis. Goulder and Mathai (2000) concluded that the effect of learning by doing on initial abatement is ambiguous but that in most numerical simulations it rises. However, in our simulations, we show that abatement is lower in the short term when including endogenous learning, compared to a model without learning calibrated on IPCC and NGFS data. Moreover, they concluded that the impact of learning by doing on the timing of abatement is very slight but that the effect on cumulative abatement can be large. We show that indeed the effect on cumulative emissions over time is quite large, but that the optimal timing of emissions is significantly impacted as well.

Next, Manne and Richels (2004) is one of the only studies looking specifically at the consequences of endogenous versus exogenous learning on the optimal path. They showed that "including learning-by-doing does not significantly alter the conclusions of previous studies that treated technology cost as exogenous". We show otherwise, that the difference between exogenous and endogenous learning does significantly alter the conclusions on optimal path.

Popp (2004) focused on the same subject, but they included endogenous technological change (in the DICE model) via investments in energy R&D. The paper focused on the differences (in terms of welfare gains, emissions trajectory, and temperature) between a scenario with endogenous R&D and a scenario with only exogenous learning. To do so, Popp (2004) used its modified version of the DICE model "ENTICE" and defined the exogenous case as having "base-case no policy" energy R&D. He found out that the welfare gains (in the long term and in total) can be quite important with endogenous learning but that the emissions trajectories and the temperature are almost identical (their optimal temperature is around 5 degrees in both cases). With our model and our parameters, we find that learning does decrease the optimal temperature (which is way smaller in our scenarios) and that the trajectories are not quite identical. However, we do not include R&D explicitly in our model.

Some disadvantages of learning-by-doing (LbD) versus R&D have been highlighted by Gillingham et al. (2008). Firstly, the reduced-form nature of LbD is a black-box where causality does not appear clearly. However, this simple formulation of TC allows us in our analytical IAM to remain simple and "transparent". Secondly, there is a lack of empirical data, and it is not straightforward to calibrate (note that this particular remark is also applicable to the R&D approach). In our analysis, the IPCC and NGFS database allows us to calibrate the model on a wide pool of scenarios.

There is a literature focusing on optimal investment in R&D and its interaction with market structure. R&D is typically undersupplied by the market because there are knowledge spillovers and the returns to R&D are not fully appropriable to the company that develops the technology. As a result, optimal policy involves both a carbon tax and subsidies (Acemoglu et al., 2012; Fischer and Newell, 2008). Most of the emissions reduction are due to the carbon tax and it is the most efficient single policy but combining the policies lower the cost significantly (Fischer and Newell, 2008). However, as pointed out by Gerlagh et al. (2009), it may be hard to target R&D subsidies correctly in the private sector, so a carbon tax (the second-best policy) could then be implemented (which would be significantly higher than in a first-best scenario). In our cost-benefit model, we do not model R&D directly, but our MAC can be understood as the marginal abatement cost function in a context of optimal or plausible R&D subsidies[1].

[1]We do not model R&D explicitely or directly: we assume an exogenous R&D process that leads to 2 different types of learning. Some R&D enable spillovers from the rest of the economy. These potential spillovers depend on time and drive our exogenous learning. Other R&D investments will develop technology that builds incrementally on knowledge that has already accumulated within green sectors but could not have been developed elsewhere. This R&D is one of the drivers of our endogenous learning. Our MAC equation (see Section 3) can therefore be understood as the marginal abatement cost function in a context of optimal or plausible R&D subsidies. More specifically, since we calibrate our model on existing bottom-up models in the IPCC and NGFS database, we use the assumptions regarding R&D in these studies (they do not mention the amount of R&D subsidies).

3 Analytical results in a simple climate model

We start from the climate model in Dietz and Venmans (2019), and add exogenous and endogenous learning. We also add inertia in Appendix.

Define abatement $a = E_{BAU} - E$, with E_{BAU} being constant business as usual emissions, expressed as CO2 equivalents.

Define cumulative emissions S, with $\dot{S} = E = E_{BAU} - a$. Assume temperature to be proportional to cumulative emissions $T = \zeta S$. We use $\zeta = 0.0006°C/GtCO_2eq$ considering that the IPCC (IPCC, 2021) estimates that an emissions budget of 1000GtCO2 will lead to 0.73°C long term warming from 2020 onwards and assuming that for each tonne of CO2 emitted, there will be on average 0.2 tonne of nonCO2 equivalent emitted $\left(\frac{0.73°C}{1200GtCO2eq} = 0.0006\right)$.

Define cumulative abatement relative to cumulative abatement at time zero as $A(S,t) = \frac{A_0 + \int_0^t a_\tau d\tau}{A_0} = \frac{A_0 + S_0 + E_{BAU}t - S_t}{A_0}$. Therefore, total cumulative abatement is $A_0 A$.

Assume the following linear marginal abatement cost function as a % of consumption,

$$MAC\% = \varphi(t) a A^{-\chi}. \tag{2}$$

The decrease of parameter φ over time represents exogenous learning. The factor $A^{-\chi}$ represents learning by doing. It corresponds to a standard linear log-log learning curve such that for every percentage increase in cumulative abatement, the marginal abatement cost, expressed as a percentage of production, decreases by χ%.

Exogenous labour-augmenting technology improves labour productivity, leading to a BAU consumption growth rate of rate g.

As in Dietz and Venmans (2019), we assume that capital is close to its steady state, and that the effect of climate on the growth rate is constant, such that the savings rate is constant. As a result, we will write the model in terms of consumption, which is proportional to production.

Consider consumption per unit of effective labour, with population growing at rate n and standardised at 1 at time zero $c = \frac{C}{e^{(n+g)t}}$. Finally, assume that climate damages are quadratic and proportional to consumption, leading to the following expression for consumption per unit of effective labour

$$c = c_{BAU_0} e^{\left(-\frac{\varphi_t}{2} a^2 A^{-\chi} - \frac{\gamma}{2}\zeta^2 S^2\right)}. \tag{3}$$

Note that γ is the damage function coefficient.

We use a utility function with constant elasticity of marginal utility. The welfare functional writes

$$max \int_0^\infty e^{-\delta t} \frac{\left(\frac{C}{e^{nt}}\right)^{(1-\eta)}}{1-\eta} dt \tag{4}$$

Defining the abstract concept $U\left(c\left(a,S,t\right)\right) = \frac{c^{1-\eta}}{1-\eta} = \frac{c_{BAU_0}^{1-\eta}}{1-\eta} e^{(1-\eta)\left(-\frac{\varphi_t}{2} a^2 A^{-\chi} - \frac{\gamma}{2}\zeta^2 S^2\right)}$ and $r = \delta - n + \eta g$ (with δ being the utility discount rate and η the negative of the elasticity of marginal utility), we can rewrite the welfare functional as

$$max \int_0^\infty e^{-(r-g)t} U dt \tag{5}$$

Then, the current value Hamiltonian is

$$H^{CV} = U - \lambda^S \left(E_{BAU} - a \right) \tag{6}$$

The first order conditions (FOC) are

$$-U_a = \lambda^S \Leftrightarrow U \left(1 - \eta\right) \left(\varphi_t a A^{-\chi}\right) = \lambda^S \tag{7}$$

$$\dot{\lambda^S} = (r-g)\lambda^S + U\left(1-\eta\right)\left(-\gamma\zeta^2 S - \frac{\chi\varphi_t}{2A_0} a^2 A^{-\chi-1}\right) \tag{8}$$

Note that the shadow price of cumulative emissions has now two components: marginal damages and learning, both of the same sign. Avoiding a unit of emissions does not only give lower damages over the future path, it also decreases abatement costs of the future path. Let's call this the "learning incentive".

3.1 MAC path

Let's define the marginal abatement cost expressed in units of consumption $MAC_t = -C_a = C_t \varphi_t a A^{-\chi} = c_t e^{gt} \varphi_t a A^{-\chi}$. We can integrate the equation 8

$$\lambda^S = \int_t^\infty e^{-(r-g)(\tau-t)} c_\tau^{1-\eta} \left(\gamma\zeta^2 S + \frac{\chi\varphi_\tau}{2A_0} a^2 A^{-\chi-1}\right) d\tau. \tag{9}$$

Substitute lambda using equation 7 gives

$$c_t^{1-\eta}\left(\varphi_t a A^{-\chi}\right) = \int_t^\infty e^{-(r-g)(\tau-t)} c_\tau^{1-\eta} \left(\gamma\zeta^2 S + \frac{\chi\varphi_\tau}{2A_0} a^2 A^{-\chi-1}\right) d\tau \tag{10}$$

Multiply by $c_t^\eta e^{gt}$,

$$MAC_t = \int_t^\infty e^{-r(\tau-t)} \left(\frac{c_\tau}{c_t}\right)^{-\eta} \left(\underbrace{c_\tau \gamma\zeta^2 S}_{c_S = marg\ damage} + \underbrace{\frac{c_\tau \chi\varphi_\tau}{2A_0} a^2 A^{-\chi-1}}_{c_A = marg\ gain\ learning} \right) d\tau \tag{11}$$

$$\begin{aligned} MAC\% &= \varphi_t a A^{-\chi} \\ &= \int_t^\infty e^{-(r-g)(\tau-t)} \left(\frac{c_\tau}{c_t}\right)^{1-\eta} \left(\underbrace{\gamma\zeta^2 S}_{\% marg\ damage} + \underbrace{\frac{\chi\varphi_t}{2A_0} a^2 A^{-\chi-1}}_{\% marg\ gain\ learning} \right) d\tau \end{aligned} \tag{12}$$

Note that since c is consumption per unit of effective labour, the factor $\left(\frac{c_\tau}{c_0}\right)^\eta$ is approximately 1 (decreasing as a rate less than 0.02%). We obtain the well known expression that at any point in time, the marginal abatement cost should be equal to the social cost of carbon. However, the social cost of carbon is not only the sum of all discounted marginal damages, it also includes the sum of all future gains from endogenous learning.

In a model with exogenous learning, we obtain the following optimality condition : MAC is equal to the present value of all future marginal damages, as in a model without learning. In a model with endogenous learning the optimality condition is: MAC equals the present value of all future damages plus marginal learning gains from a unit of abatement today. For a given temperature, the incentive to abate is larger (for a given temperature path, a marginal increase in endogenous learning, will increase the MAC).

Substituting $MAC = \frac{\lambda^S}{U(1-\eta)} ce^{gt}$ from equation 7 into equation 8 allows us to obtain the growth rate of the MAC

$$\frac{\dot{MAC}}{MAC} = \tilde{r} - \frac{\gamma\zeta^2 S}{\varphi_t a A^{-\chi}} - \frac{\chi a}{2A_0 A}. \tag{13}$$

Note that $\tilde{r} = r + \frac{\dot{c}}{c} - \frac{\dot{U}}{U}$ is the standard Ramsey consumption discount rate, resulting in the Hotelling rule. Next, applying cost-benefit, rather than cost-effectiveness, results in the second term, which flattens the price path, as shown in Dietz and Venmans (2019). The last term, which is the consequence of endogenous learning, is also negative. Therefore, it also flattens the carbon price path.

Thus, the growth rate of the social cost of carbon is unaffected by exogenous learning and reduced by endogenous learning. Furthermore, consider a model where learning is endogenous and compare it with a model that is identical, with the same MAC function at each point in time, but learning is exogenous: the carbon price and abatement will be higher at all points in time in the model with endogenous learning.

In fact, consider equation 11 at the initial time; the endogenous model has an extra positive term in the integral. By assumption, the factor $\varphi_t A^{-\chi}$ is the same for both models (identical MAC function at each point in time). As a result, abatement is larger in the endogenous model. We will confirm this numerically in Section 6.

3.2 Abatement path

Taking the derivative with respect to time of the marginal abatement cost (in utils) expressed in equation 7 (note $\dot{A} = a/A_0$), we have

$$U(1-\eta)\left(\varphi_t \dot{a} A^{-\chi} + \dot{\varphi}_t a A^{-\chi} - \frac{\chi \varphi_t a^2 A^{-\chi-1}}{A_0}\right) + \dot{U}(1-\eta)\left(\varphi_t a A^{-\chi}\right) = \dot{\lambda^S} \tag{14}$$

The second term $\dot{\varphi}_t a A^{-\chi}$ quantifies the reduction of the marginal abatement cost due to exogenous learning. Similarly, the third term $-\frac{\chi \varphi_t a^2 A^{-\chi-1}}{A_0}$ is of the same nature and has the same sign as the effect of the exogenous learning. Both can be described as a "decreasing cost effect on abatement". This effect tells us that for a given carbon price, less abatement is obtained today compared to the future, leading to a steeper abatement path.

Dividing by $U(1-\eta)$ and using equation 8 gives us a formula for the growth rate of abatement (the Euler equation).

$$\frac{\dot{a}}{a} = r - g - \frac{\dot{U}}{U} - \frac{\dot{\varphi}}{\varphi} + \frac{\chi a}{A_0 A} - \frac{1}{2}\frac{\chi a}{A_0 A} - \frac{\gamma \zeta^2 S}{\varphi a A^{-\chi}} \tag{15}$$

The first three terms correspond to the growth-adjusted discount rate. Next, the positive term $-\frac{\dot{\varphi}}{\varphi}$ represents the "decreasing cost effect on abatement" of exogenous learning. It steepens the abatement trajectory, abate less at the start, but more in the future for a given steady state temperature. The next term $+\frac{\chi a}{A_0 A}$ is the "decreasing cost effect on abatement" from endogenous learning, which dominates the "early learning incentive" ($-\frac{1}{2}\frac{\chi a}{A_0 A}$). Therefore, endogenous learning also steepens the abatement path. The last term $-\frac{\gamma \zeta^2 S}{\varphi a A^{-\chi}}$ represents increasing marginal damage costs and is added in the case of cost-benefit analysis and absent in cost-effectiveness analysis. It shows that taking the timing of the damages into account creates an incentive for a flatter abatement path (earlier abatement). The effect is larger in later periods when S approaches the steady state. Learning magnifies this term by decreasing the denominator. Note however that learning will also have an indirect effect on cumulative emissions.

Many bottom up models model endogenous learning as exogenous learning. In other words, they include the "decreasing cost effect on abatement" but exclude the "early learning incentive". They also disregard the timing of the damages by doing a cost-effectiveness analysis. Equation 13 shows that both approximations have the same effect on the growth rate of the abatement path: the path is too steep (too little abatement in the short run).

4 Methodology to isolate the exogenous and endogenous learning dynamics effects

In this section, we develop a model which allows us to isolate the dynamic effects of learning. We make marginal abatement costs identical at each point in time, such that the differences in the optimal emissions paths are only affected by the dynamic incentives of learning.

4.1 Isolating the exogenous effect

We start with a model where there is only exogenous learning. The marginal abatement cost, expressed as % of consumption is

$$MAC\%(a,t) = a\varphi^{exo}(t) \tag{16}$$

where $\varphi^{exo}(t)$ is time dependent. Let's the resulting optimal path be $M\hat{A}C, \hat{A}, \hat{a}$. Next, we fit a ninth-degree polynomial in abatement to the marginal abatement cost of the exogenous learning curve.

The new marginal abatement cost, expressed as % of consumption is then

$$M\hat{A}C\% = \left(\sum_{n=0}^{9} \varphi_n^{nolearn} \hat{a}^n\right) \tag{17}$$

Note that we take the assumption that the first derivative of U with respect to a is the same (see equation 26 in Appendix). If the optimal exogenous learning abatement path were to be followed, this would give an abatement cost function that is identical at any point in time. Hence, any deviation from the exogenous learning path is only due to different incentive structure between no learning and exogenous learning.

4.2 Isolating the endogenous effect

We start with a model where there is only endogenous learning. The marginal abatement cost, expressed as % of consumption is

$$MAC\%(a, A) = \varphi^{endo} A^{-\chi} a \tag{18}$$

where φ^{endo} is a constant. Let's the resulting optimal path be $M\hat{A}C, \hat{A}, \hat{a}$. Next, we fit a twentieth-degree polynomial in time to the learning factor of the optimal endogenous learning path $\varphi^{endo} \hat{A}_t^{-\chi} = \sum_{n=0}^{9} \varphi_n^{exo} t^n$. The new marginal abatement cost, expressed as % of consumption is then

$$M\hat{A}C\% = \left(\sum_{n=0}^{9} \varphi_n^{exo} t^n\right) \hat{a} \tag{19}$$

If the optimal endogenous learning abatement path were to be followed, this would give an abatement cost function that is identical at any point in time. Hence, any deviation is only due to different incentive structure between endogenous and exogenous learning.

5 Calibration of our model

For the calibration, we use a slightly more complicated model, which has also a speed penalty on abatement (see Appendix). We define $v = \dot{a}$ and we use $\frac{C}{C_{BAU_0}} = e^{\left(gt - \frac{\varphi_t}{2} a^2 A^{-\chi} - \frac{\theta_2}{2} v^2 - \frac{\gamma}{2} \zeta^2 S^2\right)}$. Abatement a is defined as $60GtCO2eq - Emissions$. For exogenous learning we use the functional form $\varphi_t = \varphi_\infty + (\varphi_0 - \varphi_\infty) e^{-g_\varphi t}$ in the main analysis. We fit both the total abatement cost and marginal abatement cost functions using maximum likelihood, giving equal weight to both errors and assuming that they are normally distributed. Table 1 provide the parameter estimates for fitting both total abatement costs and marginal abatement costs to the climate scenarios database of IPCC and NGFS.

As to the other parameters $(\delta, n, \eta, g, \zeta, \gamma)$, we used the values provided in Table 2. To estimate the carbon price in our numerical simulations, we use the following value for the world GDP of 2020: 84537 billions U.S. dollars (International Monetary Fund, 2021).

6 Results

For the numerical simulations, we solve the model which includes inertia[2] (see Appendix). The model, which consists of a system of 4 differential equations in 4 variables (S, a, v, λ^S), is solved as a boundary value problem with MATLAB's bvp5c function. The boundary conditions are also presented in Appendix (Equation 30).

The parameters estimates presented in Section 5 are used in our four central cases: the exogenous, endogenous and two "No learning" cases. The endogenous learning case uses the parameters of the "EndoLearn3" and the exogenous learning case corresponds to "ExoLearn3" (see Table 1). In order to improve comparability, a same value for θ_2 is used in all cases as well as an identical φ_0 at the start, except for the first "No learning" case[3]. This particular case corresponds to the "NoLearn2" fit and includes the fitted φ_0 value. However, the second "No learning2" scenario is a more traditional reference case which has the same value for φ_0 as the other cases.

Stemming from this significant difference in the starting φ_0 values, the two "no learning" cases differ tremendously, resulting in a temperature difference of 0.28 degrees at the end of the century (see Fig. 1). When the MAC curve is calibrated on current abatement costs and learning is ignored (as in No learning2), the optimal carbon price is overestimated. As to the learning cases, the slope of the MAC is higher in the endogenous case until 2090 but we see that emissions on the other hand are always lower, which result in a slightly lower temperature at the end of the century (compared to the exogenous case). Thus, the incentive stemming from learning-by-doing (which is absent in the exogenous formulation) leads to more efforts and a lower temperature increase. The difference is not huge but still noticeable. We will isolate the effect of the endogenous learning dynamic more accurately later in this section.

A sensitivity analysis is performed over the learning parameters $(\chi, A_0, g_\varphi, \varphi_\infty)$. The standard deviations from the fit are used as ranges for this analysis, except for A_0[4]. We can conclude from Fig. 2 that there is less abatement in the short term when considering learning (exogenous or endogenous), compared to the first No learning case (all the 18 learning cases of the sensitivity analysis have

[2]Note that we use Equation 29 to get the exact solution (not the approximation of Equation 30).

[3]The θ_2 value used corresponds to the average of the fitted θ_2 values of the endogenous and exogenous cases (0.00176). The same approach applies to the common starting φ_0 which is also computed as the average of the fitted φ_0 values of the endogenous and exogenous cases (0.0000492).

[4]As to A_0, the standard deviation is too large to be used as the range for the sensitivity analysis. Therefore, we consider the following values: $0.5A_0, A_0, 1.5A_0$.

Variable	NoLearn1	NoLearn2	ExoLearn	ExoLearn2	ExoLearn3	EndoLearn	EndoLearn2	EndoLearn3
φ_0	.00003528	.0000346	.00004573	.00006168	.00005106	.00004981	.00006153	.00004736
	1.490e-06	1.479e-06	2.486e-06	4.866e-06	4.772e-06	4.276e-06	4.597e-06	3.644e-06
θ_2		set at .00176			.00175115			.00178073
					.00027976			.00030203
g_φ			.00494108	.057958	.04818181			
			.00099184	.01044132	.01435699			
φ_∞				.00003369	.00003347			
				1.684e-06	1.728e-06			
χ						.18203806	.14672636	.10954442
						.0483043	.03044451	.04521848
A_0						set at 300	37.315085	100.6696
							12.373234	134.74255
N	1850	1850	1850	1850	1848	1850	1850	1848
ll	7094.4002	7120.1211	7111.7791	7636.4378	7639.8787	7112.3337	7113.744	7121.0399
bic	-14181.278	-14232.719	-14208.512	-15250.307	-15249.67	-14209.622	-14204.919	-14211.992
aic	-14186.8	-14238.242	-14219.558	-15266.876	-15271.757	-14220.667	-14221.488	-14234.08

Table 1: Parameter estimates for fitting both total abatement costs and marginal abatement costs to the climate scenarios database of IPCC and NGFS. NoLearn, ExoLearn and EndoLearn respectively represent models without learning, with exogenous learning and endogenous learning. NoLearn2 has a speed penalty $\theta_2 = 0.00176$ to mimic the learning models with speed penalty. Exolearn has a learning formula $\varphi_t = \varphi_0 e^{-g_\varphi t}$converging to zero abatement costs in the very long run. Exolearn2 and 3 have $\varphi_t = \varphi_\infty + (\varphi_0 - \varphi_\infty)\, e^{-g_\varphi t}$. Endolearn has an exogenously set $A_{Cum_0} = 300$ whereas in Endolearn2 and 3 this parameter is estimated endogenously. All models "3" have an endogenously estimated speed penalty θ_2.

Parameter	Value	Source
$\delta - n$	0.011-0.005	Drupp et al. (2018); United Nations (2017)
η	1.35	Drupp et al. (2018)
g	0.02	By assumption
ζ	0.0006	Calculation on basis of IPCC (2021)(see Section 3)
γ	0.0154	By assumption

Table 2: Assumptions concerning the parameters values for the numerical simulations

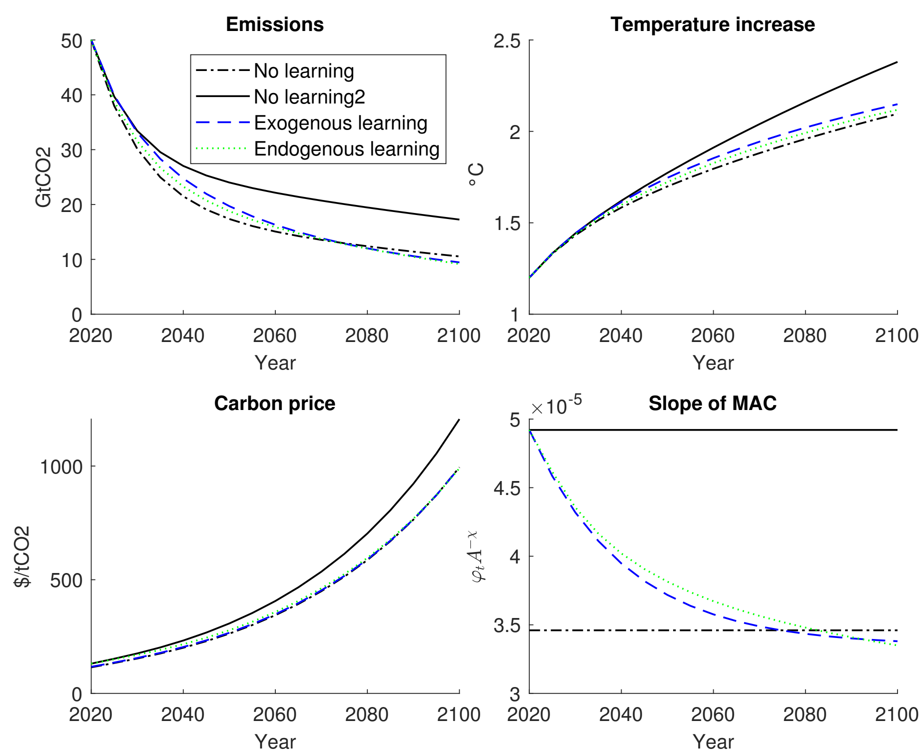

Figure 1: Central cases

a lower abatement until 2050). After 2050, more and more learning scenarios start to abate a greater amount of emissions than the No learning case. Our conclusions from Section 3.2 are verified numerically: both endogenous and exogenous learning lead to a steeper abatement path. This is due to the fact that both dynamics have a decreasing cost effect on abatement (in the equation 15 of the growth rate of abatement (the Euler equation)). As to the endogenous case, this decreasing cost effect on abatement dominates the early learning incentive.

Under most conditions (16 out of the 18 learning cases), the temperature is higher in 2100 in the scenarios including learning (than in the "No learning" case). Furthermore, note that, compared to the "No learning 2" case (which has a higher slope of MAC then No learning), all the learning cases have always a higher abatement (because of the "bad" calibration of this second reference scenario on current costs instead of on future costs as in the first No learning scenario).

In the long term (Fig. 2), we show that most of the learning cases (13/18 cases) lead to a lower peak warming, compared to No learning. The central "No learning" cases reach 3.93 and 2.87 degrees in 2500 while the learning cases are more ambitious, especially the endogenous cases (the central endogenous learning case reach 2.34 degrees). Thus, compared to a model without learning, including exogenous or endogenous learning steepens the abatement path and leads to less emissions reduction in the short term, but thanks to more efforts in the medium/long term, peak warming is lower. Moreover, slight negative emissions occur in the endogenous case (from 2270).

Next, we wish to isolate the learning effect on the optimal trajectory. Firstly, we compare the exogenous learning path with a model where there is no learning, but identical marginal abatement costs, by fitting a polynomial (as explained in Section 4). As shown in Fig. 3, there are no visible difference between the 2 paths, hence apparently the exogenous learning dynamic has no effect on the optimal trajectory. In fact, the difference in emissions is only of 0.009 GtCO2 (0.04%) in 2050 for instance. Thus, we show that as long as the model is well calibrated, including an exogenous learning feature which is dependent on time does not change the results, compared to a model where the cost of abatement depends only on the level of abatement.

Secondly, we isolate the endogenous learning dynamic effect. As explained in more details in Section 4, we fit a polynomial which is dependent on time to the learning factor of the optimal endogenous learning path and hence get a model with exogenous learning (which we compare to the endogenous learning case). As opposed to the previous exercise showing the absence of impact of the exogenous learning dynamic, Fig. 4 shows the significant impact of the endogenous learning dynamic on the optimal path. Including the endogenous learning dynamic leads to less emissions through the whole period until 2100 and hence less warming at the end of the century. For instance, there is a difference of 1.88 GtCO2 in 2050 (9.11%) and of 0.56 GtCO2 in 2100 (5.78%) while the temperature is 3.01% higher in the exogenous case in 2100 (2.184° compared to 2.118°).

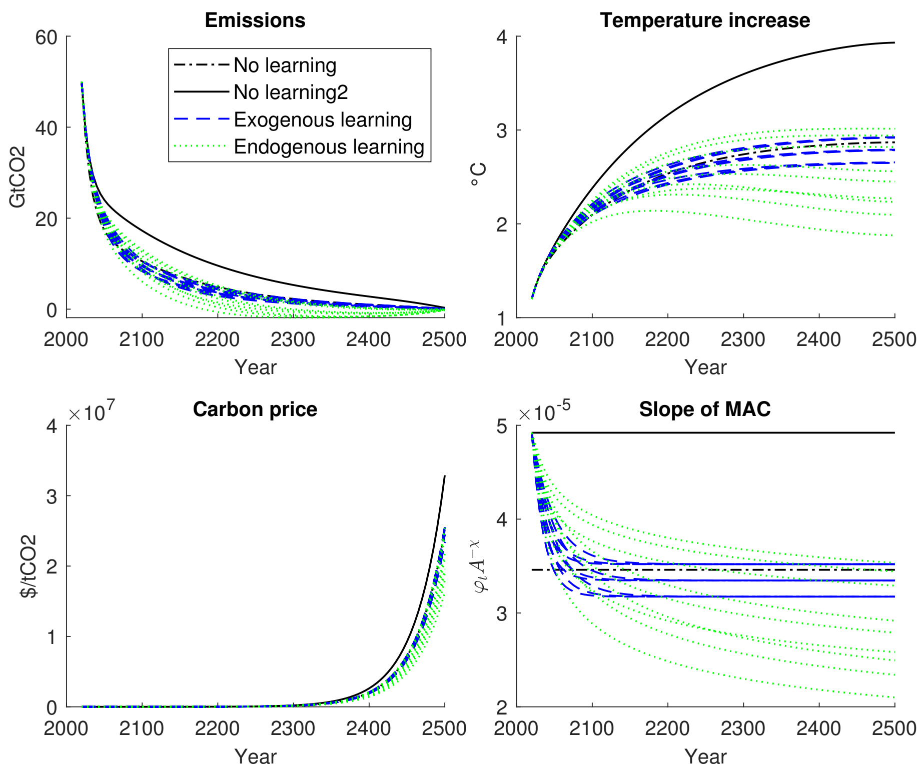

Figure 2: Sensitivity analysis on the learning parameters until 2500

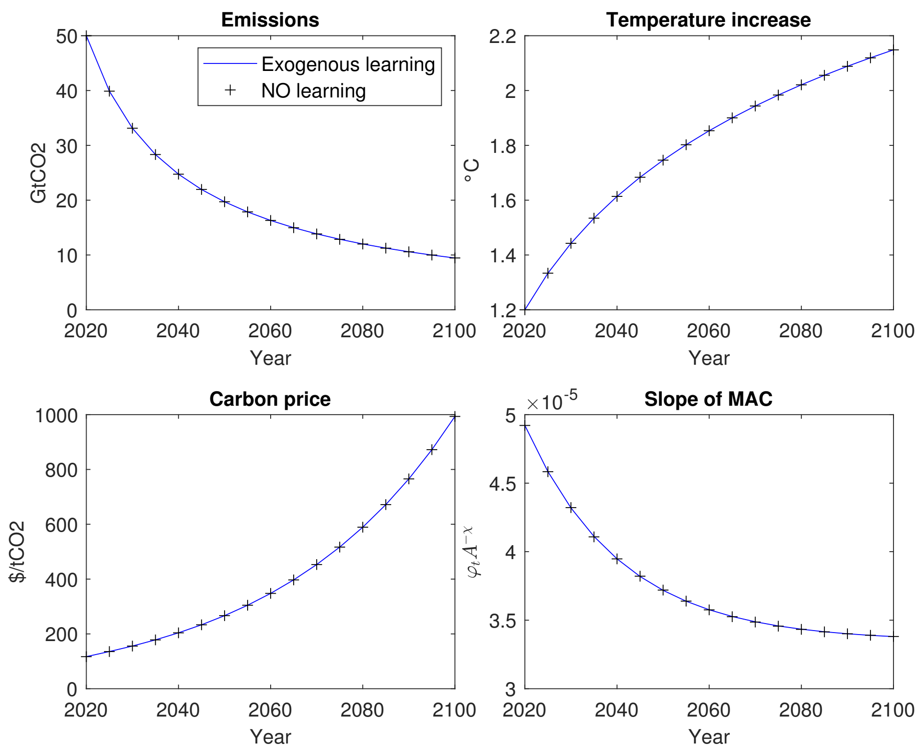

Figure 3: Exogenous learning versus no learning: fitting a polynomial in abatement to the marginal abatement cost of the exogenous learning curve

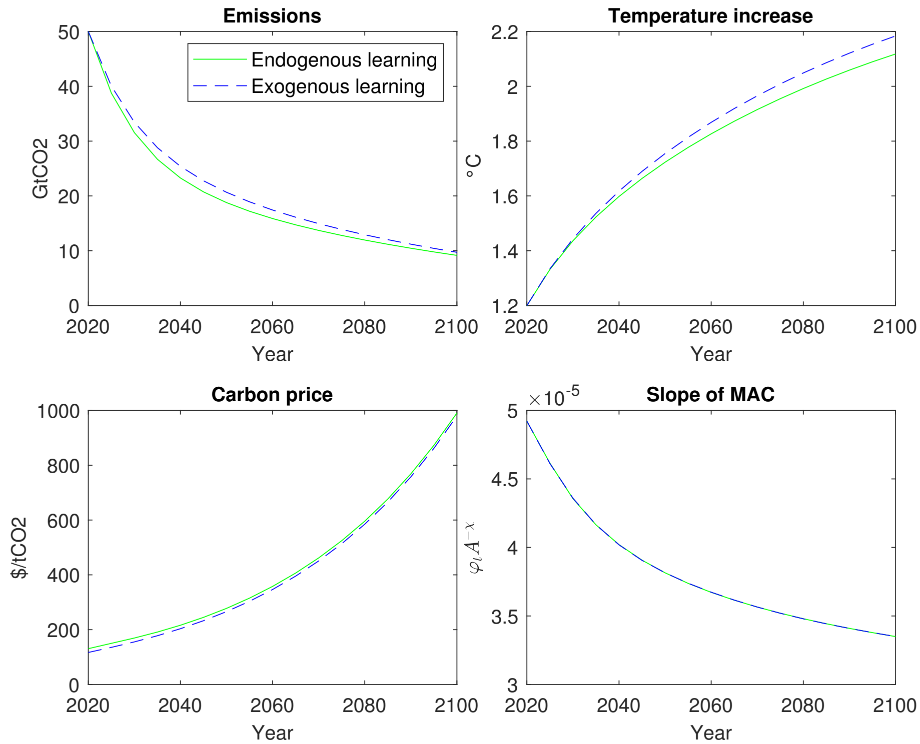

Figure 4: Endogenous learning versus exogenous learning: fitting a polynomial in time to the learning factor of the optimal endogenous learning curve

Note that if we consider a lower discount rate[5], the differences become more significant: e.g. about a 20% emissions difference in 2050 and a temperature in 2100 of 1.667° in the endogenous case and of 1.746° in the exogenous case. This finding indicates that models that do not include a proper endogenous learning dynamic (which is the case of many IAMs and bottom-up energy models) underestimate the optimal abatement throughout the century. As a consequence, policy makers and modellers should keep in mind that trajectories coming from models without endogenous learning might not be ambitious enough. In fact, they do not consider the "learning gains" which come from the fact that reducing emissions leads to less costly abatement in the future. As to the carbon price, it is always higher in the endogenous case in this century after 2020. We have confirmed our analytical findings from Section 3.1 : the carbon price and the abatement are higher in the endogenous case, compared to the exogenous case, because of the supplementary term in the optimality condition (the social cost of carbon is not only the sum of all discounted marginal damages, it also includes the sum of all future gains from endogenous learning).

[5]We consider $\delta = 0.006$ instead of $\delta = 0.011$.

7 Conclusion

Optimal emissions trajectories produced by integrated assessment and bottom-up energy models are of essential need to inform policies. In this paper, we sorted out the question of the differentiated impact of endogenous and exogenous learning on the optimal path. It was highly relevant to do so given the facts that learning dynamics matter when looking at the optimal path and that many widely used models do not include endogenous learning.

Theoretically, endogenous learning modifies the optimality condition: the social cost of carbon is not only the sum of all discounted marginal damages anymore, it also includes the sum of all future gains from endogenous learning. This leads to more abatement and a higher carbon price. Moreover, looking at the equation of the growth rate of abatement (the Euler equation), we saw that both exogenous and endogenous learning lead to a decreasing cost effect on abatement. As to the endogenous case, this decreasing cost effect on abatement dominates the early learning incentive. In fact, both learning dynamics lead to a steeper abatement path. We confirmed those findings numerically in our central cases. Furthermore, we found out that slight negative emissions occur in the endogenous case (from 2270).

In order to provide a fair and accurate comparison of the different types of learning, we first calibrated our model on the IPCC and NGFS database. Second, we developed a methodology to isolate respectively the exogenous and the endogenous effect thanks to a polynomial fit (which allow the models compared to have exactly the same slope of MAC).

This last methodology leads to a more accurate comparison. Contrary to our central cases, we compared models which have the same "costs" and only differ in terms of learning incentives. Regarding the impact of exogenous learning on the optimal trajectory, this methodology led to a different conclusion than the one we explained before (i.e., a steeper path). We showed that as long as the model is well calibrated, the isolated exogenous learning dynamic (which is time dependent) has absolutely no effect on the optimal path, compared to a model where the cost of abatement depends only on the level of abatement (and not on time). We believe that this fact has not been highlighted before and is of interest to modellers: as long as the cost calibration is accurate, including exogenous learning does not impact the results.

On the contrary, the endogenous learning dynamic has a significant impact on the optimal trajectory: it leads to less emissions throughout the entire period until 2100 and hence less warming at the end of the century. Depending on the discount rate, the impact on emissions in 2050 can be up to a 20% difference (9% in our central case) between an endogenous model and an exogenous model with exactly the same slope of MAC. As to the temperature in 2100, the endogenous learning leads to a warming of 2.118° (1.667° with a low discount rate), which is 3% (5%) lower than in the exogenous case. Thus, the common practice of modelling endogenous learning as an exogenous process underestimates the optimal abatement, leading to a higher warming at the end of the century. This is a strong conclusion for policy makers and modellers: one should keep in mind that trajectories coming from models without endogenous learning might not be ambitious enough.

References

[1] Acemoglu, D., Aghion, P., Bursztyn, L. & Hemous, D. The environment and directed technical change. *American economic review* **102**, 131–66 (2012).

[2] Dietz, S. & Venmans, F. Cumulative carbon emissions and economic policy: In search of general principles. *Journal of Environmental Economics and Management* **96**, 108–129 (2019). URL `https://www.sciencedirect.com/science/article/pii/S0095069618302122`.

[3] Drupp, M. A., Freeman, M. C., Groom, B. & Nesje, F. Discounting disentangled. *American Economic Journal: Economic Policy* **10**, 109–34 (2018). URL `https://www.aeaweb.org/articles?id=10.1257/pol.20160240`.

[4] Fischer, C. & Newell, R. G. Environmental and technology policies for climate mitigation. *Journal of Environmental Economics and Management* **55**, 142–162 (2008). URL `https://www.sciencedirect.com/science/article/pii/S0095069607001064`.

[5] Gerlagh, R., Kverndokk, S. & Rosendahl, K. E. Optimal timing of climate change policy: Interaction between carbon taxes and innovation externalities. *Environmental and resource Economics* **43**, 369–390 (2009).

[6] Gillingham, K., Newell, R. G. & Pizer, W. A. Modeling endogenous technological change for climate policy analysis. *Energy Economics* **30**, 2734–2753 (2008). URL `https://www.sciencedirect.com/science/article/pii/S0140988308000443`. Technological Change and the Environment.

[7] Huppmann, D. et al. IAMC 1.5C Scenario Explorer and Data hosted by IIASA. *Integrated Assessment Modeling Consortium and International Institute for Applied Systems Analysis* (2019). URL `data.ene.iiasa.ac.at/iamc-1.5c-explorer`.

[8] International Monetary Fund. World economic outlook database (2021).

[9] IPCC. Climate change 2021: The physical science basis. contribution of working group i to the sixth assessment report of the intergovernmental panel on climate change. *[Masson-Delmotte, V., P. Zhai, A. Pirani, S.L. Connors, C. Pean, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews, T.K. Maycock, T. Waterfield, O. Yelekci, R. Yu, and B. Zhou (eds.)]. Cambridge University Press. In Press* (2021).

[10] Goulder, L. H. & Mathai, K. Optimal co2 abatement in the presence of induced technological change. *Journal of Environmental Economics and Management* **39**, 1–38 (2000). URL https://www.sciencedirect.com/science/article/pii/S0095069699910896.

[11] Grubb, M. *et al.* Taking time seriously: Implications for optimal climate policy (2020).

[12] Grubb, M., Wieners, C. & Yang, P. Modeling myths: On dice and dynamic realism in integrated assessment models of climate change mitigation. *Wiley Interdisciplinary Reviews: Climate Change* **12**, e698 (2021).

[13] Manne, A. & Richels, R. The impact of learning-by-doing on the timing and costs of co2 abatement. *Energy Economics* **26**, 603–619 (2004). URL https://www.sciencedirect.com/science/article/pii/S0140988304000441. EMF 19 Alternative technology strategies for climate change policy.

[14] NGFS. Climate scenarios database (2021). URL https://www.ngfs.net/ngfs-scenarios-portal/.

[15] Popp, D. Entice: endogenous technological change in the dice model of global warming. *Journal of Environmental Economics and Management* **48**, 742–768 (2004). URL https://www.sciencedirect.com/science/article/pii/S0095069603001220.

[16] Rogelj, J. et al. Mitigation pathways compatible with 1.5c in the context of sustainable development, in "special report on global warming of 1.5c". *Intergovernmental Panel on Climate Change, Geneva* (2018). URL http://www.ipcc.ch/report/sr15/.

[17] United Nations. World population prospects: The 2017 revision, key findings and advance tables. *Technical Report Working Paper No. ESA/P/WP/248* (2017).

[18] Wright, T. P. Factors affecting the cost of airplanes. *Journal of the aeronautical sciences* **3**, 122–128 (1936).

Appendix: model with inertia

Assume a penalty for abatement speed (as a % of GDP) of $\theta_1 v + \frac{\theta_2}{2} v^2$. It represents increasing repurposing/stranding assets costs, bottlenecks in fast green R&D as well as macro-economic, labour market and financial adjustment costs. Note that we set $\theta_1 = 0$ in our calibration.

The welfare functional writes

$$max \int_0^\infty e^{-(\delta-n)t} \frac{c^{(1-\eta)}}{1-\eta} \tag{20}$$

$$max\int_0^{\infty} e^{-(\delta-n+(\eta-1)g)t}\frac{c_{BAU_0}^{1-\eta}}{1-\eta}e^{(1-\eta)\left(-\frac{\varphi_t}{2}a^2A^{-\chi}-\theta_1 v-\frac{\theta_2}{2}v^2-\frac{\gamma}{2}\zeta^2S^2\right)} \tag{21}$$

Define the abstract concept $U = \frac{c_{BAU_0}^{1-\eta}}{1-\eta}e^{(1-\eta)\left(-\frac{\varphi_t}{2}a^2A^{-\chi}-\theta_1 v-\frac{\theta_2}{2}v^2-\frac{\gamma}{2}\zeta^2S^2\right)}$ and $r=\delta-n+\eta g$

The current value Hamiltonian (with sign switch to obtain a positive shadow value) is

$$H^{CV} = U - \lambda^S\left(E_{BAU}-a\right)+\lambda^a v \tag{22}$$

The FOC are

$$-U_v = \lambda^a \tag{23}$$

$$\dot{\lambda^S} = (r-g)\,\lambda^S + U_S \tag{24}$$

$$\dot{\lambda^a} = (r-g)\,\lambda^a - U_a - \lambda^S \tag{25}$$

With the 2 equations of motion, we have a system of 5 equations with 5 unknowns. In order to obtain a system of only 4 differential equations, we will substitute equation 23 in 25.

Derivatives that will be useful:

$$\begin{aligned}
U_S &= U\,(1-\eta)\left(-\gamma\zeta^2 S - \frac{\chi\varphi_t}{2A_0}a^2A^{-\chi-1}\right)\\
U_a &= U\,(1-\eta)\left(-\varphi_t a A^{-\chi}\right)\\
U_v &= U\,(1-\eta)\left(-\theta_1-\theta_2 v\right)
\end{aligned} \tag{26}$$

Time derivative of $U_v(S,a,v,t)$ writes [6]

$$\begin{aligned}
-\dot{U}_v(S,a,v,t) &= -\left(U_{vS}E + U_{va}v + U_{vv}\dot{v} + U_{vt}\right)\\
&= U\,(1-\eta)\,\theta_2\dot{v}+ \qquad (27)\\
&+ U\,(1-\eta)^2\left(\theta_1+\theta_2 v\right)\\
&\left\{\left(-\gamma\zeta^2 S - \frac{\chi\varphi_t}{2A_0}a^2A^{-\chi-1}\right)\left(E_{BAU}-a\right)+\left(-\varphi_t a A^{-\chi}\right)v \qquad (28)\right.\\
&\left.+\left(-\theta_1-\theta_2 v\right)\dot{v}+\frac{\varphi_t}{2}a^2A^{-\chi}\left(\frac{\chi E_{BAU}}{A_0 A}+g_\varphi\right)\right\}
\end{aligned}$$

Inserting equations 23, 26 and 28 in equation 25, we obtain

[6] E_{BAU} will cancel out, unless it changes over time. The 3 last lines are negligible, they correspond to $\dot{U}(1-\eta)(\theta_1+\theta_2 v)$.

$$\dot{v} = \frac{1}{\theta_2 - (1-\eta)(\theta_1 + \theta_2 v)^2} \Bigg\{ (\theta_1 + \theta_2 v)(r - g) + \varphi a A^{-\chi} - \frac{\lambda^S}{U(1-\eta)}$$
$$- (1-\eta)(\theta_1 + \theta_2 v) \left[-\gamma \zeta^2 S (E_{BAU} - a) + \varphi a A^{-\chi} \left(\frac{\chi}{2 A_0 A} a^2 - v + \frac{1}{2} a g_\varphi \right) \right] \Bigg\} \tag{29}$$

For a very good approximate solution, write
$-\dot{U}_v = U(1-\eta)\theta_2 \dot{v} + \dot{U}(1-\eta)(\theta_1 + \theta_2 v)$.
This results in

$$\dot{v} = \frac{1}{\theta_2} \left[\left((r - g) - \frac{\dot{U}}{U} \right) (\theta_1 + \theta_2 v) + \varphi a A^{-\chi} - \frac{\lambda^S}{(1-\eta) U} \right] \tag{30}$$

$\frac{\dot{U}}{U} = (1-\eta) \frac{d}{dt} \left(-\frac{\varphi_t}{2} a^2 A^{-\chi} - \theta_1 v - \frac{\theta_2}{2} v^2 - \frac{\gamma}{2} \zeta^2 S^2 \right)$ is the effect of abatement and climate damages on the growth rate. At the start, abatement reduces the growth rate, whereas after a few decades, damages reduce the growth rate. The effect is fairly constant (Dietz and Venmans, 2019), around 0.02%.

We now have a system of 4 differential equations in 4 variables S, a, v, λ^S, which can be solved as a boundary value problem using bvp5c in MATLAB.

The boundary conditions are

$$\begin{aligned} S(0) &= S_0 \\ a(0) &= a_0 = E_{BAU} - E_0 \\ a(\infty) &= E_{BAU} \\ v(\infty) &= 0 \end{aligned} \tag{31}$$

We find the optimal price using

$$p = \frac{\lambda^S}{c^{-\eta}} = \frac{\lambda^S}{U(1-\eta)} c \cong c \left[\varphi_t a A^{-\chi} + (r - g)(\theta_1 + \theta_2 v) - \theta_2 \dot{v} \right] \tag{32}$$

(omitting the term with $\dot{U}$).