No News is Good News in Climate Change?

Reyer Gerlagh and Matti Liski*

April 9, 2013

Abstract

We develop a tractable climate-economy model for evaluating a price for carbon when climate change impacts are non-existing – they may arrive later. We identify the general-equilibrium value of learning (or not learning) the economic losses from impacts. The value justifies increasing carbon prices due to increasing vulnerability to impacts when the global economy expands, although the assessment of climate change becomes more optimistic. The model produces a belief distribution for the social cost of carbon that can be calibrated to match a comprehensive survey of previous estimates. Even if climate change progresses without impacts for the coming century, the carbon price continues to increase as suggested by the "climate policy ramp" of existing evaluations. The quantitative assessment justifies climate policies for a rational skeptic.

(JEL classification: H43; H41; D61; D91; Q54; E21. Keywords: carbon tax, learning, climate change)

1 Introduction

"Estimating impacts has been the most difficult part of all climate science"
—William W.D. Nordhaus, EAERE lecture 2012

Climate change and the resulting global warming is a complex phenomenon involving more uncertainties than perhaps any other environmental problem. Yet, there is relatively

*Gerlagh <r.gerlagh@uvt.nl> is at the economics department of the Tilburg University. Liski <matti.liski@aalto.fi> is at the economics department of the Aalto University, Helsinki. We thank participants at the CESifo area conference "Energy and Climate Economics" 2012 for helpful comments and discussions, and James Corbishley for diligent research assistance.

good hard evidence on the fundamentals of climate change, including such processes as the diffusion of carbon dioxide between atmosphere, oceans and biosphere (e.g, Hoos et al. 2001). The warming itself is unequivocal to the scientists as well as the role of anthropogenic emissions in the process (see, e.g., IPCC Fourth Assessment Report, 2007). However, similar hard evidence on the socio-economic impacts that follow from the ongoing climate change is lacking. In contrast with data on greenhouse gas concentrations in the atmosphere or on the global mean land-ocean temperature, we have little or no quantitative information on how changes in the climate will impact our economies, although there is extensive research on what such impacts might be; see Tol (2009) for a comprehensive survey on methods and results.

Most evaluations of the social cost of carbon build on middle-of-the-road assumptions on climate change impacts, commonly expressed in terms of GDP losses, presumed to be real and existing at the point in time where the climate policies are designed. Climate-economy models such as DICE (Nordhaus, 2007) combine the impact scenarios with interdisciplinary climate-research inputs to obtain a monetized value for the social cost of releasing one unit of carbon emissions to the atmosphere — there is a pressing demand for such a number as it is required, for example, in the cost-benefit analyses to assess regulations across wide domains (Greenstone et al., 2011). Clearly, the carbon price evaluations depend critically on the realism of the impact assessments; that is, on the part of climate science we know very little about. In this paper, we start with the observation that climate change impacts are *de facto* not experienced currently and, to address this problem, we develop an approach for evaluating the carbon price even without observable impacts.

To evaluate the carbon price when impacts are absent at the time of policy-making, we develop a novel and tractable climate-economy model that, in contrast with large-scale simulation models, allows a transparent quantitative assessment of the carbon price under the gradual learning of impacts. The resulting carbon price formula is detailed enough for reproducing the baseline predictions of the comprehensive climate-economy models — it forms a solid basis for testing the sensitivity of current carbon price evaluations to the fact that impacts will be revealed only gradually over time. The result is the first tractable general-equilibrium carbon price formula that can quantify the effect of learning on the climate policies. We build on Golosov et al. (2011) for the general-equilibrium climate-economy interactions, and on Gerlagh and Liski (2012) for a detailed but tractable description of the global carbon cycle for the dependence of climate change on the anthropogenic emissions. The approach incorporates the idea that the fundamentals of climate change are relatively well understood but it remains, however, uncertain if climate change will ever have a significant impact on our economies. For illustration, the Greenland ice-sheet has not yet melted, and even if the process of warming of the climate is well understood, little is known about the consequences of a 2-degree warming — whether it will lead to such melting and what the social-economic impacts of such a possible event will be.

To device a conservative test against existing policy proposals, we assume that impacts of climate change are currently non-existing. There is uncertainty about the time at which impacts will start to materialize at such levels that they provide hard evidence. That is, our approach describes uncertainty for the arrival date of impacts, as a stochastic process separate from the uncertainty with respect to the magnitude of damages. We consider two basic possible states of the world: a benign scenario in which no (or low) climate damages occur, and a more pessimistic scenario in which (high) climate damages will come about at some future date. The subjective probability of each state depends on the observations up to that date. There are infinitely many possible scenarios, including one for no arrival of impacts. When climate change progresses and no very large climate change impacts occur, beliefs are updated: it becomes more likely that impacts of a given temperature increase will never become very large, and thus "no news is good news".

Is no news then good news for climate policy makers? That is, if the climate impacts remain low or non-existing, should climate policies then gradually loosen over time? Our learning dynamics are biased towards supporting this outcome and yet the result is opposite: the monetized value for the carbon price in 2015 is even somewhat higher than for a sure-thing best-guess (median) estimate for damages. Moreover, the future development of the carbon price over the next 100 years closely follows the gradual tightening of policies, advocated as the "climate policy ramp" by Nordhaus (2007).

That the gradual tightening of climate policies makes sense even if climate optimism increases over time, follows from the connection between optimal carbon prices, beliefs and the general-equilibrium. The model generates a parametric distribution function for possible carbon prices which we match with existing estimates for the distribution, as presented by Tol (2008, 2009) — this distribution sets the initial carbon price level and a subjective probability, as held by the profession, for both the severity of climate damages and the possible time when they occur. Learning over time then involves updating of the beliefs of the "climate policy experts" as captured by our distribution — the calibrated parameters imply that the learning rates are low, indicating that in climate change optimism comes slowly, even if no damages occur. This persistence of beliefs is one part of the explanation for our climate policy ramp in the absence of impacts.

The remaining part of the explanation follows from the general-equilibrium setting. Building on Golosov et al. (2011) and Gerlagh and Liski (2012), the optimal general-equilibrium carbon price is proportional to the income level that in most scenarios is expected to increase dramatically, due to the rise of the middle class in major emerging economies. Our case, following the IPCC (2000) business-as-usual scenario, implies global per capita incomes that increase tree-fold during the coming century. Then, even if climate policy makers become more optimistic over time, the severity of potential impacts increases due to the expanding scale of the potential losses. This increasing economic vulnerability together with low rates of learning explains the "climate policy ramp".[1]

Our approach is different from earlier contributions on learning in climate change. The literature modeling the learning of climate impacts has focused on the structural uncertainties of the climate system, including those related to the climate sensitivity (Kelly and Kolstad 1999 and Leach 2007) or unknown thresholds leading to tipping points (Lemoine and Traeger 2010). Our research departs from this literature in two main ways. First, our target is to make progress in a field previously dominated by simulation models by developing analytically tractable and thus transparent policies; however, the model is rich in details to capture the main dynamics of the global carbon cycle. We follow Golosov et al. (2011) in modeling the economy but Gerlagh and Liski (2012) for the details of the carbon cycle and also emission abatement options.[2]

Second, our model of learning is about economic impacts only: it is taken as given that the climate change as a phenomenon is understood, but the economic losses from a given climate change are not known. This allows us to keep the climate description equivalent to that, for example, in DICE, and thereby focus on the learning of economic uncertainties. The model of learning is simple enough to allow for closed-form evaluations

[1]The supporting potential climate event that justifies the initial carbon price consistent with Nordhaus (2007) is equivalent to a GDP loss of about 15 per cent at temperatures that are $3°C$ above pre-industrial level. Such an event is clearly catastrophic, but not a "tail event", in the sense of Weitzman (2009) where policies become undefined since, effectively, it is not possible to transfer wealth to the high consequence events. In our case, the optimal carbon price prior to the possible bad news is effectively implementing a consumption smoothing path across the expected outcomes, and therefore we are in the domain of "weak tail dominance" where policies ultimately converge (Nordhaus, 2010), although this converges is extremely slow.

[2]We cite both Gerlagh and Liski (2012) and (2013); the former is a longer working paper version that contains a detailed description of the energy sector that is needed in the quantitative analysis of the current paper.

and sharp results — we assume that the economy starts without observations for the climate damages and then future observations are considered experiments with a potential *persistent* arrival of the damage when temperatures increase. Despite the simplicity, the learning captures the evolution of the initial assessment that the climate change will ultimately lead to damages, as well as the general-equilibrium implications of the change over time in assessment. In the concluding section, we show that the results can arise also in a more comprehensive learning model capturing gradual learning from, for example, extreme weather events.

The paper is structured as follows. In Section 2, we first explain the emissions-temperature response that follows from the detailed description of the global carbon cycle. The response is needed for obtaining the full externality cost of current emissions. The calibration of the carbon cycle comes from Gerlagh and Liski (2012). In Section 3, we then introduce the consumption choice model that justifies the derived externality cost in general equilibrium; this model is effectively Brock-Mirrman (1972) adapted to climate change. In Section 4, we introduce the learning dynamics that maintain the analytic nature of optimal policies. Section 5 introduces the calibration and the quantitative assessment.

2 Emissions–temperature response

We assume that the physical relationships in the climate system are known with certainty – only the impacts will be uncertain. At the heart of these relationships is the emissions-temperature response that captures the delay with which current emissions cause future changes in temperatures. To introduce this response, let D_t be a measure of the global mean temperature increase above the pre-industrial levels t. Current emissions, denoted by z_t, affect temperatures at time $t + \tau$ according to a known function $\mathcal{R}(\tau)$:

$$\frac{dD_{t+\tau}}{dz_t} = \mathcal{R}(\tau) > 0. \tag{1}$$

Function $\mathcal{R}(\tau)$ captures the delays in the temperature response; when calibrated with physical data describing the carbon cycle, $\mathcal{R}(\tau)$ has a non-linear shape with a peak about 60-70 years after the date of the emissions, and it also has a fat right tail. In Fig. 1, we illustrate $\mathcal{R}(\tau)$ as implied by our calibration based on Gerlagh and Liski (2013), and also as implied by the Norhaus' DICE model (2007). The Figure shows the life-path of the temperature responses following from one unit of emissions at time $t = 0$, relative to the counterfactual where the temperature adjusts immediately.[3]

Gerlagh and Liski (2013) derive a closed form for $\mathcal{R}(\tau)$ as follows. There is a linear relationship between concentrations and D_t in the climate equilibrium (steady state), captured by parameter π: a one-unit increase in the steady-state atmospheric CO_2 stock leads to a π-unit increase in the steady-state level of D_t. Outside steady state, there is a delay in the effect from concentrations to temperatures, and this delay is captured by parameter $0 < \varepsilon < 1$: a one-unit increase in emissions increases the next period concentrations one-to-one but temperatures only $\varepsilon\pi$ -units.

The non-linearity of the temperature response, as depicted in Fig. 1, follows from the break-down of the CO_2 stock to reservoirs such as atmosphere, oceans and biosphere that have differing decay rates for the CO_2 that they contain. Such a reservoir system can be described by a system of "atmospheric boxes" (Maier-Reimer and Hasselman 1987). Let $\mathcal{I}$ denote the set of boxes, with share $0 < a_i < 0$ of annual emissions entering each box $i \in \mathcal{I}$, and $0 \leqslant \eta_i < 1$ its carbon depreciation factor. A three-box representation will be sufficiently rich to capture the time-scale of the climate response to emissions.

This description leads to a closed-form for an emissions-damage response (see Gerlagh and Liski (2013) the derivation): the impact of emissions at time t on temperatures at time $t + \tau$ is

$$\frac{dD_{t+\tau}}{dz_t} = \mathcal{R}(\tau) = \sum\nolimits_{i \in \mathcal{I}} a_i \pi \varepsilon \frac{(1-\eta_i)^\tau - (1-\varepsilon)^\tau}{\varepsilon - \eta_i} > 0, \tag{2}$$

where the geometric terms $(1 - \eta_i)^\tau$ and $(1 - \varepsilon)^\tau$ characterize the delays in carbon concentration and temperature adjustments, respectively.

Parameter η_i captures, for example, the carbon uptake from the atmosphere by forests and other biomass, and oceans. The term $(1-\eta_i)^\tau$ measures how much of carbon z_t is in box i after τ periods, and the term $-(1-\varepsilon)^\tau$ captures the slow temperature adjustment. The limiting cases can be helpful. Consider one CO_2 box, so that the summation over the boxes can be ignored. If atmospheric carbon-dioxide does not depreciate at all, $\eta = 0$, then the temperature slowly converges at speed ε to the long-run equilibrium climate sensitivity π, giving $\theta_\tau = \pi[1-(1-\varepsilon)^\tau]$. If atmospheric carbon-dioxide depreciates fully, $\eta = 1$, the temperature immediately adjusts to $\pi\varepsilon$, and then slowly converges to zero, $\theta_\tau = \pi\varepsilon(1-\varepsilon)^{\tau-1}$. If temperature adjustment is immediate, $\varepsilon = 1$, then the temperature

[3] The counterfactual response assumes that (i) the emitted CO2 remains in the atmosphere forever and (ii) the emitted CO2 has immediate full temperature effects. The deviation from the DICE temperature response is explained below, after the introduction of the carbon cycle representation.

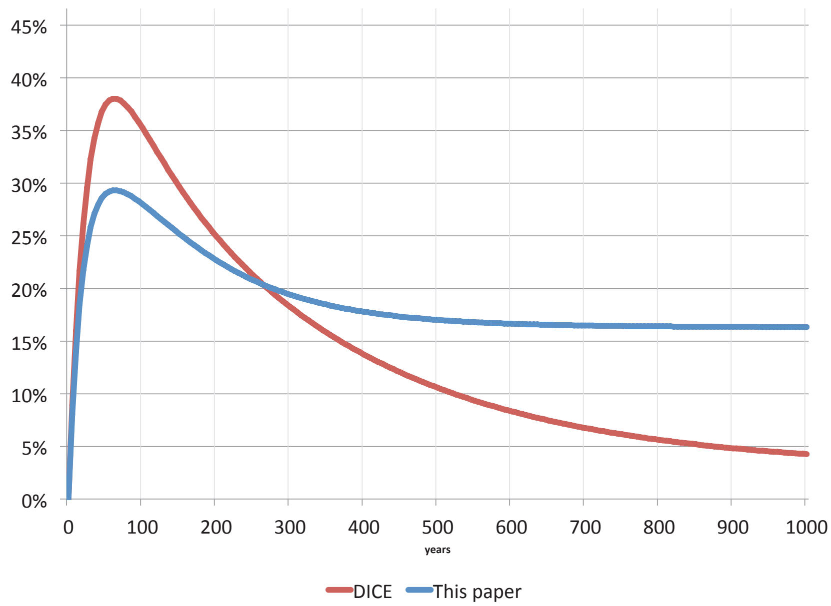

Figure 1: Emissions-temperature responses. The paths are presented relative to the counterfactual response that assumes that (i) the emitted CO2 remains in the atmosphere forever and (ii) the emitted CO2 has immediate full temperature effects. Responses shown for the current paper and DICE 2007 where the parameters of our model are set to match to DICE carbon cycle representation.

response function directly follows the carbon-dioxide depreciation $\theta_\tau = \pi(1-\eta)^{\tau-1}$. If temperature adjustment is absent, $\varepsilon = 0$, there is no response, $\theta_\tau = 0$.

The physical data on carbon emissions, stocks in various reservoirs, and the observed concentration developments are used to calibrate a 3-box carbon cycle representation leading to the following emission shares and depreciation factors per decade:[4]

$$
\begin{aligned}
a &= (.163, .184, .449) \\
\eta &= (0, .074, .470).
\end{aligned}
$$

Thus, about 16 per cent of carbon emissions does not depreciate while about 45 per cent has a half-time of one decade. Following Nordhaus (2001), we set $\pi = .0156$ [per $TtC0_2$, see Gerlagh and Liski (2013)].[5] We assume $\varepsilon = .183$ per decade, implying a

[4]Some fraction of emissions enters the ocean and biomass within a decade, so the shares a_i do not sum to unity.

[5]This choice is obtained by interpreting D directly as a variable causing damages and assuming, as

global temperature adjustment speed of 2 per cent per year. These quantitative choices parametrize the emissions-temperature response that is depicted in Figure 1. [6]

3 General-equilibrium carbon price: no uncertainty

The advantage of our $\mathcal{R}(\tau)$ is that it allows obtaining the explicit social cost of carbon, while maintaining a description of the climate system that is consistent with comprehensive climate models — this description proves useful also when the damages from temperature increases are not known with certainty. Let Δ_t denote the loss, measured in utils, per unit of temperature increase D_t at time t . Below we introduce uncertainty and learning regarding Δ_t but, for the time being, assume that $\Delta_t = \Delta > 0$ is a constant.

To obtain the carbon price, that is, the social cost of current carbon emissions z_t, consider the effect of emissions at t on period $t+\tau$ utility:

$$-\frac{du_{t+\tau}}{dz_t} = \Delta \frac{dD_{t+\tau}}{dz_t} = \Delta \mathcal{R}(\tau).$$

The full loss of utils per increase in temperatures as measured by $D_{t+\tau}$ is thus a constant given by Δ for any future $t+\tau$, giving the social cost of carbon emissions z_t at time t, appropriately discounted to t, as

$$\begin{aligned}
-\sum_{\tau=1}^{\infty} \delta^{\tau} \frac{du_{t+\tau}}{dz_t} &= \Delta \sum_{\tau=1}^{\infty} \delta^{\tau} \frac{dD_{t+\tau}}{dz_t} \\
&= \Delta \sum_{i\in\mathcal{I}} \frac{\beta a_i \pi \varepsilon}{\varepsilon - \eta_i} \sum_{\tau=1}^{\infty} \delta^{\tau}(1-\eta_i)^{\tau} - \delta^{\tau}(1-\varepsilon_j)^{\tau} \\
&= \Delta \sum_{i\in\mathcal{I}} \frac{\delta \pi a_i \varepsilon}{[1-\delta(1-\eta_i)][1-\delta(1-\varepsilon)]} \\
&= h. \qquad (3)
\end{aligned}$$

Value h compresses the present-value utility costs of current emissions to a number; it will help to identify the currently optimal carbon price. We can justify this externality-cost calculation in a general-equilibrium by introducing a consumption choice model following

in Nordhaus (2001), that doubling the steady state CO_2 stock leads to 2.6 per cent output loss. For this reason, the interpretation of π is "climate damage sensitivity" rather than "climate sensitivity". The Appendix of Gerlagh and Liski (2013) explicates these steps.

[6] In Figure 1, the main reason for the deviation from DICE is that DICE assumes an almost full CO_2 storage capacity for the deep oceans, while large-scale ocean circulation models point to a reduced deep-ocean overturning running parallel with climate change (Maier-Reimer and Hasselman 1987). The scientific literature suggests that the CO2 millennial depreciation in the DICE carbon cycle is too optimistic (Archer and Brovkin, 2008).

Brock and Mirman (1972), and by introducing economic losses from climate change as in Golosov et al. (2011).[7] Assume that the utility function is logarithmic in consumption c_t, and linear in intangible damages associated with climate change:[8]

$$u_t = \ln(c_t) - \Delta_u D_t, \tag{4}$$

where the potential direct utility losses per temperature increase are denoted by $\Delta_u \geqslant 0$.

Let k_t denote total capital that is assumed to depreciate in one period, say, in a decade. Combining log-utility for consumption with a production function where capital contribution takes the Cobb-Douglas form, with $0 < \alpha < 1$,

$$y_t = k_t^{\alpha} A_t \exp(-\Delta_y D_t) \tag{5}$$

leads to the workhorse model in analytical macroeconomics for the consumption choices over time. The contribution of fossil-fuels, that is, carbon inputs z_t, as well as labor l_t enter through the function $A_t = A(t, l_t, z_t)$ that captures the details of the energy sector of the economy. The income losses due to climate change arise as reduced output, and depend on the history of emissions through the state variable D_t, defined through the emissions-temperature response $\mathcal{R}(\tau)$. The multiplier, $\Delta_y > 0$, is a constant damage coefficient.

A convenient feature of the structure (4)-(5) is that the optimal policies become free of the energy-sector details: the current policy for savings and emissions can be determined when the current state of the economy is known.[9] Here we merely assume that the output is differentiable, increasing and strictly concave in emissions z_t. The structure introduced

[7] The assumption of GDP losses together with the assumption that abatement is not a separate choice but only related to the use of carbon inputs, is an important departure from the Nordhaus DICE modeling tradition. Golosov et al. (2011) develop this "real business cycle model of climate change". Our analysis follows Gerlagh and Liski (2013) where we develop the carbon cycle with realistic delays for this model, and also introduce a detailed energy sector that allows calibration of productivities such that the model can produce the IPCC basic scenarios as the benchmark outcome.

[8] Golosov et al. (2011) do not have direct utility losses from temperature increase; including them changes essentially nothing in the analysis.

[9] The details of the energy sector will affect the future development of the economy and thus the future states of the economy and future policies. For this reason, we will introduce more structure to the energy sector in Section 5.2.

by (4)-(5) implies that the optimal allocations $\{c_t, k_t, z_t\}_{t\geqslant 0}$ that maximize[10]

$$\sum\nolimits_{t=0}^{\infty} \delta^t u_t$$

will be characterized by policies where constant share $0 < g = \alpha\delta < 1$ of the gross output is invested,

$$k_{t+1} = g y_t, \tag{6}$$

and where the climate policy is defined through

$$\frac{\partial y_t}{\partial z_t} = h(1-g) y_t. \tag{7}$$

Note that since $\frac{\partial u}{\partial c} = 1/c = 1/(1-g)y$, expression (7) implies that the utility-weighted carbon price $\frac{\partial y_t}{\partial z_t}\frac{\partial u}{\partial c} = h$ remains constant over time. Intuitively, the marginal output gain from increasing carbon use at time t, having consumption value $\frac{\partial u}{\partial c}$, should be equalized with the future utility-cost of current emissions, given by h. Using our previous expression (3) for h, we can express the optimal carbon price, denoted by $\tau_t = \frac{\partial y_t}{\partial z_t}$, as

$$\tau_t = (1-g) y_t \Delta \sum\nolimits_{i\in\mathcal{I}} \frac{\delta \pi a_i \varepsilon}{[1-\delta(1-\eta_i)][1-\delta(1-\varepsilon)]} \tag{8}$$

where

$$\Delta = \Delta_u + \frac{\Delta_y}{1-g} \tag{9}$$

This description of the optimal policy now gives a structural interpretation for the loss in utils, Δ, that we used to derive the full utility cost of current emissions in (3): it incorporates the direct utility loss Δ_u as well as the current and future output losses, captured by $\Delta_y/(1-g)$, from a current temperature increase.[11]

Proposition 1 *(Gerlagh and Liski, 2013). Given the economy described by emissions-temperature response (2), preferences (4), and technologies (5), the optimal carbon price (8) is proportional to income, with proportionality depending only on δ, Δ, and the carbon cycle parameters in (2). Given loss Δ, the same tax is optimal for any division between utility and production losses satisfying (9).*

[10] Note that we do not scale the objective with labor to avoid having time trends in the policy variables. This implies no loss of generality; the solution for the case where the total, rather than average, utility is maximized is available on request.

[11] The latter term takes this form since a current drop in output propagates to subsequent periods through savings. See Gerlagh and Liski (2013) for the explicit derivation.

We build on this general-equilibrium model in the analysis of learning. Learning of climate impacts will change the climate policy variable over time, $h = h_t$, so that through (7) there will be learning-induced changes in the carbon price, $\tau_t = (1 - g)y_t h_t$.

Throughout the quantitative analysis, we assume 10-year periods; the first year is '2010' corresponding to period 2006-2015. We assume only output losses from climate change and thus set $\Delta_u = 0$, to maintain an easy comparison to earlier studies. We take the Gross Global Product as 600 Trillion Euro [$Teuro$] for the decade, 2006-2015 (World Bank, using PPP). Throughout we assume a capital share of $\alpha = .3$ and one per cent pure rate of annual time preference, implying $\delta = .90$ for decadal periods and resulting in savings $g = .27$.

Normalizing the output loss parameter at unity, $\Delta = 1$, gives us the Nordhaus (2007) baseline where a temperature rise of 3 Kelvin leads to about 2.7 per cent loss of output.[12] This, together with our carbon cycle, results in a carbon price of 22 EUR/tCO_2, equivalent to about 105 USD/tC, for 2010. This estimate is higher than the Nordhaus baseline (2007) because of our lower pure rate of time preference; this choice facilitates the calibration in Section 5.1.[13]

4 Carbon price learning

Our approach to climate change uncertainty is dichotomous: the climate change as described through the emissions-temperature response $\mathcal{R}(\tau)$ is assumed to be perfectly known but the impacts, as captured by Δ, are non-existing and it remains uncertain whether Δ will ever turn positive. This approximates the asymmetry of the availability of evidence on the physical facts of the climate change, and, on the other hand, on the socio-economic impacts.

How should the learning of damages be conceptualized when observations or signals regarding the severity of the problem caused by temperature increases are lacking? Here

[12] To clarify the units, the damages are measured per Teraton of CO2 [TtonCO2], and the 3-Kelvin rise follows from doubling the CO2 stock. We have chosen the value of π such that the normalization $\Delta = 1$ gives the Nordhaus case.

[13] Note that 1 tCO2 = 3.67 tC, and 1 Euro is about 1.3 USD. Our number 105 USD/tC is almost precisely equal to the DICE-2007 carbon price when the elasticity of substitution parameter and the pure rate of time preference are both chosen to take value 1, as in our analytical model. The number appearing in Nordhaus (2007), that is $35USD/tC$, can be matched by setting 2.7 per cent pure rate of time preference. However, Tol's data, used in the calibration below, does not exist for this value of time preference.

we develop a simple model for learning that preserves the tractability of the carbon prices formulas introduced in the previous section. We want to model learning as a phenomenon where potentially very little hard evidence on damages accumulates when the climate change progresses.[14]

4.1 Beliefs

The economy can be in two states, $I_t \in \{0, 1\}$. If $I_t = 0$, no damages have been experienced by t. If $I_t = 1$, damages have appeared, and once $I_t = 1$, then $I_{t+\tau} = 1$ for all $\tau \geq 0$. The damage at time t, in utils, is $\Delta_t = \Delta I_t$ per unit of temperature measure D_t. Thus, once damages appear, the policies can be determined exactly as in Proposition 1. With no experience of damages prior to or at t, the experienced damage is zero. Note that if $I_t = 1$, then the total damage experienced depends on state of the physical system, that is, on the total temperature increase reached at time t: the damage equals ΔD_t at t and $\Delta D_{t+\tau}$ for all subsequent periods $\tau > 0$. Since D_t is extremely persistent in past emissions (see Fig. 1), there is an implied irreversibility of impacts that contributes to the cautiousness of policies.

The hazard rate for damages, denoted as p, is the probability that damages start and $I_t = 0$ moves to $I_{t+1} = 1$. The hazard rate is a given constant but unknown to the policy maker. p has has a discrete prior distribution: it can either take value $p = 0$ or $p = \lambda$. The hazard rate can depend on the degree of climate change as measured by D_t, that is, we can allow that only for periods where $D_t > \overline{D} \geqslant 0$, the state can switch. We consider this extension in Section 4.3, and assume now learning in all periods by setting $\overline{D} = 0$.[15]

Since there are no prior climate experiments, we do not know the value of p. We assume that there is a subjective prior probability $\mu_0 > 0$ for a positive hazard rate, $p = \lambda$, which implies the probability for eventual climate impacts:

$$\begin{aligned} 1 - \mu_0 &= \Pr(\lim_{t\to\infty} I_t = 0) = \Pr(p = 0) \\ \mu_0 &= \Pr(\lim_{t\to\infty} I_t = 1) = \Pr(p = \lambda > 0). \end{aligned}$$

Let μ_t denote the posterior probability that $p = \lambda$ conditional on no learning by time

[14] Alternatively, gradual learning can follow, for example, through extreme whether events; we explain the extension to this direction in the concluding section.

[15] For example, $\overline{D}$ can correspond to 2-degrees Celsius warming, but since we have little information about the learning thresholds, we will set $\overline{D} = 0$ in the calibration. Also, the solution of the model can be easily extended to the case of different temperature brackets having different hazard rates.

t. Each period where $D_t > \overline{D} = 0$, but where no damages have appeared so far, $I_t = 0$, climate change runs an experiment. If the outcome is $I_{t+1} = 1$, which happens with probability $\mu_t \lambda > 0$, we have learned that $p = \lambda$, so $\mu_{t+1} = 1$. If the outcome is $I_{t+1} = 0$, we have not learned the state of nature with certainty, but we update the beliefs μ_{t+1}. We can then write the Bayesian updating rule as[16]

$$\begin{aligned} \mu_t &= \Pr(p = \lambda \,|I_t = 0\,) \\ &= \frac{\mu_0(1-\lambda)^t}{\mu_0(1-\lambda)^t + 1 - \mu_0}. \end{aligned} \tag{10}$$

which is the probability that climate change damages will ultimately arrive even though such damages have not been experienced by time t. Note that μ_t declines over time: "no news is good news". The assessment of the distribution for damages becomes more optimistic over time.[17]

4.2 Carbon price distribution

The model generates a distribution of damages. Let Z be a stochastic variable, measuring the full future utility cost from increasing current emissions z_t marginally: Z can take the values $Z_1, Z_2, ...$, where Z_τ is the current social cost of carbon if damages appear for the first time, precisely at period $t + \tau$. Thus, Z_τ characterizes the present-value marginal utility damages of current emissions z_t, assuming that the damage indicator I_t remains at zero for all periods prior to $t + \tau$ but then turns positive. Proceeding as in Section 3, and using the emissions-temperature response from Section 2, we can obtain the present-value of such delayed damages in closed-form:

$$\begin{aligned} Z_\tau &= \Delta \sum\nolimits_{t=\tau}^{\infty} \delta^\tau \mathcal{R}(\tau) \\ &= \Delta \sum\nolimits_{(i,j)} \frac{\delta \pi a_i b_j \varepsilon_j}{\varepsilon_j - \eta_i} \delta^\tau \left(\frac{(1-\eta_i)^\tau}{[1-\delta(1-\eta_i)]} - \frac{(1-\varepsilon_j)^\tau}{[1-\delta(1-\varepsilon_j)]} \right). \end{aligned}$$

[16] Note that $\Pr(p = \lambda \,|I_t = 0\,) \times \Pr(I_t = 0) = \Pr(p = \lambda \cap I_t = 0)$. The probability that there has been no news by time t is $\Pr(I_t = 0) = \mu_0(1-\lambda)^t + 1 - \mu_0$. The probability that there has been no news by time t and that $p = \lambda$ is $\Pr(p = \lambda \cap I_t = 0) = \mu_0(1-\lambda)^t$. Combining gives the equation.

[17] One could argue that impacts must ultimately arrive for a sufficiently severe climate change. While the model can be extended to include temperature brackets where impacts arrive almost surely, it is also reasonable to think that, for example, a long period of 2-degrees warming without impacts is evidence for not having impacts at such temperatures. Even if one considers "no news is good news" learning to be biased, this bias is consistent with the idea of having a conservative test against the climate policy ramp, as explained in the Introduction.

Given our model of learning, we find for the distribution of Z that

$$\Pr(Z = Z_\tau | I_t = 0) = \Pr(I_\tau = 1 \cap I_{\tau-1} = 0 | I_t = 0)$$

which gives the probability that damages turn positive exactly after τ periods when the current time t subjective belief for the climate problem is μ_t. To find the corresponding cumulative distribution function for the social cost of carbon, denoted by $F_t(Z)$, we first establish the probability that the damage has revealed itself at period t, irrespective of if t it is the first time:

$$\begin{aligned}
\Pr(I_t &= 1) = (1-\mu_0)\Pr(I_t = 1|p = 0) + \mu_0 \Pr(I_t = 1|p = \lambda) \\
&= \mu_0[1 - \Pr(I_t = 0|p = \lambda)] \\
&= \mu_0[1 - \Pr(I_1 = ... = I_t = 0|p = \lambda)] \\
&= \mu_0[1 - (1-\lambda)^t]
\end{aligned}$$

We can generalize this to expectations at period t,

$$\Pr(I_{t+\tau} = 1 | I_t = 0) = \mu_t[1 - (1-\lambda)^\tau]$$

so that the distribution for the carbon price is then given by

$$\begin{aligned}
F_t(Z) &= \Pr(Z \leq Z_\tau | I_t = 0) = \Pr(I_{t+\tau-1} = 0 | I_t = 0) \\
&= 1 - \mu_t + \mu_t(1-\lambda)^{\tau-1}.
\end{aligned}$$

We can use this distribution to determine the social cost of carbon at time t as dependent on beliefs μ_t.

Proposition 2 *Conditional on no experience of impacts by time t ($I_t = 0$), the previous-period distribution of the optimal carbon price per consumption $F_{t-1}(Z)$ stochastically dominates the current distribution $F_t(Z)$. The social cost of carbon as measured by $h_t = \mathbb{E}_t Z$ declines over time conditional on $I_t = 0$. Moreover,*

$$\begin{aligned}
h_t &= \mathbb{E}_t Z = \textstyle\sum_{\tau=1}^{\infty} \delta^\tau \mathbb{E}_t \frac{du_{t+\tau}}{dz_t} \\
&= \mu_t \Delta \textstyle\sum_{i \in \mathcal{I}} \left\{ \frac{\delta \pi a_i \varepsilon}{[1-\delta(1-\eta_i)][1-\delta(1-\varepsilon)]} \right. \\
&\quad \left. - \frac{\delta(1-\lambda)\pi a_i b_j \varepsilon_j}{[1-\delta(1-\lambda)(1-\eta_i)][1-\delta(1-\lambda)(1-\varepsilon_j)]} \right\}.
\end{aligned}$$

Proof. It is straightforward to see that the expected change in damages associated with current emissions are equal to

$$\begin{aligned} h_t &= \mathbb{E}_t \Delta \textstyle\sum_{\tau=1}^{\infty} \delta^{\tau} I_{t+\tau} \dfrac{dD_{t+\tau}}{dz_t} \\ &= \Delta \textstyle\sum_{\tau=1}^{\infty} \delta^{\tau} \Pr(I_{t+\tau}=1|I_t=0)\mathcal{R}(\tau) \\ &= \mu_t \Delta [\textstyle\sum_{\tau=1}^{\infty} \delta^{\tau}\mathcal{R}(\tau) - \sum_{\tau=1}^{\infty}(1-\lambda)^{\tau}\delta^{\tau}\mathcal{R}(\tau)]. \end{aligned}$$

Using our temperature-response function leads to the expression for h_t. Decreasing carbon prices and stochastic dominance follow from μ_t decreasing over time. ∎

The result gives a closed-form expression for the optimal carbon price policy depending both on the climate system parameters and on the current belief of the damage distribution given by (μ_t, λ, Δ). Recall that, from (7), the optimal general-equilibrium carbon price is the income-weighted future utility-cost of current actions so that $\tau_t = (1-g)y_t h_t$, giving the learning-adjusted carbon price expression, analogous to (8), as

$$\begin{aligned} \tau_t = (1-g)y_t\mu_t\Delta \textstyle\sum_{i\in\mathcal{I}} \Big\{ & \dfrac{\delta\pi a_i \varepsilon}{[1-\delta(1-\eta_i)][1-\delta(1-\varepsilon)]} \\ & - \dfrac{\delta(1-\lambda)\pi a_i b_j \varepsilon_j}{[1-\delta(1-\lambda)(1-\eta_i)][1-\delta(1-\lambda)(1-\varepsilon_j)]} \Big\}. \end{aligned} \tag{11}$$

This result is the key to the observation that the "climate policy ramp", that is, the gradually tightening carbon price policy over time, can follow even with increasing climate optimism over time: in a growing economy, the income-weighted social cost of carbon can increase, although the optimism as captured by the declining h_t increases. The economy becomes more exposed to losses from climate change.

Limiting cases can be revealing. At time $t = 0$, where the subjective belief of damages is given by μ_0, if damages are almost surely observable, the optimal initial policy prior to experimentation is the full information policy weighted with the subjective probability for damages:

$$\lambda \approx 1 \Rightarrow h_0 \approx \mu_0 h$$

where h is defined in (3). However, if damages do not appear at time $t = 1$, then conditional on no news, the carbon policy is

$$\lambda \approx 1, I_1 = 0 \Rightarrow h_1 \approx 0$$

because the subjective assessment μ_1 drops to zero by the updating rule (10). The carbon policy variable drops from its maximum value to its minimum value in just one period, because no news reveals the true state of the world precisely when $\lambda \approx 1$. The effect is particularly dramatic if the subjective prior μ_0 assumes that climate is causing damages almost surely, $\mu_0 \approx 1$. On the other hand, for any given μ_0 at time $t = 0$, if climate change damages are hard to observe, then

$$\lambda \approx 0 \Rightarrow h_0 \approx 0$$

simply because the climate change is a problem with a non-significant rate of appearances in all cases. Note that in this case there will be no learning either: the subjective assessment μ_t in (10) almost does not change over time.

Outside the extreme values for (μ, λ) discussed so far, for intermediate values we find that the carbon policies as captured by h_t will initially be high but h_t will decline as t increases conditional on no news by t. This pattern of carbon pricing fully incorporates the understanding of future learning; the early prices reflect our initial beliefs. Thus, for each point in time in the future, we can consider two separate carbon price levels: one for the good news situation, and another for the situation where the damage has revealed itself at a prior point in time. These two carbon price paths diverge from each other over time for reasons that are obvious from the discussion above.

4.3 Learning thresholds

Before moving to the quantitative analysis, let us consider the situation where we expect learning only for climate change above a certain threshold level, covering situations where the degree of climate change determines the intensity of experimentation. Suppose learning takes place only above the temperature threshold, $D_t \geq \bar{D}$, corresponding, for example, to 1 or 2 degrees Celsius above the pre-industrial temperature levels.

Proposition 3 *Let h_0 be the carbon policy in Proposition 2 for initial belief μ_0, and assume that temperatures generate information on damages only if $D_t \geq \bar{D}$. Let h_t be the optimal policy path if no damages are observed until period t. Let $D_0 < \bar{D}$ and t' be the first period such that $D_{t'} \geqslant \bar{D}$. Then, prior to t', the carbon policy variable h_t, increases over time: $h_t < h_{t+1}$ for $0 < t < t'$.*

Proof. Let T be the periods t for which $D_t \geqslant \bar{D}$, if no damages are observed. It is straightforward to see that the expected damages associated with current emissions, in utility terms, satisfy

$$\begin{aligned}
h_t &= \Delta\mathbb{E}_t \sum_{\tau\in T} \Pr(I_\tau = 1|I_{\tau-1} = 0) \textstyle\sum_{s=\tau}^{\infty} \delta^s \dfrac{dD_{t+s}}{dz_t} \\
&= \Delta\mathbb{E}_t \sum_{\tau\in T} \Pr(I_\tau = 1|I_{\tau-1} = 0) \textstyle\sum_{s=\tau}^{\infty} \delta^s \mathcal{R}(s) \\
&< \Delta\mathbb{E}_t \sum_{\tau\in T} \Pr(I_\tau = 1|I_{\tau-1} = 0) \textstyle\sum_{s=\tau-1}^{\infty} \delta^s \mathcal{R}(s) \\
&= \Delta\mathbb{E}_{t+1} \sum_{\tau\in T} \Pr(I_\tau = 1|I_{\tau-1} = 0) \textstyle\sum_{s=\tau}^{\infty} \delta^{s-1} \mathcal{R}(s-1) \\
&= h_{t+1}
\end{aligned}$$

The first line follows as with certainty $I_t = 0$, for $0 < t < t'$. The fourth line follows as beliefs do not change between t and $t+1$. ■

As long as no information can be obtained, policy becomes more strict, in utility terms, over time until the temperatures start generating information. Note again that since the actual carbon tax is a multiple of income, the tax implied by h_0 for $D_t < \bar{D}$ will be growing over time at a rate exceeding the growth of the economy. Recall that our emissions-temperature response implies that the temperature peak for a given emissions impulse lags 60-70 years behind the date of emissions: the learning effects described here may start several decades after emissions have set in motion climate change. Meanwhile, optimal policies are characterized by constant beliefs, but by potentially sharply increasing carbon prices.

5 Quantitative assessment

5.1 Matching Tol

Richard Tol (2009) conducted a comprehensive survey of the existing estimates for the social cost of carbon. From the sample of 232 estimates he derived a distribution for the carbon price measured in 1995 USD/tC, for various time discount rates used in the studies. We match the carbon price distribution $F(Z)$ implied by our learning model to that in Tol to obtain estimates for the three parameters (μ, λ, Δ) above.

The exercise provides a basis for our illustration in that the initial carbon price level will be consistent with the views expressed by the experts; the updating of the beliefs is then given by the model interpretation. The distribution in Tol (2009) arises from multifaceted differences in the underlying original studies — arguably, the heterogeneity is driven by different subjective views on the fundamentals of the climate problem. In contrast, in our interpretation of the social cost distribution the heterogeneity is coming purely from different possible outcomes for the arrival date of the damage. This uncertainty is parametrized by μ_0 and λ; for given parameters, we can find the cumulative distribution for the arrival times, and then translate this into a distribution for the social cost of carbon, comparable with that in Tol.

Consider now Figure 2 which depicts the blue linear spline curve connecting the 33, 50, 67, 90 and 95 percentiles of carbon prices, expressed in 2010 EUR/tCO_2, as reported by Tol.[18] We focus on Tol's sample corresponding to 1 percent pure rate of time preference.[19] There is a mass point at zero, corresponding to 20 per cent of the assessment indicating insignificant or positive climate change impacts.[20]

We calibrate the climate system and economic parameters as in Section 3, and then fit our cumulative damage distribution function $F(Z)$ by choosing the initial prior μ_0, the hazard rate λ, the damage parameter Δ. We set $\overline{D} = 0$, assuming that any level of temperature increase produces information, that is, there is no delay in the information acquisition.

We can match our cumulative distribution with Tol's by choosing the triple (μ_0, λ, Δ). This can be achieved either by minimizing the errors at the reported percentile points, or, more directly, by matching the means and the end-points of the distributions. Both approaches are almost outcome-equivalent; we followed the latter approach to allow for the interpretation that the initial assessment of the carbon price equals exactly the assessment held by the profession. Moreover, to avoid giving too much weight to a few extreme cost estimates in the sample, we truncated the fitted distribution at 93 EUR/tCO_2 by setting $\Delta = 4.2$, that is, this is by factor four higher than the middle-of-the-road damage assumed in Nordhaus (2007) — the implied output loss is then about 10.7 per cent from doubling the CO_2 stock, if climate impacts materialize. To match the lower end of the distribution, we set $\mu_0 = .8$, meaning an initial 20 per cent assessment of no negative climate impacts. Finally, we choose λ to match Tol's mean value for the carbon price which is 32.7 for 2010 EUR/tCO_2 (his Table 2, 2009). We can obtain value $\lambda = .07$ such that the initial carbon price implied by our model exactly matches 32.7. The resulting cumulative distribution is depicted in Figure 2 as the green line.

We have forced one structural interpretation of the existing cost distribution to obtain strong conclusions for the subjective views on the likelihood of climate change damages.

[18] Note that 1 tCO2 = 3.67 tC, and 1 Euro in 2010 is about at parity with 1 USD in 1995.

[19] Tol reports distributions 0, 1 and 3 per cent discount rates. Our analysis of the 3-percent case produced very similar qualitative results; the levels of the policy variables systematically lower.

[20] This number we inferred from Tol (2008).

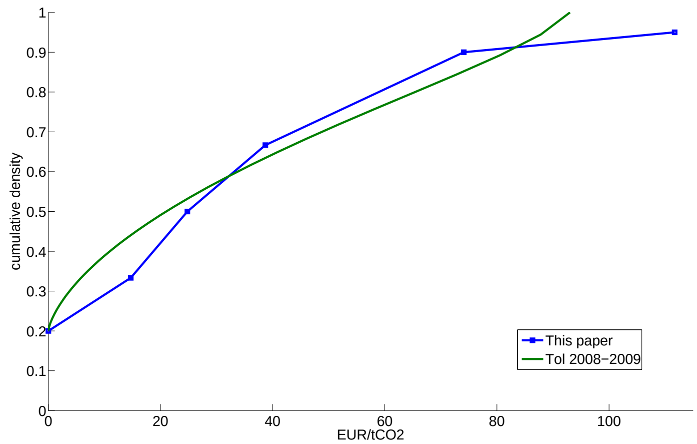

Figure 2: Fitting cumulative distribution with Tol's (2009) distribution.

First, the damage parameter, $\Delta = 4.2$, implies that a 3-Kelvin temperature increase leads to a 10.7 per cent output loss if the impacts arrive at such temperature levels. Second, $\lambda = .07$ means that information is generated very slowly – 7 per cent probability of learning per decade. A geometric distribution, $\lambda = .07$ per decade, means that the expected arrival time for a severe climate change damage event is about 140 years. After 100 years without damages, the posterior for the eventual impact arrival is still 65 per cent.

5.2 Climate-economy adjustment paths

By the above calibration procedure, we have set the optimal initial carbon price for 2010 at 32.7 EUR/tCO_2 but how should the carbon price develop over time? We consider now the optimal carbon policies using a climate-economy model from Gerlagh and Liski (2012) that determines the development of basic macro variables over time. The model has the structure introduced in Section 3 so that savings are given by a policy where fraction $g = \alpha\delta$ of output is saved in a period, and carbon policies are given by the carbon price per consumption h_t that we derived above. We calibrate savings to 25 per cent to maintain consistency with comparable climate-economy models; we set time discounting to 1 per cent, which leads to a relatively low capital share of output. Climate policy variable h_t follows from the carbon cycle calibration and from the parametrization of the learning process described above.

While the policy variables (g, h_t) can be separately solved and calibrated, the climate-economy adjustment paths will depend also on the details of the energy sector. Consider a production function as in (5) but further specified to

$$\begin{aligned} y_t &= k_t^{\alpha}[A_t(l_{y,t}, e_t)]^{1-\alpha} \exp(-\Delta_y D_t) \\ A_t(l_{y,t}, e_t) &= \min\{A_{y,t} l_{y,t}, A_{e,t} e_t\} \end{aligned}$$

where the overall labor-energy composite $A_t(l_{y,t}, e_t)$ takes a CES form with extreme low elasticity of substitution between labor in the final-good sector $l_{y,t}$ and energy e_t. That is, we consider a Leontief structure for the labor-energy composite.[21] Final good and energy productivities $A_{y,t}$ and $A_{e,t}$ are calibrated so that the model matches the business-as-usual (BAU) quantities with the A1F1 SRES scenario from the IPCC (2000). Energy e_t uses labor: the core allocation problem is how to allocate a given total labor l_t at time t between final output $l_{y,t}$, fossil-fuel energy, $l_{f,t}$, and non-carbon energy, $l_{n,t}$. Thus the energy and climate policy steers the labor allocation $(l_{y,t}, l_{f,t}, l_{n,t})_{t\geq 0}$ and thereby the quantities of fossil-fuel, $e_{f,t}$, and non-carbon energy, $e_{n,t}$. Both energy sources are intermediates, summing up to the total energy input:

$$e_t = e_{f,t} + e_{n,t}.$$

We assume that $e_{f,t}$ can be produced with constant-returns to scale technology using labor $l_{f,t}$ and the fossil-fuel z_t,

$$e_{f,t} = \min\{A_{f,t} l_{f,t}, B_t z_t\},$$

where $A_{f,t}$ and B_t describe productivities. The fuel resource is not a fixed factor and commands no resource rent; by this assumption, our focus is on the "coal phase", as coined by Golosov et al. (2011), where the fuel resource relevant for long-term climate policies is in principle unlimited. In contrast, the non-fossil fuel energy production is land-intensive and subject to diminishing returns and land rents (Fischer and Newell, 2008):

$$e_{n,t} = \frac{\varphi + 1}{\varphi}(A_{n,t} l_{n,t})^{\frac{\varphi}{\varphi+1}},$$

where $\varphi > 0$ describes the elasticity of supply from this sector, given the labor cost.

The model structure described here can reasonably well capture the two main adjustment channels to carbon policies: energy savings that typically feature the early decades of the adjustment, and then decarbonization that is needed to meet the long-run climate targets. The use of a Leontief aggregation for energy and final goods implies that we

[21] This assumption implies that early emissions are reduced through energy savings but long-run targets are achieved through decarbonization.

focus on the long-run decarbonization. All parameter choices are the same as in Gerlagh and Liski (2012), excluding those related to the modeling of learning of damages and the discount fact which we set to reflect 1 per cent annual discounting. Time periods are decades.[22]

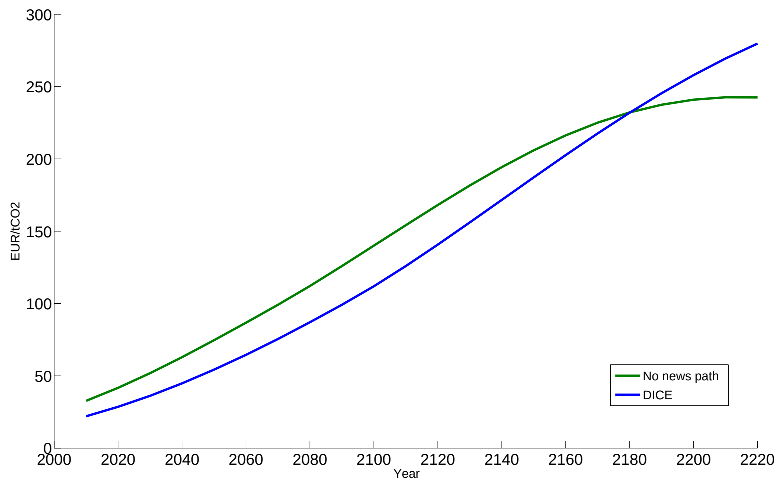

Figure 3: The optimal carbon price path conditional on no news on damages

5.2.1 "Climate policy ramp"

Let us now look at the optimal time path for the carbon price *conditional on no learning*, that is, we focus first on the evolution of the carbon pricing policy when no news on climate impacts arrive. Note that without impacts, the economy is unaffected by the climate change but, since the carbon policies are in place, emissions will be reduced below the business-as-usual path. The optimal carbon price path is depicted as the green path in Figure 3 over the coming century and beyond (we come to the other path shortly). Strikingly, for the learning parameters derived, it takes close to 200 years for beliefs to become optimistically enough to support the carbon price to start declining – despite the ultimate decline, the social cost of carbon virtually never dies out. It is practically impossible to have an affirmative assessment from not observing such an event that severe damages will not ultimately arrive.

For the shape of the carbon price path, recall that the optimal carbon price is proportional to income, $\tau_t = (1-g)y_t h_t$. In table 1, we decompose the effect the income,

[22] See the Appendices of Gerlagh and Liski (2012) for the temporal solution of the labor allocation as well as for the details of the calibration. See also the supplementary material of that paper.

$\tau_t = (1 - g)$, and the utility-cost, h_t, to the carbon price over time.

where the global income about doubles during the first 50 years, and then again more than doubles in the next 50 year period, after which the growth starts to level off. The scenario for the total global per-capita income growth is driven by the rise of the emerging economies, and also consistent with the IPCC baseline predictions.

Figure 3 also depicts the carbon price path that corresponds to the base calibration with certainty, based on Nordhaus' DICE (2007) middle-of-the-road damage, corresponding to $\Delta_y = 1$ – our model tracks very closely the DICE outcome when damages are assumed to exist from the start. For 2010, with 1 per cent annual discounting this path gives about 22 EUR/tCO_2 as the optimal price which exceeds Nordhaus' (2007) baseline policy; assuming 2.7 per cent time discount rate leads to his initial number almost precisely. The blue line captures the climate policy ramp, that is, the gradualism in tightening of the policies over the coming century and beyond. Our climate policy proposal, without actual experienced damages, has the same shape, and hence the main result of the paper follows: the policies should become tighter as the climate change progresses even if the impacts become more uncertain. The fact that our policy path has a higher level is a consequence of the calibration procedure; the mean value of estimates in the literature for the social cost of carbon lies above the Nordhaus' number. Assuming that the potential high-damages are about 8 per cent for a 3 Kelvin temperature increase, brings the optimal policy ramp close to the certainty policy, with a similar shape of the policy ramp for the coming decades.

	income	externality share	carbon price
2010	437	.075	33
2050	1065	.070	75
2100	2215	.063	140
2150	3719	.055	206
2200	5123	.047	241

Table 1: Decomposition of the carbon price to income (Trillion 2010 EUR) and to the externality share of income. Carbon prices in EUR/tCO_2 year 2010.

5.2.2 Virtual carbon price

It may seem surprising that carbon prices under learning, as depicted in Figs. 3, reach such high levels, despite no actual damage taking place. Obviously, in addition to the income growth development, the persistent tightness of the climate policy is supported by the possibility of real damages that may arrive at any time period. The virtual carbon price captures the economic meaning of the threat: it is the carbon price at time t that would be socially optimal if bad news arrived at time t. Fig. 4 depicts the virtual carbon price path for the near and longer terms. Note that the price path is "virtual" because it is drawn against the economy that does not, but could, experience the damage. The starting level is given by our calibration at 93 EUR/tCO_2, as this is the highest price estimate that we applied to the immediate arrival of impacts. The virtual price increases for a long period of time reflecting the expanding world economy.

As the level of the virtual price path indicates —it reaches levels exceeding the carbon price without news by a factor of ten— the bad news outcome can be characterized as a catastrophic outcome.

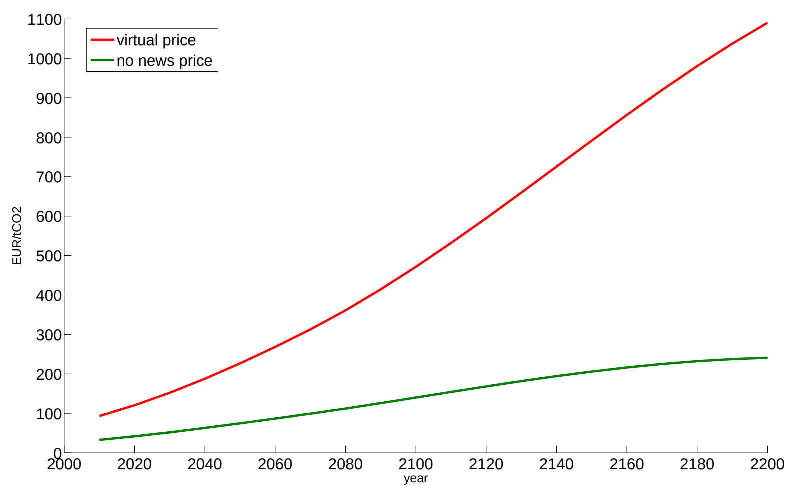

Figure 4: The virtual carbon price

6 Concluding Remarks

The learning model used in the analysis is stylized — it does not allow a gradual arrival of information, say, learning of damages from extreme weather events. However, the model can be easily extended to this direction, and under certain assumptions, qualitatively similar carbon price dynamics can follow. The key observation is that if the precision of the estimates for the climate impacts increases over time, and if a lack of precision leads to a carbon price mark-up, then the assessment of the utility-costs from current emissions in most cases declines over time, while the expanding economy still puts pressure on carbon price increases.

Consider the following model of learning about the true damage parameters Δ_y and Δ_u. Assume that priors are normally distributed random variable with mean m_y (m_u, resp.) and variance σ_y^2 (σ_u^2, resp.). Signals are the realizations of damages that come from the true distributions, but initially we cannot tell apart damages from weather volatility and those from more persistent climate impacts. We experience output losses given by

$$\exp(-\Lambda_{y,t} D_t)$$

where we observe the state of the climate D_t, and output losses contain a stochastic signal for the persistent damage sensitivity

$$\Lambda_{y,t} = \Delta_y + \varepsilon_{y,t}$$

with $\varepsilon_{y,t} \sim N(0, \sigma_\varepsilon^2)$ and i.i.d. across periods and also independent of Δ_y. While we observe $\Lambda_{y,t}$ but cannot tell apart the contribution of the noise and the true damage that has an initial prior $\Delta_y \sim N(m_{\Delta,y}, \sigma_{\Delta,y}^2)$ with $m_{\Delta,y} > 0$. Thus, in expectations, temperature causes output losses but there can also be temporary positive productivity shocks, $\varepsilon_{y,t} < 0$. Let $H_{y,t} = \Lambda_{y,1}, ..., \Lambda_{y,t}$ be the history of observed damages. Then, in this setting, we can apply the normal learning rule to see that after sufficiently many observations, $\mathbb{E}\Lambda_y(\cdot)$ converges to the initially unknown true mean Δ_y. However, for such learning, the improving estimate will have no effect on policies, as, obviously, there will be no expected trend in learning. Also notice that the model allows for positive effects of climate change on output. To capture the effect of improving estimates on policies, consider the intangible damages that we can include in the periodic utility as having a log-normal distribution,

$$u_t = \ln(c_t) - \exp(\Lambda_{u,t}) D_t$$

where

$$\Lambda_{u,t} = \Delta_u + \varepsilon_{u,t}$$

with zero-mean normal realizations $\varepsilon_{u,t}$ that are i.i.d. across periods and also independent of Δ_u. Here, too, the initial prior is normal, $\Delta_u \sim N(m_{\Delta,u}, \sigma_{\Delta,u}^2)$. Thus, again, the realized (experienced) damage depends on the unknown damage-generating process and on the noise term. As above, we obtain the expected intangible damage, after a given history of observed damages. We can easily calibrate such a model with a lognormal distribution of intangible damages.

However, in this set up, there is a difference between the tangible damages that have, in utility terms, normal distribution and the intangible damages that have, in utility terms, log-normal distribution. For the intangible damages, the expected damages have a skewed distribution with a fat tail for large damages, so that

$$-\mathbb{E}\{\exp(\Lambda_{u,t+1})D_{t+1}\} = -\exp(\Lambda_u(H_{u,t}) + \frac{1}{2}var[\Lambda_u\,|H_{u,t}])D_{t+1} \tag{12}$$

where $var[\Lambda_u\,|H_{u,t}]$ is the conditional variance of the intangible damage after history $H_{u,t}$ of observations. By the normal learning rule, after t observations, this takes the form:

$$var[\Lambda_u\,|H_{u,t}] = \frac{1}{th_{\varepsilon_u} + h_{\Delta_u}} + \frac{1}{h_{\varepsilon_u}} \tag{13}$$

The expected damage hence depends on the distribution: the optimal policy h_t will depend on the distributional assumptions and on the details of learning. Moreover, (13) implies that the fat-tail effect on damages declines over time.

If all damages follow log-normal distribution as in (12), our carbon price distribution is also log-normal since τ_t is a linear function of the damage. This allows us to follow similar steps as for the simple analysis presented in the main text, to conclude that the effect of uncertainty on carbon pricing declines over time.

References

[1] Fischer C. and R. Newell, 2008, Environmetnal and technology policies for climate mitigation, J of Environmental and Economic Management 55:142-162.

[2] Gerlagh, R., and M. Liski, Gerlagh,Carbon Prices for the Next Thousand Years (2013). Available at http://hse-econ.fi/liski/

[3] Gerlagh, R., and M. Liski, Gerlagh,Carbon Prices for the Next Thousand Years (June 28, 2012). CESifo Working Paper Series No. 3855. Available at http://hse-econ.fi/liski/

[4] Golosov, M., J. Hassler, P. Krusell, A. Tsyvinski (2011), Optimal taxes on fossiel fuel in general equilibrium, NBER WP 17348.

[5] Hasselmann, K., S. Hasselmann, R Giering, V Ocana, H v Storch (1997), Sensitivity of optimal CO_2 emissions paths using a simplified structural integrated assessment model (SIAM), Climatic Change 37: 345-387

[6] Hassler, J., P. Krusell, and C. Olovsson (2011), Energy-saving technical change, mimeo.

[7] Hooss G, R. Voss, K Hasselmann, E Maier-Reimer, F Joos (2001), A nonlinear impulse response model of the coupled carbon cycle-climate system (NICCS), Climate Dynamics 18: 189-202

[8] Houghton, R.A. 2003.Revised estimates of the annual net flux of carbon to the atmosphere from changes in land use and land management 1850-2000. Tellus55B(2):378-390.

[9] IPCC, Intergovernmental Panel on Climate Change (2000), Special Report on Emissions Scenarios, edited by N. Nakicenovic and R. Swart, Cambridge Univ. Press, Cambridge, U.K.

[10] Kelly, D.L., and C.D. Kolstad (1999), Bayesian learning, growth, and pollution, Journal of economic dynamics and control 23, 491-518.

[11] Krieglera, Hall, Helda, Dawson, and Schellnhuber (2009), "Imprecise probability assessment of tipping points in the climate system", Proceedings of the National Academy of Sciences, vol 106: 5041-5046.

[12] Leach, A. J. (2007), ÔThe climate change learning curveÕ, Journal of Economic Dynamics and Control 31, 1728 -1752.

[13] Lemoine, D. M. and Traeger, C. P. (2010), Tipping points and ambiguity in the integrated assessment of climate change, Department of Agricultural and Resource Economics, UC Berkeley, Working Paper Series 1704668.

[14] Maier-Reimer E. and K.Hasselman (1987), Transport and storage of CO2 in the ocean - an inorganic ocean-circulation carbon cycle model, Climate Dynamics 2, 63-90.

[15] Nordhaus, W. D.(2007), A Review of The Stern Review on the Economics of Climate Change, Journal of Economic Literature, 45 (3), 686-702.

[16] Nordhaus, W. D.(2010). Economic Policy in the Face of Severe Tail Events, Journal of Public Economic Theory, 14 (2), 2012, pp. 197- 219.

[17] Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M. Tignor and H.L. Miller (eds), "Climate Change 2007: The Physical Science Basis", Technical Summary, Table TS2.5, footnote b.

[18] Stern, N. (2006), "The economics of climate change: the Stern review", Cambridge, UK: Cambride University Press.

[19] Tol, R. (2009), The economic effects of climate change, Journal of economic perspectives, 23(2): 29-51.

[20] Weitzman, M. (2007). A Review of The Stern Review on the Economics of Climate Change, Journal of Economic Literature, 45 (3), 703-724.

[21] Weitzman, M. (2009), On modeling and interpreting the economics of catastrophic climate change, The Review of Economics and Statistics 91(1), 1-19.