This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Gerlagh, Reyer; Liski, Matti

Carbon prices for the next hundred years

Published in:
ECONOMIC JOURNAL

DOI:
[10.1111/econj.12436](https://doi.org/10.1111/econj.12436)

Published: 01/03/2018

Document Version
Peer reviewed version

Please cite the original version:
Gerlagh, R., & Liski, M. (2018). Carbon prices for the next hundred years. *ECONOMIC JOURNAL*, 128(609), 728–757 . [10.1111/econj.12436]. <https://doi.org/10.1111/econj.12436>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Carbon prices for the next hundred years

Reyer Gerlagh and Matti Liski*

July 8, 2016

Abstract

This paper examines the socially optimal pricing of carbon emissions over time when climate-change impacts are unknown, potentially high-consequence events. The carbon price tends to increase with income. But learning about impacts, or their absence, decouples the carbon price from income growth. The price should grow faster than the economy if the past warming is not substantial enough for learning the true long-run social cost. It grows slower than the economy as soon as the warming generates information about events that could have arrived but have not done so. A quantitative assessment shows that the price grows roughly at the rate of the economy for the next 100 years.

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

*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 two anonymous referees and the editor (Cripps) for valuable comments helping to restructure the manuscript. Helpful comments on earlier versions from Derek Lemoine, and audiences at the CESifo area conference “Energy and Climate Economics”, Paris School of Economics, Berlin TU, EAERE Toulouse, and Helsinki are gratefully acknowledged. We have also benefitted from discussions with Lint Barrage, David Hemous, Sverre Jensen, Terry Iverson, Anthony Millner, and Nicolas Treich. We thank James Corbishley for research assistance. Matti is grateful for funding from the Academy of Finland project 140776/Ficca, and from the Aalto University project SAGA.

1 Introduction

A price for carbon is a monetary measure of the social cost that follows from releasing a unit of carbon dioxide to the atmosphere, based on expected climate-change impacts. However, there is little or no quantitative information on how climate change impacts our economies, although there is extensive research on what such impacts might be.¹ Rather, the social cost measures build on *beliefs* about impacts that will be updated when the “climate experiment” generates actual, potentially catastrophic, impacts. But this can take a long period of time; the past century of carbon emissions has not yet led to precise estimates, and another 50-100 years may pass without additional hard evidence on the ultimate consequences of current emissions.² Over such long horizons, policies on global warming have to deal with fundamental changes in the economy as much as with forthcoming information on how seriously the economy will be impacted.

Economic growth stands out as a major driver of global change when considering a period such as the next 50-100 years. For the coming century, global income is expected to grow by multiple factors, in part due to the rise of the middle class in major emerging economies.³ The US government has recently developed estimates for the carbon price, for regulatory purposes, assuming that the global GDP increases by a factor that varies between five and seven in this time-span (see Greenstone et al., 2013).⁴ The future economy grown five to seven times bigger, and with almost a century of extra climate-change experience, prices emissions differently — but how exactly?

Increasing incomes can lead to larger stakes, that is, bigger economic losses per temperature increase; thereby, income growth drives a gradual tightening of policies over time, the “climate policy ramp” (Nordhaus 2007; Golosov et al. 2014).⁵ Further support

¹See IPCC (2014) for a survey on methods and results. There is a growing empirical literature on how climate impacts various sectors of the economy (e.g., Deschenes and Greenstone, 2007, and Schlenker and Roberts, 2009, Dell, Jones, and Olken, 2012).

²Among climate researchers, the delays in learning the impacts are widely accepted. For example, Roe and Baker (2007) establish that, because of positive feedback mechanisms of the climate system, it is unlikely that we will better understand the temperature sensitivity to emissions in the near future.

³See, for example, the IPCC Special Report on Emissions Scenarios (2000), U.S. Climate Change Science Program (2007), Stanford Energy Modeling Forum (for example, in Weyant et al. 2006).

⁴More recently, Cai et al. (2015b) evaluate the social cost of carbon under stochastic growth, with stochastic total factor productive calibrated to match empirical data. See also Jenssen and Traeger (2014).

⁵There are other arguments such as green technological change for not following gradualism but rather a jump-start in emissions pricing (van der Zwaan et al. 2002; Gerlagh, Kverndokk and Rosendahl 2009; and Acemoglu, Aghion, Bursztyn, and Hemous, 2012).

for the policy ramp is given by expectations that mitigation options become cheaper, and that climate-change losses are increasing more than proportionally with temperature change (Wigley et al., 1996). Additionally, when economic growth levels off, capital returns will diminish (Piketty and Zucman, 2014), and thus the relative returns of the climate investments increase.

But the policy ramp, as such, does not describe carbon prices for the situation where climate impacts are fundamentally unknown, including the possibility of a climate tipping point with large or catastrophic economic impacts. In fact, the critics of the climate-economy models,⁶ used for evaluating the policy ramp, are concerned that the models do not incorporate the main reason for having a carbon price, that is, the climate change “unknowns” (Pindyck 2013).⁷

The carbon price development is a necessary input in practical policy making; for example, plans for phasing out a fleet of polluting vehicles, power generating plants, or energy inefficient buildings depend on the assumed time path for the carbon price over the decades to come. Taking the possibility of tipping points as the main reason for the carbon price, the question arises if the policy makers should still assume that the carbon price grows roughly at the same rate as the economy? An emerging literature uses numerical stochastic climate economy models, rather than versions of the deterministic models criticised by Pindyck (2013), to systematically evaluate if and how the policy ramp is affected by tipping points. One recent study finds that the initial carbon price is relatively insensitive to the possibility of a tipping point but it should grow faster than in the deterministic case (Lemoine and Traeger 2014); others find an opposite result, that the carbon price starts considerably higher and grows slower over time (Lontzek et al. 2015).

Our contribution is an analytical climate-economy model with closed-form policy rules for pricing uncertain high-consequence events. The model is detailed enough for a quantitative assessment of the optimal carbon price path and thus comparable with

⁶Most evaluations of the social cost of carbon build on a set of middle-of-the-road assumptions on climate change impacts, commonly expressed in terms of GDP losses, and then use climate-economy models such as DICE, FUND, or PAGE (see Greenstone, Kopits, and Wolverton (2013) for a succinct description and references) that combine the impact assumptions with background scenarios to obtain a monetised value for the social cost.

⁷For many economists, such climate uncertainties and the implied low-probability but high-consequence events, which cannot be ruled out by new information any time soon, have become the prime argument for having a price for carbon (Weitzman, 2009, 2011, 2013). See also Lange and Treich (2008), and Heal and Millner (2014) for surveys on uncertainties in climate-change economics.

the computational approaches; yet, the model is explicit about the key drivers of the policy ramp. Similarly as Golosov, Hassler, Krusell, and Tsyvinski (2014) develop the analytical foundations of climate policies building on a workhorse model in analytical macroeconomics (Brock and Mirman, 1972), we develop a tractable pricing rule for the tipping points.

The optimal pricing of tipping points is sensitive to past experience, with the following dichotomy. First, we find that the carbon price grows faster than the economy, if the past experience cannot generate information that is substantial enough for learning the true long-run cost of the emissions. For example, 1 degree Celsius warming above the preindustrial levels may be taken as a safe limit, but this also implies that temperatures below this limit generate no information about the future events. As the economy approaches such higher temperatures, the potential regime shift damage becomes increasingly more relevant to the policy maker and the optimal carbon tax grows faster than implied by a model with smooth annual economic damages from climate change. Second, as soon as warming is sufficient to generate information about events that could have arrived but have not done so, the carbon price grows slower than the economy and may even decline. Thus, conditional on the possibility of a regime shift but no occurrence, the event becomes less likely over time and the optimal carbon tax falls relative to a smooth model. This learning effect captures the idea that if the economy grows over periods such as the next 50-100 years without verifiable economic climate-change impacts, there can be increasing optimism. There is thus a downward pressure in the carbon price from this source that coexists with the upward pressure from the growing economic stakes. The question is then which pressure dominates?

In a quantitative assessment, we address the question if there can, in principle, be a case for increasing optimism that reverses the ramp implied by growing economic stakes. We calibrate the beliefs regarding the impacts, learning rates, and the potential economic losses to the choices made in the numerical stochastic climate economy models (Lemoine and Traeger, 2014; Lontzek et al., 2015). It turns out that there is a fairly representative set of choices for the “triple” (beliefs, learning rates, impacts). To overturn the ramp where the carbon price grows roughly at the rate of the economy for the next 100 years, one would have to dramatically deviate from the typical values for the triple; in particular, the planner would have to rule out severe climate impacts by orders of magnitude faster than what is implied by the scenario in our explorative calibration.

The implications of this observation are not minor: as we show, the extreme persistence of beliefs about ultimate impacts can lead to policies that fully decarbonise

the economy, even without actual experience of such climate-change impacts. Take, for example, the 2015 United Nations Climate Change Conference outcome that calls for decarbonisation by the end of the century. We assess how big the unknown event would have to be to justify this target as an optimal policy. The potential impact on the economy should be by a factor 10 larger than the central impact assessments in the literature.

Methodologically, the results provide a bridge between the climate-economy models (Integrated Assessment Models, IAMs) based on middle-of-the-road impact estimates, and their critics expressing concerns that these models do not cover “unknowns” that should be the main reason for having a carbon price (Pindyck 2013).⁸ We reconcile the views in a model supporting a carbon price ramp, not based on moderate climate change damages that arise smoothly over time, but through a description of uncertain high-consequence events and belief updating. Effectively, the setting developed here becomes a macroeconomic experimentation model with learning of unknown arrival rates as in Keller, Rady and Cripps (2005). The description of the macroeconomy builds on the Brock-Mirman model (1972), following Golosov et al. (2014); however, we introduce climate change differently through a hidden state that determines whether a negative productivity shock can hit the economy in the future. Also, we adopt a rich emissions-temperature response, including the delays between emissions and potential impacts.⁹ Such delays are necessary for a quantitative assessment that has some hope to be comparable to those in the comprehensive climate-economy models.

Our results complement those in the literature that introduces uncertainties and anticipated learning into the integrated assessment models. For example, Kelly and Kolstad (1999) and Leach (2007) analyse the speed of learning under smooth climate change and find that learning is very slow. Jensen and Traeger (2013) quantify the impact of such slow learning on the social cost of carbon; they find that the impact on the current policies is not significant. Kelly and Tan (2015) find that some aspects of learning are faster under fat-tailed climate risk. Thus, while our approach is analytical and the focus is on regime shifts, the substantial lessons are consistent with the literature on smooth

⁸By the nature of our quantitative exercise, we rule out “tail events”. The supporting potential high-damage climate event that justifies the estimated initial carbon price is equivalent to a GDP-loss of about 10 per cent at temperatures that are 3 degrees Celsius above the pre-industrial level. Such an event is economically significant 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; see, for example, Nordhaus (2010) and Millner (2013) for discussion.

⁹This extension builds on Gerlagh and Liski (2016).

learning.

The paper is structured as follows. In Section 2, we first explain the basic planning problem, and introduce the learning and climate dynamics. In Section 3, we characterise the optimal policies. In Section 4, we calibrate the model and perform the quantitative assessment. Section 5 concludes. The online supplementary file contains a program for reproducing the quantitative assessment and the graphs in the text.¹⁰

2 The climate-economy model

2.1 The basic setting

We consider a climate-economy planning problem where production possibilities at time t depend on capital k_t inherited, and potentially also on the full history of carbon input use,

$$s_t = (z_0, \dots, z_{t-1}).$$

Given k_t and history s_t at time period t , consumption, c_t , and carbon inputs, z_t , are chosen to maximise the expected discounted utility

$$\max \mathbb{E}_t \sum_{\tau=0}^{\infty} \delta^\tau u_{t+\tau} \quad (1)$$

where $0 < \delta < 1$ is the discount factor and $u_{t+\tau}$ is the periodic utility, specified below. The chosen allocations must satisfy

$$c_t + k_{t+1} = y_t, \quad (2)$$

with $y_t = f_t(k_t, s_t, z_t)$ denoting the output at time t . Losses due to climate change depend on the history of emissions s_t through variable D_t that is a measure of the global mean temperature increase above the pre-industrial levels at time t . We assume that this measure is a function of history s_t ,

$$D_t = \sum_{\tau=1}^t \mathcal{R}(\tau) z_{t-\tau} \quad (3)$$

where the weights $\mathcal{R}(\tau)$ define the “emissions-temperature response”. That is, current emissions z_t affect temperatures at some later time $t + \tau$ according to a known response function $\mathcal{R}(\tau)$:

$$\frac{dD_{t+\tau}}{dz_t} = \mathcal{R}(\tau) > 0. \quad (4)$$

¹⁰Follow the link https://www.dropbox.com/sh/7meos655j14jh5p/_dlr8X_FHI

Below, in equation (9), we provide a parametric form for the response function, connecting to the fundamentals of the climate problem. The key characteristic of $\mathcal{R}(\tau)$ is the considerable delay of the response following an impulse of emissions; it has a non-linear shape peaking several decades after the date of the emissions, and a fat tail of almost permanent impacts. Simplistically, there is no uncertainty about $\mathcal{R}(\tau)$; the response serves the purpose of introducing delays to the potential impacts on the economy that, in turn, will be uncertain.

Output is given by production function

$$y_t = k_t^\alpha A_t(z_t) \exp(-\Delta_{y,t} D_t), \quad (5)$$

where $0 < \alpha < 1$, and the contribution of carbon inputs z_t enter through function $A_t(z_t)$ that captures the energy sector of the economy as well as the total factor productivity. Formally, we treat carbon inputs z_t as reproducible, and thereby do not impose an upper limit for the total cumulative use. The assumption is motivated by the size of carbon deposits in the form of coal that exceeds the absorptive capacity of the atmosphere.¹¹ The analytical results do not require a specific form A_t ; we postpone the detailed discussion of A_t to Section 4.3. Here, we merely assume that carbon input z_t has a positive but diminishing marginal product.

Losses from the temperature increase arise potentially from two sources. First, they can lead to reduced output, through the negative productivity impact in (5), as in most applied climate-economy models (e.g., Nordhaus, 2008). This impact depends on the full history of emissions, determining current temperature D_t through (4), and damage coefficient $\Delta_{y,t} \geq 0$ that is a stochastic variable. For illustration, assuming constant $\Delta_{y,t} = \Delta_y > 0$, the output loss per unit of emissions equals $[1 - \exp(-\Delta_y \mathcal{R}(\tau))]$, τ periods after the date of emissions. For the life-path of the impact, see Fig. 1, discussed in detail below.

Second, we allow a direct impact on periodic utility that we define as

$$u_t = u(c_t) - \Delta_{u,t} D_t, \quad (6)$$

where $u(c_t) = \ln(c_t)$, $\Delta_{u,t} \geq 0$.¹² While the current and past impacts are known, the

¹¹See Gerlagh (2011) and van der Ploeg and Withagen (2012) for detailed analysis. Golosov et al. (2014) consider the effect of resource scarcity on climate policies for a parametric class of preferences and technologies coming close to the ones used in this paper.

¹²The impact of climate change on output can be equivalently represented as direct utility losses in our framework, a result that we want to illustrate by this formation. Our main interest is in the productivity

sequence of future impacts

$$\{\Delta_{y,t+\tau}, \Delta_{u,t+\tau}\}_{\tau>0} \quad (7)$$

is not known to the policy maker.

The next Section specifies how the policy maker learns the future impacts in (7). But first we stop to discuss the critical assumptions of the main model structure that originates back to Brock-Mirman (1972) who first introduced the core of this analytical consumption choice model. There are four main critical assumptions: (i) logarithmic consumption utility; (ii) full one period depreciation of capital, see eq. (2); (iii) emissions-temperature response that we have denoted by $\mathcal{R}(\tau)$; and (iv) exponential impact of climate on output. Barrage (2014), in a supplement to Golosov et al. (2014), has numerically assessed the loss of generality implied by assumptions (i) and (ii) in a context similar to ours. In the end, the question is how well this set of assumptions, that leads to analytical policy rules, approximates a more general case that requires numerical approaches. Log utility implies a relatively low preference for consumption smoothing over time, thereby making the decision maker “patient”, increasing the initial carbon price level when compared to a case where the utility function has more curvature. As long as we stay in the expected utility framework, this overshooting effect vanishes if we adjust the time discount rate. We make such an adjustment when calibrating the model in Section 4.1.¹³

In contrast, the one period depreciation assumption tends to decrease the carbon price level and its growth, since it implies a lower growth of the economy than in the case where some capital survives to the next period. In our analysis, one period has a length of a decade, which partly offsets the assumption.¹⁴ But, again, it is possible to make adjustments in the calibration procedure to almost exactly offset the full 100 per cent depreciation, closing the gap between the predictions of the numerical models with partial depreciation and analytical models with full depreciation.¹⁵

impact; a comprehensive analysis of the directly utility impact would need to take a stand on how the relative value of market and non-market goods depends on climate change, in the spirit of Hoel and Sterner (2007).

¹³See Fig. 6 in the Appendix and the related discussion in the main text of Section 4.1. However, deviations from the expected utility setting, for example, separation of attitudes towards risk and consumption smoothing, cannot be handled with small adjustments but require a different modelling approach. See Jenssen and Traeger (2014) and Traeger (2014,2015) for analysis of this issue in the climate and economy context.

¹⁴Arguably, from a business cycle perspective, 10-year periods are rather long for TFP shocks. In climate change, the relevant time horizon for planning is long enough to justify decadal time steps.

¹⁵In the Appendix, we use Fig. 6, cited in the footnote above, to illustrate numerically that the desired

Assumption (iii), the emissions-temperature response, has a simplistic form in Golosov et al. (2014). This analytical approximation can perform very poorly, when compared to the mainstream (numerical) climate-economy models, a result that we demonstrate in Gerlagh and Liski (2016). We follow the formulation in Gerlagh and Liski (2016), and introduce it in Section 2.3.

Deviations from assumption (iv), the exponential loss of output, is extensively analysed in Van den Bijgaart et al. (2016). Combined with the other assumptions, the exponential loss implies unitary elasticity of output losses with respect to income, from a given temperature increase, as will be seen as we derive the policy rules. Clearly, the property is central to the feature of the model that the carbon price increases in lock-step with income.¹⁶ Van den Bijgaart et al. (2016) show that the analytical policies remain close to those obtained from numerical models, but can deviate considerably under extremely convex or concave representations of losses. The assumption made in the current paper can be seen to represent the central case.

2.2 Learning dynamics

There are several conceivable approaches to the process that generates the future climate-change impacts in (7), specifying how the decision maker forms expectations $\mathbb{E}_t\{\Delta_{y,t+\tau}|\Omega_t\}$ and $\mathbb{E}_t\{\Delta_{u,t+\tau}|\Omega_t\}$, where Ω_t is the current information set. In one approach, the climate generates “experience” and thus evidence on a continual basis through events such as hurricanes, hot-summer spells, or perhaps a long-period of stable climatic conditions; whether experienced events are due to climate change or within the normal variation is initially a matter of beliefs. We are not interested in gradual learning but in tipping points that are hard to learn before they arrive. We assume the economy starts with no experienced losses but may irreversibly enter a climate-economy state where potentially catastrophic damages occur. Thus, the decision-maker learns by (not) observing damages, which allows updating the beliefs on the ultimate arrival of such damages.

Specifically, there are two climate-economy 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 coefficient at time t for output, affecting production in (5), is $\Delta_{y,t} = \Delta_y I_t$, where $\Delta_y > 0$ is a constant, independent of time. Similarly, the

adjustment involves recalibrating the capital factor share and the size of the initial capital stock; the analysis detailing the analytics of the adjustments is provided in the online Appendix.

¹⁶More precisely, in this paper, the property holds when learning of climate-change impacts is slow or if the impacts are known for sure.

direct utility loss, entering (6), is $\Delta_{u,t} = \Delta_u I_t$.¹⁷

How does the economy make the transit from $I_t = 0$ to $I_{t+1} = 1$? We can think of an explicit statistical model where the tipping point temperature is unknown but with known prior distribution; as soon as the critical temperature is exceeded, the event occurs. This approach has been recently used in Lemoine and Traeger (2014); it originates back to Tsur and Zemel (1995, 1996). In this approach, experimentation is extremely informative: the planner is 100% certain that the critical threshold has not been exceeded at time t if the event has not occurred by time t . Another common approach assumes event hazard rates that depend on the relevant state of the system, temperature in our case, without explicitly invoking a threshold. Thus, stabilising the state leads to a stabilised but positive hazard rate for the event.¹⁸ This latter assumption features in recent papers by Cai et al. (2013) and Lontzek et al. (2015); see also Clarke and Reed (1996), Polasky et al. (2011), and Sakamoto (2014).

One may argue that the appropriate statistical model lies between these two approaches. The planner might have triggered the event in the past, even though it has not revealed itself yet. This implies delays in learning the event. In models with exogenous (state-dependent) hazards, this feature is (implicitly) captured since the event arrival probability remains positive in the long run; it might still happen, even after the stabilisation of the relevant state. Yet, the current beliefs should explicitly depend on the past experiments such as the thresholds tried in the past; in this approach, the state dependent hazard cannot be an exogenous function but depends on the history of past experiments. We seek to make progress towards modelling such delays in information arrival, by allowing a role for temperature levels in the hazard rate determination.

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 unknown to the policy maker. For the main part of the analysis, we assume that p has a discrete prior distribution: it can either take value $p = 0$ or $p = \lambda > 0$, where constant λ can be interpreted as the intensity of experimentation. Later, we allow λ to change over time, and also depend on critical temperature levels.

¹⁷The model describes the delay between emissions and temperatures through the response function $\mathcal{R}(\tau)$, but it does not explicate that tipping points can have their own very slow dynamics. Some tipping point, once triggered, might unfold their impacts only gradually over time. If the Greenland icesheet becomes unstable, it may still take centuries for the sea level to rise. For rich tipping-point dynamics, see Cai et al. (2015a).

¹⁸The property of a persistent hazard rate marks an important distinction between this approach and the threshold model, where such persistent hazard rate is impossible.

We assume a subjective prior probability $\mu_0 > 0$ for a positive hazard rate, $p = \lambda$. The probability for eventual climate impacts satisfies:

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

Thus, $\mu_0 > 0$ can be interpreted as the decision maker's initial belief that there will be long-run climate impacts. To illustrate, the assumptions in Lemoine and Traeger (2014) imply that $\mu_0 \approx .8$, if the temperature is expected to increase to about 3.3 degrees Celsius.¹⁹ In our quantitative analysis, we use this number as our benchmark.

Let μ_t denote the posterior probability that $p = \lambda$ at time t , conditional on no impacts having yet occurred by time t , $I_t = 0$. Each period 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 the beliefs are updated to μ_{t+1} . We can write the Bayesian updating rule as²⁰

$$\begin{aligned} \mu_t &= \Pr(p = \lambda | I_t = 0) \\ &= \frac{\mu_0(1 - \lambda)^t}{\mu_0(1 - \lambda)^t + 1 - \mu_0} \end{aligned} \tag{8}$$

which is the probability that climate change impacts 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.²¹ The choices in triple (μ_t, λ, Δ) describe the current beliefs, the underlying stochastic process for damages, and the size of damages, respectively.²² The triple thus describes the subjective part of carbon pricing, the informativeness of the experiment, and the economic stakes involved.

¹⁹See their fn. 11.

²⁰Note 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.

²¹One 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 major impacts is evidence for not having major 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.

²²We define Δ in Remark 2

Variants of the learning dynamics considered here are common in other fields of economics but some features of the setting deserve attention. Malueg and Tsutsui (1997) were among the first to consider learning of unknown Poisson rates in an R&D race; see also, for example, Keller, Rady, and Cripps (2005), and Bonatti and Hörner (2011). In this literature, new information is generated by periodic effort; no current effort means no new information. In our setting, one could also introduce effort for information acquisition so that news about impacts could arrive separately from experiencing them. However, in climate change it seems less natural to assume that pure research could produce robust information about how the physical reality interacts with the economy, without actual experienced impacts. The basic model introduced here connects the experience and learning in a stark way.²³

2.3 Climate dynamics

Learning is delayed but there is also a mechanical delay in how the economic stakes depend on emissions, due to climate-system inertia. The delay between the current emissions and future temperatures, denoted by $\mathcal{R}(\tau)$ in Section 2.1, is captured by a somewhat detailed analytical representation; nevertheless, it is a necessary input to the quantitative assessment if one is interested in results comparable to those in the literature using numerical climate-economy models. We build on a closed-form for $\mathcal{R}(\tau)$ that is derived in Gerlagh and Liski (2016), Theorem 1.

Remark 1 *Consider a carbon diffusion process, described by a set of impulse-responses $\mathcal{I}$, with fraction $0 < a_i < 0$ of emissions having decay rate $0 \leq \eta_i < 1$, $i \in \mathcal{I}$. For temperature sensitivity π and adjustment speed ε , the impact of emissions at time t on temperatures at time $t + \tau$ is*

$$\frac{dD_{t+\tau}}{dz_t} = \mathcal{R}(\tau) = \sum_{i \in \mathcal{I}} a_i \pi \varepsilon \frac{(1 - \eta_i)^\tau - (1 - \varepsilon)^\tau}{\varepsilon - \eta_i} > 0. \quad (9)$$

To explain this result, we outline first the two main determinants of the response: the carbon cycle and the relationship between carbon concentrations and temperatures. The carbon cycle refers to a diffusion process of carbon between reservoirs of carbon, such as

²³ It should be noted that an explicit model of learning that combines the threshold model with learning delays should keep track of the experiments performed in the past: each past temperature increase is a new experiment whose outcome becomes known only with a delay. For tractability, we ignore such history dependence; Proposition 2 below approximates this feature. This question is the topic of ongoing research by Liski and Salanié (2015).

those in the atmosphere, oceans and biosphere. Obviously, the atmospheric reservoir is the one relevant for climate warming but the other reservoirs are relevant for the delays and persistencies of changes in the atmospheric stock. Assuming a linear diffusion, the system can be de-coupled by eliminating interactions between the reservoirs, leading to an isomorphic system of separable impulse-responses for carbon stocks (Maier-Reimer and Hasselman 1987). The shares and decay rates have intuitive meanings, discussed just below, and they follow from the physical description of the system of carbon reservoirs.²⁴

The carbon cycle is relatively well understood in natural sciences but the relationship between temperatures and carbon concentrations is fundamentally uncertain (see, for example, Roe and Baker, 2007). Acknowledging these complications, we note that economic impacts introduce yet another layer of fundamental uncertainty; we focus on this uncertainty and make the following simplistic assumptions on the determinants of the climate equilibrium. Emissions z_t increase the atmospheric CO_2 stock, through the carbon cycle, and there is a linear relationship between the steady state atmospheric CO_2 stock and the steady state level of D_t . This relationship is 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 CO_2 stocks one-to-one but the direct temperature increase is only $\varepsilon\pi$ -units

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 under decay i still lives after τ periods, and the term $-(1 - \varepsilon)^\tau$ captures the slow temperature adjustment. The limiting cases can be helpful. Consider one CO_2 reservoir. 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 $\mathcal{R}(\tau) = \pi[1 - (1 - \varepsilon)^\tau]$. If atmospheric carbon-dioxide depreciates fully, $\eta = 1$, the temperature immediately adjusts to $\pi\varepsilon$, and then slowly converges to zero, $\mathcal{R}(\tau) = \pi\varepsilon(1 - \varepsilon)^{\tau-1}$. If temperature adjustment is immediate, $\varepsilon = 1$, then the temperature response function directly follows the carbon-dioxide depreciation $\mathcal{R}(\tau) = \pi(1 - \eta)^{\tau-1}$. If temperature adjustment is absent, $\varepsilon = 0$, there is no response, $\mathcal{R}(\tau) = 0$.

When multiplying temperature measure D_t by given output-loss coefficient $\Delta_y > 0$, we can interpret the emissions-temperature response as an emissions-damage response.

²⁴The true diffusion process is non-linear (Joos et al. 2013); the linear representation should be seen as an approximation.

Figure 1: Emissions-damage response. The path depicts the output loss associated with $1TtCO_2$ impulse of carbon at time $t = 0$ for $\Delta_y = 1$ and $\pi = .0156$.

Fig. 1 shows the life path of damages (percentage of total output) caused by an impulse of one Teraton of Carbon [$TtCO_2$] in the first period.²⁵ The output loss is thus measured per $TtCO_2$, and it equals $1 - \exp(-\Delta_y \mathcal{R}(\tau))$, τ periods after the impulse. The non-monotonicity of the response, as depicted in Fig. 1, captures well the climate impact dynamics, for example, in DICE-2007 (Nordhaus, 2008).

The physical data on carbon emissions, stocks in various reservoirs, and the observed concentration developments can be used to calibrate a three-reservoir carbon cycle representation; we choose the following emission shares and depreciation factors per decade:²⁶

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

Thus, about 16 per cent of carbon emissions does not depreciate while about 45 per cent has a half-time of one decade. We assume $\varepsilon = .183$ per decade, implying a global temperature adjustment speed of 2 per cent per year. Normalizing the output loss parameter at unity, $\Delta_y = 1$, and setting $\pi = .0156$ [per $TtCO_2$, see Gerlagh and Liski (2016)] is consistent with the Nordhaus (2008) baseline where a temperature rise of 3 degrees Celsius

²⁵One $TtCO_2$ equals about 25 years of global CO_2 emissions at current levels (40 $GtCO_2/yr.$)

²⁶Some fraction of emissions depreciates within one decade from the atmosphere, and therefore the shares a_i do not sum to unity. The choices here are based on Gerlagh and Liski (2016) but similar representative numbers can be found in the scientific literature; see, e.g., Maier-Reimer and Hasselman (1987).

leads to about 2.7 per cent loss of output.²⁷ These quantitative choices parametrise the emissions-temperature response that is depicted in Figure 1. In the calibration below, we choose Δ_y to match previous studies on tipping points; throughout, $\Delta_y = 1$ refers to Nordhaus’ baseline.

3 Optimal policies

The economic problem defined through (1)-(6), together with learning and dynamics in Sections 2.2 and 2.3, is essentially the same as in Brock-Mirman (1972) consumption-choice model. The dynamic programming arguments leading to the optimal saving policies are well known in analytical macro-economics (Sargent, 1987).²⁸ The state vector is (k_t, s_t, Ω_t) , where information set Ω_t includes I_t and the current belief μ_t . Because of the log-utility for consumption, full capital depreciation in one period, and Cobb-Douglas capital contribution, the optimal savings can be expressed as a share of the gross output that is left as capital for the next period:

$$g = \alpha\delta. \quad (10)$$

Anticipating the quantitative analysis, we will match capital share α and savings g with the representative values, leaving time preference δ as the adjusting parameter. Thus, in this parametric class, savings can be evaluated without paying attention to climate policies; formally, welfare (i.e., value function) will be separable in the contributions of k_t and s_t . Therefore, the climate policy analysis can be conducted by taking savings g as given and by tracking the direct utility impacts of the potential loss from climate change.²⁹ It proves useful to aggregate both the potential output and direct utility losses into one measure:

Remark 2 *For $I_t = 1$, the present-value loss of utils from marginal climate change at time t is*

$$\Delta \equiv - \sum_{\tau=0}^{\infty} \delta^{\tau} \frac{du_{t+\tau}}{dD_t} = \Delta_u + \frac{\Delta_y}{1-g}. \quad (11)$$

²⁷To clarify the units, the damages are measured per Teraton of CO₂ [TtonCO₂], and the 3 degrees Celsius rise follows from doubling the CO₂ stock. We have chosen the value of π such that the normalisation $\Delta_y = 1$ gives the Nordhaus case. For this reason, the interpretation of π is “climate damage sensitivity” rather than “climate sensitivity”.

²⁸These, and the derivation of the optimal climate policy rule, are provided in the Appendix using dynamic programming. Yet, the analysis in the main text is self-contained.

²⁹See also Golosov et al. (2014) or Gerlagh and Liski (2016).

Thus, output and direct utility losses are comparable in terms of their effect on utility; for convenience, we will use Δ as an aggregate measure of both losses. For the proof of the remark, consider the effect of temperature $D_{t+\tau}$ on utility in period $t + \tau$ when $I_t = 1$ (climate impacts have arrived). Recall that the consumption utility is $\ln(c_{t+\tau}) = \ln((1 - g)y_{t+\tau}) = \ln(1 - g) + \ln(y_{t+\tau})$ so that, through the exponential output loss, the consumption utility loss is given by $\partial \ln(c_{t+\tau}) / \partial D_{t+\tau} = -\Delta_y$. As there is also the direct utility loss, captured by Δ_u in (6), the full loss in utils at $t + \tau$ is

$$-\frac{du_{t+\tau}}{dD_{t+\tau}} = \Delta_y + \Delta_u.$$

But, part g of the output loss at $t + \tau$ also propagates through savings to period $t + \tau + 1$ and further to periods $t + \tau + n$ with $n > 0$, so that the full loss of utils, discounted to time t and denoted by Δ , is given by (11). In our calibration, we quantify Δ_y and set $\Delta_u = 0$, for ease of comparison with the literature.

3.1 Climate policies if impacts arrive

Saving rule g is not affected by the potential climate shock, but the climate policy rule depends on state I_t . We solve the climate policy by working “backwards”; we find the optimal policy in contingency $I_t = 1$, and then in $I_t = 0$. Thus, first set $I_t = 1$, and consider the social cost of current carbon emissions z_t , obtained from the effect of emissions at t on a stream of future utilities. The full loss of utils per increase of temperatures as measured by $D_{t+\tau}$, caused by z_t at time t , when discounted to t with factor $0 < \delta < 1$, is denoted by h . It follows with the aid of (9) and (11):

$$\begin{aligned} h &\equiv -\sum_{\tau=1}^{\infty} \delta^{\tau} \frac{du_{t+\tau}}{dz_t} \\ &= \Delta \sum_{\tau=1}^{\infty} \delta^{\tau} \frac{dD_{t+\tau}}{dz_t} = \Delta \sum_{\tau=1}^{\infty} \delta^{\tau} \mathcal{R}(\tau) \end{aligned} \tag{12}$$

$$\begin{aligned} &= \Delta \sum_{i \in \mathcal{I}} \frac{a_i \pi \varepsilon}{\varepsilon - \eta_i} \sum_{\tau=1}^{\infty} \delta^{\tau} (1 - \eta_i)^{\tau} - \delta^{\tau} (1 - \varepsilon_j)^{\tau} \\ &= \delta \Delta \pi \frac{\varepsilon}{1 - \delta(1 - \varepsilon)} \sum_{i \in \mathcal{I}} \frac{a_i}{1 - \delta(1 - \eta_i)}. \end{aligned} \tag{13}$$

The present-value utility costs of current emissions can thus be compressed to a number, h , that will be an input to the determination of the currently optimal carbon price. The first term on the right-hand side, $\delta \Delta \pi$, describes the utility loss associated with one emission unit when steady state damages would happen immediately at the next period. The second term discounts damages because of the time-delay associated with

temperature adjustment. The third term with the summation describes the persistence of damages as the atmospheric CO_2 stock decays slowly.³⁰

Remark 3 *Conditional on $I_t = 1$, the optimal carbon price is*

$$\tau_t = \frac{\partial y_t}{\partial z_t} = (1 - g)y_t \delta \Delta \pi \frac{\varepsilon}{1 - \delta(1 - \varepsilon)} \sum_{i \in \mathcal{I}} \frac{a_i}{1 - \delta(1 - \eta_i)}. \quad (14)$$

Thus, if impacts arrive, the optimal carbon price in (14) develops in lock-step with income, with the degree of proportionality depending on δ , Δ , and the climate-system parameters. This rule differs from that in Golosov et al. (2014) mainly because of the carbon cycle and loss parameter Δ that includes both utility and output losses. In particular, the same tax is optimal for any division between utility and production losses satisfying (11).

For the proof, as noted above, the payoff implications of temperature changes are separable from the payoff implications of changes in capital. Then, the climate policy can be found by looking at the present-value of future utility costs of current emissions, holding savings fixed, and balancing these future impacts with the current utility-weighted marginal product of carbon: $\frac{\partial y_t}{\partial z_t} \frac{\partial u}{\partial c} = h$. Since $\frac{\partial u_t}{\partial c_t} = 1/c_t = 1/(1 - g)y_t$, we can express the optimal carbon price as

$$\tau_t = \frac{\partial y_t}{\partial z_t} = h(1 - g)y_t \quad (15)$$

which, after using the previously derived h , gives the result.

This parametric class for preferences and technologies effectively implies a unit elasticity of losses with respect to economic stakes. There is an emerging set of papers that focus on relaxing this assumption, among other issues, in numerical evaluations of carbon taxes (Barrage, 2014; Van den Bijgaart, 2016; Rezai and Van der Ploeg, 2015). The assumption represents an intermediate position in the following sense. Some economic climate-change losses, such as decreased agricultural yields in tropical areas, are likely to increase less than one-to-one with income, as the share of the agricultural sector tends to decrease when income grows. At the same time, as these agricultural impacts are expected to be more severe in the currently warm-climate and less-developed countries, the share of damages in world-wide income will increase when those economies grow at

³⁰The last term shows that the effective discounting of the utility impacts interacts with the decay of the carbon. In particular, if all carbon reservoirs have a positive decay, the present value of impacts remains bounded even when $\delta \rightarrow 1$. But, some fraction of carbon is persistent (Section 2.3), so that the present value becomes unbounded for sufficiently low discounting. See Gerlagh and Liski (2016).

rates larger than the world-wide average growth rate. Also, the monetary evaluation of economically intangible impacts such as ecological losses are expected to increase more than proportionally with income (Mendelsohn et al., 2006; Mendelsohn et al., 2012).

3.2 Climate policies before impacts

Once damages appear, the policies can be determined exactly as in Proposition 3. Prior to their appearance, the model generates a parametric distribution for the time when damages occur. Let Z be the stochastic variable, measuring the full future utility cost from increasing current emissions z_t marginally. Let $h_t = \mathbb{E}_t Z$ be the expected present value of future utility losses associated with one unit of current emissions. Z can take the values $Z_1, Z_2, \dots$, 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 characterises the present-value marginal utility losses from 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 3.1, and using the emissions-temperature response from Section 2.3, we can obtain the present-value of such delayed utility losses in closed-form:

$$\begin{aligned} Z_\tau &= \Delta \sum_{s=\tau}^{\infty} \delta^s \mathcal{R}(s) \\ &= \Delta \sum_{i \in \mathcal{I}} \frac{\pi a_i \varepsilon}{\varepsilon - \eta_i} \delta^\tau \left(\frac{(1 - \eta_i)^\tau}{[1 - \delta(1 - \eta_i)]} - \frac{(1 - \varepsilon)^\tau}{[1 - \delta(1 - \varepsilon)]} \right). \end{aligned}$$

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 utility losses, denoted by $F_t(Z)$, we first establish the probability that the damage has revealed itself at period t , irrespective of whether t is the first time:

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

We can generalise 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 Z is then given by

$$\begin{aligned} F_t(Z_\tau) &= \Pr(Z \leq Z_\tau | I_t = 0) = \Pr(I_{t+\tau-1} = 0 | I_t = 0) \\ &= 1 - \mu_t + \mu_t(1 - \lambda)^{\tau-1}. \end{aligned} \quad (16)$$

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

Theorem 1 *Conditional on no experience of impacts by time t ($I_t = 0$), the previous-period distribution of the social cost of carbon $F_{t-1}(Z)$ 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,*

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

Proof. The expected utility losses from current emissions are equal to

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

Using our temperature-response function leads to the expression for h_t . Decreasing carbon prices measured in utils and stochastic dominance follow from (16) and μ_t decreasing over time. ■

The result gives the social cost of carbon in closed form; in particular, the cost depends on (μ_0, λ, Δ) that are the three critical parameters in our calibration and quantitative assessment. The first term on the right in the definition of h^l equals h that is the full information policy variable, defined in (13). Intuitively, the second term subtracts from h the present value of damages that in expectations do not occur, substituting $\delta(1 - \lambda)$ for the discount factor. Now, the optimal carbon price translates the utility losses to money, and equals the income-weighted future utility-cost of current actions, analogous to (14). More formally, the current utility-weighted gain from increasing emissions ($\frac{\partial y_t}{\partial z_t} u'_t$) should be equated with the *current perception* of the future marginal loss in utils ($\mu_t h^l$):

Proposition 1 *The optimal learning-adjusted carbon price is*

$$\tau_t = \mu_t h^l (1 - g) y_t. \quad (17)$$

This form for carbon pricing is the key to the main result of this paper: the tax is no longer developing in lock-step with income, as was the case with full information. Below, in the next Section, we allow the belief dependent part of the rule develop more generally so that the intensity of learning, as measured by λ , can be a function of temperature, or perhaps a function of time. But constant λ , assumed so far, is useful for understanding the mechanisms at work.³¹

Limiting cases reveal the mechanisms at work. Consider time $t = 0$, where the subjective belief of damages is given by $\mu_0 < 1$. If damages are almost surely observable, $\lambda \nearrow 1$, the optimal initial policy prior to experimentation is the full information policy, weighted with the subjective probability for damages, $h^l \rightarrow h$. However, if damages do not appear the next period, $I_1 = 0$, then the subjective assessment μ_1 drops to zero by the updating rule (8) as beliefs become very optimistic, and the carbon price drops to zero, $\mu_1 \searrow 0, h_1 = \mu_1 h^l \searrow 0$. In this case, “no news” reveals the true climate-economy state precisely. On the other hand, if climate change damages are not easily observable, $\lambda \searrow 0$, climate change is a problem with a non-significant rate of appearances in all cases and carbon prices are low, $h^l \searrow 0$. But this case also implies that climate experiments are not very informative; there will be no learning, and the subjective assessment μ_t in (8) remains almost unchanged over time.

The experience-sensitive “climate policy ramp” thus depends on the intensity of learning. Our calibration matches this intensity with the representative values used in the recent climate science and economics literature to evaluate how quickly the tax path can be expected to depart from the income path.

3.3 Carbon price growth rate

Above, learning led to carbon prices growing slower than the economy. This is consistent with the recent numerical findings presented in Lontzek et al. (2015), but not with Lemoine and Traeger (2014). As explained in Section 2.2, these two contributions illustrate two distinct views on how information about extreme events is expected to arrive.

³¹Equation (17), while simple, makes the current tax to depend on the full state of the economy through y_t . Moreover, τ_t looks only at the cost side of current emissions by giving the money-metric social cost per ton of emissions; the deepness of the cuts in emissions induced depends on further details of the energy sector (provided in Section 4.3).

Now, we generalise our specification so that the same model can cover both cases; that is, carbon prices can also grow faster than the economy.

For the idea of the generalisation, note that if the model with constant hazard rate had already been running for centuries, then the existence of damages would most likely have become known by now. So, why cannot we ignore the possibility of high impacts? The puzzle comes from the constant experimentation intensity, which we relax below. Suppose now that a climate system that varies within the natural boundaries does not lead to catastrophes. Catastrophes — and learning — take place only above a temperature threshold, $D_t \geq \bar{D}$, corresponding, for example, to 1 or 2 degrees Celsius above the pre-industrial temperature levels.³² When learning requires exceeding such thresholds, the carbon price should grow faster than the economy, as follows:

Proposition 2 *Assume that temperatures generate information on impacts only if $D_t \geq \bar{D} > 0$. Let $D_0 < \bar{D}$ and $t' < \infty$ be the first period such that $D_{t'} \geq \bar{D}$. Then, prior to t' , the expected current-value utility impact of emissions increases over time: $h_t < h_{t+1}$ for $0 < t < t'$.*

Proof. Let T be the set of periods τ such that $D_\tau \geq \bar{D}$. The expected utility losses for $0 < t < t'$ satisfy

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

³²Records show that over the last million years global temperatures have a cycle of glacials and interglacials, spanning a temperature range of about 10 degrees Celsius. The current Holocene interglacial is at the upper end of the range. Global warming is expected to bring temperatures out of the range observed over a million years.

The second line follows because $I_t = 0$ with certainty for $0 < t < t'$. The inequality follows as we subtract one period to take one period of delay away. The fifth line follows as beliefs do not change between t and $t + 1$. ■

As long as no information can be obtained, no damages will occur but policy h_t becomes more strict over time as the expected first appearance of damages comes closer. In spirit, while formally different, this mechanism captures the source of policy tightening in Lemoine and Traeger (2014).³³

Proposition 3 *For $0 < t < t'$, defined in Proposition 2, the optimal carbon tax grows faster than the economy.*

Mechanically, since the actual carbon tax is a multiple of income, the tax implied by h_t for $D_t < \bar{D}$ will be growing over time at a rate exceeding the growth of the economy, by Proposition 2. The tightening of policies continues until the temperatures start generating information; thereafter, the optimal carbon tax grows slower than the economy as in the previous Section. Further, 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 of effects from temperatures can start several decades *after* the emissions that caused climate change to break through the threshold. Thus, the shape of the emissions-temperature response, $\mathcal{R}(\tau)$ is thus not only important as a measure of the development over time for the potential shock on the economy; it also dictates how quickly the climate experiment can become informative. This is not unheard of in climate science (Roe and Baker, 2007); the instrumental value of current emissions to generate sharper estimates before 2050 on the possibility of severe climate change impacts is, by the nature of climate change, very limited.³⁴

We can generalise our formula to describe the basic learning model and the threshold model at one go. Consider that the learning intensity increases over time, along with global temperature rise; that is, λ_t becomes time-dependent.³⁵ We now rewrite the

³³In that paper, the support for the tipping point temperature is uniform so that the hazard rate for the event increases with a longer emissions history.

³⁴But, to be conservative in the quantitative analysis, we eliminate this delay in learning and allow better information to work against the climate policy ramp from the start.

³⁵For tractability, we do not write λ_t as explicitly dependent on the temperature change. A fully specified temperature-dependent experimentation intensity requires further conceptual steps; see the discussion in Section 2.2 and fn. 23.

Bayesian updating of the beliefs, in equation (8), through

$$\begin{aligned}\mu_t &= \Pr(\lambda_t > 0 | I_t = 0) \\ &= \frac{(1 - \lambda_t)\mu_{t-1}}{1 - \lambda_t\mu_{t-1}}\end{aligned}$$

so that the full utility impact described in Theorem 1 becomes

$$h_t = \mu_t \Delta[\sum_{\tau=1}^{\infty} \delta^{\tau} \mathcal{R}(\tau) - \sum_{\tau=1}^{\infty} \delta^{\tau} \mathcal{R}(\tau) \prod_{s=1}^{\tau} (1 - \lambda_{t+s})].$$

The equation describes both the basic learning model and the learning threshold model. For λ_t increasing over time, there are two opposing effects. First, μ_t declines over time as long as $\lambda_t > 0$, pressing down variable h_t . Second, the most-right variable within the product increases with t , and as it enters negatively, this will increase h_t . The above propositions inform us that for λ_t almost constant, the first effect dominates so that carbon policy h_t becomes less stringent over time, while for sharply increasing λ_t the policy will tend to become more demanding over time, with carbon prices rising faster than income.

3.4 The effect of discounting

Theorem 1 decomposes the utility impacts of emissions into a post-learning carbon price and the net present value of impacts in the early periods, in which a catastrophe is relatively unlikely to happen. In this sense, there is a delay in the time structure of expected damages. The economic consequence of this feature is that carbon prices based on uncertain catastrophes will be more sensitive to discounting, compared to carbon prices based on a certain damage profile. In technical terms, the carbon price is log-convex in the discount factor. We state this result in Proposition 4 below. The proposition is conditional on a discount factor that is not too low. If δ is less than a half, then the planner values the entire future less than twice the immediate next period, and the delay structure of climate damages, which extends to multiple periods, is less important.

We write h_{λ}^l for h^l defined in Theorem 1, adding subscript λ to differentiate between the uncertain social cost, $0 < \lambda < 1$, and the deterministic cost, which can equivalently be labeled as $\lambda = 1$.

Proposition 4 *The uncertain social cost is more sensitive to discounting than the deterministic carbon price:*

$$\frac{1}{h_{\lambda}^l} \frac{\partial h_{\lambda}^l}{\partial \delta} > \frac{1}{h_{\lambda=1}^l} \frac{\partial h_{\lambda=1}^l}{\partial \delta}$$

if $\delta > \frac{1}{2}$.

The proposition is intuitively clear, but the proof requires some tedious algebra (see the Appendix). Our model captures the caution due to delayed discrete events, and such caution increases with lower time discounting. To the best of our knowledge, this result has not been noted in the literature. The result is different from but in spirit similar to that in Traeger (2015) who finds for a continuous-state model that the belief updates act as persistent shocks, making the policies sensitive to pure time discounting.

4 Quantitative assessment

4.1 Calibration

We assess quantitatively if the carbon tax path, the climate policy ramp, is sensitive to the assumption that climate-change impacts are unknown, potentially extreme events. We calibrate the model so that it remains comparable to the mainstream quantitative models that assume immediate and moderate damages, if we assume such damages in our model. Setting $(\mu_0, \lambda, \Delta) = (1, 1, 1)$, losses are immediate and moderate. In particular, $\Delta = 1$ refers to Nordhaus (2008) baseline where a temperature rise of 3 degrees Celsius leads to about 2.7 per cent loss of output (Section 2.3). Let us first hold on to the idea of immediate and moderate damages, $(\mu_0, \lambda, \Delta) = (1, 1, 1)$, to facilitate the introduction of other than learning-related parameters.

Reasonable savings are in the range 20-30 per cent of output ($.2 < g \leq .3$); the capital share of output should be close to 30 per cent, $\alpha = .3$. We work with 10-year periods. Then, by the optimal savings rule $g = \alpha\delta$, the annual pure rate of time preference ranges between one and three percent, that is, δ per decade ranges between .74 and .9, respectively, to keep savings in the range above. To compensate for the fact that log utility implies low preferences for consumption smoothing and thus tends to make the decision maker “patient” under income growth, we do not want to choose the lowest time discount rate as our main case. We set the annual discount rate to .02 so that $\delta = .82$ and $g = .25$. By the Ramsey rule, this comes close to the Nordhaus DICE (2008) choices where the elasticity of marginal utility is two and the time discount rate equals one per cent.³⁶

We can now immediately calculate the optimal initial tax for moderate damages $(\mu_0, \lambda, \Delta) = (1, 1, 1)$, and compare the tax with Nordhaus’ number. If we take the Gross

³⁶We demonstrate in the Appendix, in Fig. 6, that the properties of the DICE policy path are preserved under the utility function transformation. See also fn. 42.

Global Product as 600 Trillion Euro [*Teuro*] for the decade, 2006-2015 (World Bank, using PPP), we obtain that the tax should be 10 EUR/tCO₂, close to the number appearing in Nordhaus (2007), that is 35*USD/tC*.³⁷ This gives confidence that the benchmark output of the model is in the ballpark.³⁸

The main analysis is about contrasting the policy outcomes under two scenarios, one with smooth certain damages, the other with uncertain but large damages. We obtain the policy outcome for high-consequence damages through an explorative calibration of (μ_t, λ, Δ) , based on the recent quantitative literature on the effect of tipping points on climate policies. Then, this outcome is compared with the benchmark where damages are immediate and moderate. The parametric representation of the economy remains in all other aspects the same.

In recent years, several studies have presented quantitative estimates for the effect of catastrophic irreversible events in the climate system on the social costs of carbon.³⁹ We seek to find representative values for (μ_0, λ, Δ) , based on this literature.

The most critical parameter is the learning intensity, λ . In the literature, the decadal learning rate, as captured by our λ , ranges between .02 and .1.⁴⁰ We are interested in a conservative test that allows learning to work against the upward pressure from income

³⁷Note that 1 tCO₂ = 3.67 tC.

³⁸In the Appendix, Fig. 7, we confirm that the entire tax time paths of the model and the DICE are within 10 per cent of each other for the next century.

³⁹Candidates for such events are changes in the Atlantic Gulf Stream, melting of the Greenland Ice Sheet, the collapse of the West Antarctic Ice Sheet, the dieback of the Amazon rainforest; or an increase in the amplitude of the El Niño Southern Oscillation (Kriegler et al. 2009).

⁴⁰Lemoine and Traeger (2014) model tipping points as irreversible shifts in the system that occur upon crossing an unknown temperature threshold. The tipping points are uniformly distributed over a range of potential temperature increases, so that the hazard rate at which the tipping point is crossed becomes proportional to the speed of temperature increase. The current global temperature increase of about 0.2 Kelvin per decade (in their model) leads to about a 5 per cent per decade probability of passing the threshold. This would suggest that $\lambda = .05$ per decade. Lontzek et al. (2015), describe the appearances of climate catastrophes through a state-dependent process, without invoking a threshold explicitly. In their model, the expected global warming path leads to a hazard rate that increases over time, from about 0.25 to 1 per cent per year, and a cumulative probability of about 50 per cent for such major event by 2200. Cai et al. (2013) and Cai et al. (2015a) also deploy a state-dependent process, but they also assume a 1 Kelvin global warming threshold, below which the climate system is stable. Above 1 Kelvin, the climate system can shift into a bad state, with the hazard rate increasing proportionally with the global temperature anomaly. In a world with 2 Kelvin global warming, the decadal probability for tipping into the bad state is about 1.5 to 2 per cent. Translated to decadal rates, this literature suggest arrival rates between .02 and .1.

growth in the climate policy ramp; therefore, we assume relatively high learning rate from the set of values in the literature but yet not extreme. We set $\lambda = .08$ per decade in our main case. A geometric distribution with this arrival rate per decade means that the expected arrival time for a severe climate change damage event is about 130 years.

As to the initial belief if the event ultimately occurs, μ_0 , the literature provides only indirect guidance. As explained in Section 2.2, we can impute from Lemoine and Traeger (2014) that $\mu_0 = .8$ roughly consistent with their assumptions.⁴¹ We thus assume that there is an initial assessment of 80 per cent chance for the event to finally arrive. Although the other studies, such as Lontzek et al. (2015), do not allow us to recover μ_0 directly, choices $\lambda = .08, \mu_0 = .8$ mean that the prior decadal hazard rate, $\mu_0\lambda$ is 6.4 per cent, in the middle of the assumed range in Lotzek et al. (2015). With 80 per cent initial assessment and 8 per cent arrival probability per decade, after 100 years without damages, the posterior for the eventual impact arrival μ_t is still 64 per cent. Thus, the parameter choices imply that the beliefs are relatively persistent.

Finally, we must specify the economic stakes if the event takes place. Looking at the literature cited above, climate change damages associated with major climate regime changes are up to 10 times as large as estimates for damages associated with immediate and moderate impacts, so that Δ falls in the range 1 to 10. We set $\Delta_y = 4$. The maximum damages are by factor four higher than the middle-of-the-road damages assumed in Nordhaus (2008) — the implied output loss is then about 10.7 per cent from doubling the CO_2 stock. The choice of Δ is important; for example, it determines if the economy moves towards clean technologies before learning the climate impacts (see Section 4.3). The choice with $\Delta_y = 4$ is a useful starting point, as the initial level of the carbon price almost coincides with the one under immediate damages and $\Delta_y = 1$. Then, as analysed in the next Section, we can see for how long the uncertain tipping-point carbon price stays close to the policy ramp under moderate but certain damages; the conclusions for the shape of the policy path do not depend on the precise value of Δ_y .

4.2 The climate-policy ramp

Consider now the development of the optimal carbon price over time, with the parameter choices above. Given the closed-form formula for the carbon price in (17), one approach is to conjecture future output or income levels, say, in 2050 and, conditional on no observed

⁴¹The overall tipping probability in Lemoine and Traeger depends on the starting level of the temperatures; we can justify μ_0 between .7 and .8.

impacts by that time, obtain the future carbon price for that state of the world. However, future states of the world result partly from past policy decisions; carbon pricing decisions have an effect on the current output, and through investments, also on future incomes. For consistent policy scenarios, in Section 4.3 just below, we specify and calibrate a structure for the energy sector and total productivities through $A_t(z_t)$ in the production function. The scenarios presented here are based on this specification.

The benchmark for our assessment is the “Climate policy ramp” (dotted line in Fig. 2), based on Nordhaus’ DICE (2007) middle-of-the-road damage estimate, corresponding to $\Delta_y = 1$; that is, damages are known and smooth. For the first period 2020, capturing decade 2015-2025, the benchmark sure-loss policy path gives ca. 14 *EUR/tCO₂* as the optimal price which is almost identical to what DICE produces for this period under this choice for discounting and preferences.⁴² This middle-of-the-road sure-loss path involves a tightening of the policies over the coming century, typical for most no-uncertainty climate-policy assessments.

We now look at the optimal time path for the carbon price for *high potential damages, but conditional on not observing these damages*; that is, we consider the evolution of the policy when future impacts are potentially severe, $\Delta_y = 4$, as determined by the calibration procedure above, but when no news on climate impacts arrive. Without impacts, the economy is unaffected by climate change.⁴³ The optimal carbon price is depicted as a solid line in Fig. 2 over the coming century and beyond. The two climate policies — one with immediate damages based on the central estimate, and the other with high but only potential damages and gradual updating of beliefs to the no-news situation — have the same shape for the first century. The difference between the two is only 13 per cent by the end of the century. Strikingly, for this particular learning scenario, it takes close to 200 years without observed climate damages for beliefs to become optimistic enough for the carbon price to decline. The social cost of carbon declines very slowly.⁴⁴

⁴² We demonstrate this in detail in the online Appendix. First, in Fig. 6 we show that the DICE utility function can be changed to log-utility, while adjusting suitably the discount rate, without much affecting the baseline climate policy ramp. Second, after the first step, we can trust that the properties of DICE are preserved under the utility function transformation, and thus we produce the climate policy ramp with DICE for our specification (log utility and 2 per cent time discounting). In Fig. 7, we show that our model matches with DICE policy ramp very well, with the same parameters. The striking similarity of the two paths follows since our carbon cycle and temperature delay dynamics come close to those in DICE.

⁴³ But, since the carbon policies are in place, emissions and output will be reduced below the business-as-usual path.

⁴⁴ In the Appendix, we report the bad news carbon price for the same time horizon: for time t , the bad

Figure 2: The carbon price for a sure income loss of 2.7 per cent from doubling the carbon concentration ($\mu = 1; \Delta_y = 1$), and for uncertain damages, conditional on no news on damages ($\mu_0 = 0.8; \lambda = 0.08, \Delta_y = 4$).

To assess the shape of the carbon price path, we decompose its level into its two main components. Recall that the optimal carbon price is proportional to h_t capturing the expected utility losses from current emissions, and to income y_t ; $\tau_t = h_t(1 - g)y_t$, $h_t = \mu_t h^l$. See Table 1, for the contribution of income (y_t) and learning (μ_t) to the carbon price.⁴⁵ Expected income growth is prodigious; in our evaluation, based on the IPCC scenarios (see Section 4.3 below), income rises five-fold during the coming century. Such an estimate is not unheard of, and is driven by an increasing population and the rise of the middle class in emerging economies. The development of beliefs is captured through μ_t in the Table. Observing no major climate damages over the coming century, leads to substantial increase in optimism, but, as is evident from the Table, it is the changing scale of the global economy that matters for the development of carbon pricing. Stabilising carbon prices at the initial level — thus ruling out a climate policy ramp completely — would require that the climate experiment is by orders of magnitude more informative than considered here. The assessment of the probability of major utility

news price is the optimal price conditional on observing the damage in period t . For the first century, it is about four times the pre-learning optimal price; by the end of the second century, the gap between the two prices is 50 per cent larger. The bad news price measures how severely the economy is hit if the bad news arrive.

⁴⁵It is illuminating to consider the units of measurement for the utility loss measure $h_t = \mu_t h^l$, which has the same unit as the constant in the legend of the table: years per emissions. The variable h_t measures the life-time equivalent of welfare that is lost per unit of emissions. For the year 2010, annual emissions are about .04 TtonCO₂, implying $.75 \times .04 = .03$ years of expected life-time destroyed by these emissions.

losses, as captured by μ_t , would need to decline by 50 per cent by 2050. For the belief updating to achieve this, the calibrated λ we would need to be by factor four larger; it would require a large deviation from the parameter choices used in the literature, and probably also a very different climate-system description.

	income [T€/yr]	beliefs, μ_t [.]	carbon price [€/tCO ₂]
2020	78	.79	14
2050	146	.74	25
2100	304	.65	46
2150	510	.55	67
2200	703	.45	76

Table 1: Decomposing the contribution of income and learning to the carbon price. Multiplying the first column and the second column, with a constant $h^l(1 - g) = 0.68$ [$yr/TtCO_2$], gives the last column.

4.3 Decarbonisation

We come to the end of the quantitative assessment by looking into the specific structure for the economy’s production function (5), used for the quantitative conclusions above. The structure of production allows us to address one final and important question: will the perception of the social cost, as quantified above, lead to the decarbonisation of the economy, that is, will emissions ultimately decline to zero, if the carbon price is implemented as a carbon tax?⁴⁶ We also perform sensitivity analysis on how severe the potential impact on the economy, as captured by Δ_y , would have to be for the decarbonisation to take place.

The specification is set up for a transparent calibration of the total and energy sector

⁴⁶We will not explicitly consider a decentralised economy with tax policies in place; however, it follows from the first principles that the planning problem considered in this paper can be decentralised.

productivities. We assume

$$y_t = k_t^\alpha [\mathcal{A}_t(l_{y,t}, e_t)]^{1-\alpha} \exp(-\Delta_{y,t} D_t) \quad (18)$$

$$\mathcal{A}_t(l_{y,t}, e_t) = \min \{A_{y,t} l_{y,t}, A_{e,t} e_t\} \quad (19)$$

$$e_t = e_{f,t} + e_{n,t} \quad (20)$$

$$e_{f,t} = \min \{A_{f,t} l_{f,t}, B_t z_t\} \quad (21)$$

$$e_{n,t} = \frac{\varphi + 1}{\varphi} (A_{n,t} l_{n,t})^{\frac{\varphi}{\varphi+1}} \quad (22)$$

$$l_t = l_{f,t} + l_{n,t} + l_{y,t}. \quad (23)$$

There are time-trends for total labor l_t , and for labor productivities (A_y, A_e, A_f, A_n) in output, total energy, fossil-fuel energy, and non-carbon energy production, respectively. Total energy, e_t , depends effectively only on labor allocation at time t : the core allocation problem in the energy sector 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 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, $e_t = e_{f,t} + e_{n,t}$. The allocation outcome depends only on time and carbon inputs; labor-energy composite $[\mathcal{A}_t(l_{y,t}, e_t)]$ then defines the total factor productivity term $A_t(z_t)$ in the economy's production function (5).

Labor-energy composite $\mathcal{A}_t(l_{y,t}, e_t)$ takes a Leontief form capturing an extremely low elasticity of substitution between labor in the final-good sector $l_{y,t}$ and energy e_t . By this assumption, we avoid unrealistically deep early reductions of emissions through substitution of labor inputs, and thereby approximate the energy sector capital adjustment delays; see also Hassler, Krusell and Olovsson (2012).⁴⁷

In (21), we assume that $e_{f,t}$ can be produced with a constant-returns to scale technology using labor $l_{f,t}$ and the fossil-fuel z_t , where $A_{f,t}$ and B_t describe productivities. The fuel resource is not a fixed factor and commands no resource rent; as in Golosov et al. (2014), the fossil-fuel resource is in principle unlimited. In contrast, in equation (22), where $\varphi > 0$ describes the elasticity of supply from the non-carbon sector; the non-fossil fuel energy production is land-intensive and subject to diminishing returns and land rents (as in Fischer and Newell, 2008).

Without carbon policy, $\tau = 0$, the labor market allocation can be solved in closed

⁴⁷A cost of the assumption is that the output-energy intensity remains fixed so that energy savings cannot arise as a source of emissions reductions; the reduction path for emissions in Fig. 3 is achieved through decarbonisation, that is, substitution away from carbon energy.

form; thus, we can invert the model to map from quantities $(l, y, e_f, e_n)_{t \geq 0}$ to productivities $(A_y, A_e, A_f, A_n)_{t \geq 0}$ (see Appendix for the solution, and the supplementary material for the quantitative values). To express all energy in carbon units, we set $B_t = 1$, leaving us three distinct energy productivities (A_e, A_f, A_n) . We match the business-as-usual (BAU) quantities $(y, e_f, e_n)_{t \geq 0}$ with the A1F1 SRES scenario from the IPCC (2007). Population follows a logistic growth curve based on World Bank forecasts. Population in 2010 is set at 6.9 [billion], while the maximum population growth rate is chosen such that in 2010 the effective population growth rate per decade equals 0.12 [/decade]. The maximum expected population (reached at about 2200) is set at 11 [billion]. Using these calibrated productivity trends, we produce the adjustment path of the economy for the optimal policies in Figs 2-5.

In Fig. 3 we depict the development of the energy sector when the economy does not experience climate-change impacts but faces the optimal no news tax as shown in Fig. 2. All energy is measured in CO_2 -equivalents; then, “carbon energy” gives directly the carbon dioxide emissions per decade. The economy, when optimally planned, does not achieve full decarbonisation targets: the share of non-carbon energy of the total increases over time, reaching 37 per cent by 2100 and 60 percent by 2200. Such a persistent share for carbon energy leads to carbon concentrations about 600 ppmv by the end of the current century, well beyond safe targets recommended by the IPCC.⁴⁸ Stated differently, output loss $\Delta_y = 4$, that is, 10.7 per cent from doubling the CO_2 stock, together with increasing optimism about the absence of final damages, is too low to justify optimal full decarbonisation.

For comparison, Fig. 4 provides the same evaluation of the energy sector development but now for the middle-of-the-road sure-loss climate impact, following Nordhaus’ (2007) with the carbon price depicted in Fig. 2 ($\mu = 1$; $\Delta_y = 1$). Not surprisingly, because the carbon prices for known and unknown damages are almost equivalent in Fig. 2, the decarbonisation path with smooth damages is not considerably different from that under known damages for the first century. However, with smooth damages, the carbon price keeps on growing with the economy and therefore the decarbonisation finally takes place by the end of the second century.

The 2015 United Nations Climate Change Conference negotiated an agreement that calls for zero carbon emissions by the end of the current century. We can now quantify the economic losses that justify the decarbonisation target. We return to the model

⁴⁸The supplementary material, under the link https://www.dropbox.com/sh/7meos655j14jh5p/_dlr8X_FHI, reports the climate-related implications of the scenarios.

Figure 3: Partial decarbonisation without observed climate impacts, $\Delta_y = 4$. Energy per decade measured in Teratons of CO_2 equivalents.

where losses are unknown and ask how large should the losses be for the decarbonisation to take place by the end of the century? In Fig. 5 shows the case with $\Delta_y = 10$, holding all else equal for the economy. We observe that the non-carbon energy increases its share to 100 per cent by the end of the century. Clearly, the precise timing is sensitive to the assumed substitution possibilities between carbon and non-carbon energy. Yet, on reflection, the decarbonisation without impacts is not unreasonable; the optimal carbon price levels imply substantial value providing incentives for the transition. Current CO_2 emissions exceed 30 Gigatons annually, while the annual world output is about 60 trillion euro. A worldwide carbon price of 100 EUR/tCO_2 then represents about 5% of the value of the output, giving a ballpark idea of the decarbonisation incentives. Most climate-economy models, including Nordhaus DICE-2007, produce decarbonisation at such price levels that we, in our model with $\Delta_y = 10$, reach by 2080, explaining the decarbonisation in Fig. 5. Under smooth damages benchmark ($\Delta_y = 1$), similar price levels are reached a century later, consistent with the decarbonisation in Fig. 4.⁴⁹

⁴⁹Recall that our energy sector model assumes a low factor substitutability to approximate an explicit capital structure adjustments. The approximation entails the following feature: carbon energy would make a come-back if the future social cost of carbon sufficiently declines. Such come-backs happen after the time horizon of our main interest, and they are best analysed with an explicit capital structure in the energy sector.

Figure 4: Decarbonisation with observed climate impacts for a sure income loss of 2.7 per cent from doubling the carbon concentration ($\mu = 1$; $\Delta_y = 1$). Energy per decade measured in Teratons of CO_2 equivalents.

5 Concluding Remarks

We have developed a tractable climate-economy model that allows a stylised but transparent and self-contained quantitative assessment of the optimal carbon price when the impacts of climate change are unknown and can be learned only gradually over time. Such climate change unknowns are often used to motivate a price for carbon, but some argue that they also render mainstream climate-economy modelling useless (Pindyck, 2013) because the mainstream models assume moderate and known damages from climate change. We believe that thoughtful climate-economy models can provide a very useful structure for quantifying what kind of events are unknown enough to justify an immediate action, and, in contrast, when it is better to adopt a wait-and-see strategy.

There are two forces that distinguish the temporal evolution of the carbon price in a regime shift model from that in a smooth damage model. First, the potential regime shift becomes increasingly more relevant to the policy maker when the economy approaches temperature levels at which the regime shift can potentially occur, provided current temperatures cannot yet trigger such events. This effect makes the carbon tax to increase faster than under smooth damages. Second, at temperatures levels that are high enough for experimentation, the event becomes less likely over time (conditional on no occurrence) and the optimal carbon tax falls relative to a smooth damage model. This is the learning effect.

Figure 5: Decarbonisation without observed climate impacts, $\Delta_y = 10$. Energy per decade measured in Teratons of CO_2 equivalents.

In our quantitative assessment, we allowed the regime shift to occur at all temperature levels: the learning effect works against the climate policy ramp from the start. Yet, we found that the high-risk carbon price path need not be that different from the mainstream policy ramp. Both the shape and level of the carbon tax very much follows Nordhaus (2008), when the information-related parameters are matched with those used in the recent numerical studies, and the size of the potential impact on the economy is intermediate. The learning effect is relatively slow in reducing the optimal carbon tax. This, then, leaves the economic development as the main driver of the policy ramp; the carbon price is developing in almost lock-step with the size of the economy in the coming decades.

The result that the level of the carbon price equals that obtained by the smooth-damages models is, however, not robust. The size of the potential impact on the economy is pure speculation. It is here where the model can become a useful tool for translating the economic meaning of the climate policy targets: how big should the economic loss from the potential event be to justify the full decarbonisation of the economy? We found that high-consequence damages that are by factor four larger than typical middle-of-the-road impacts cannot justify full decarbonisation. The impacts must be by a factor 10 larger than the middle-of-the-road estimates for the optimal greening of the economy to take place by the end of the century.

For wider implications for future research, it could be interesting to explore how countries' climate policies are shaped by own and other countries' experiences of climate impacts through their effect on beliefs regarding the country-specific ultimate impacts.

In general, it is well-received that past experiences affect current policy-decisions through beliefs. There are long traditions for beliefs and learning in the strand of literature on the design of optimal monetary policy; however, the issue of incorporating beliefs as determinants of the key macroeconomic choices together with a quantitative assessment has been recently developing (see, e.g. Buera et al. 2011). The topic seems relevant in the context of climate change: policy makers are domestically motivated but are certainly learning from experiences elsewhere.

References

- [1] Acemoglu D., Aghion P., Bursztyn L., Hemous D (2012). The environment and directed technical change, *American Economic Review* 102, 131–166.
- [2] Barrage, L. (2014). Sensitivity Analysis for Golosov, Hassler, Krusell, and Tsyvinski (2014): Optimal Taxes on Fossil Fuel in General Equilibrium. Supplementary Material, *Econometrica*, 82(1): 41-88.
- [3] Bonatti, A., and J. Hörner (2011). Collaborating, *American Economic Review*, 101(2): 632-63.
- [4] Brock, W. A., and L. J. Mirman (1972). Optimal economic growth and uncertainty: The discounted case, *Journal of Economic Theory*, vol. 4(3), pages 479-513, June.
- [5] Buera, F.J, A. Monge-Naranjo, and G. E. Primiceri (2011). Learning the Wealth of Nations, *Econometrica* 79(1), 1-45.
- [6] Cai, Y, K. L. Judd, and T. S. Lontzek (2013). The social cost of stochastic and irreversible climate change, NBER working paper 18704.
- [7] Cai, Y., K.L. Judd, T.M. Lenton, T.S. Lontzek, D. Narita (2015a). Environmental tipping points significantly affect the cost-benefit assessment of climate policies, *PNAS* 112:4606-4611.
- [8] Cai, Y., K.L. Judd, T.S. Lontzek (2015b). The Social Cost of Carbon with Economic and Climate Risks, *arXiv:1504.06909v2*.
- [9] Clarke, H.R. and W.J. Reed (1994). Consumption/pollution tradeoffs in an environment vulnerable to pollution-related catastrophic collapse, *Journal of Economic Dynamics and Control*, 18, 991- 1010.

- [10] Deschenes, O., and M. Greenstone (2007). The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather. *American Economic Review*, 97(1): 354-385.
- [11] Dell, M., B. F. Jones, and B. A. Olken 2012. "Temperature Shocks and Economic Growth: Evidence from the Last Half Century." *American Economic Journal: Macroeconomics*, 4(3): 66-95.
- [12] Fischer C. and R. Newell (2008). Environmental and technology policies for climate mitigation, *Journal of Environmental and Economic Management* 55:142-162.
- [13] Gerlagh, R. (2011). Too much oil, *CESifo Economic Studies*, Vol. 57, 1/2011, 79-102.
- [14] Gerlagh, R., and M. Liski, Climate policies with non-constant discounting (2016). Available at <http://hse-econ.fi/liski/> Longer working paper: Carbon Prices for the Next Thousand Years (June 28, 2012). CESifo Working Paper Series No. 3855.
- [15] Gerlagh, R., S. Kverndokk, and K. E. Rosendahl (2009), Optimal Timing of Climate Change Policy: Interaction between Carbon Taxes and Innovation Externalities. *Environmental and Resource Economics* 43(3): 369-390.
- [16] Greenstone, M., E. Kopits, A. Wolverton (2013). Estimating the Social Cost of Carbon for Use in U.S. Federal Rulemakings: A Summary and Interpretation, *Review of Environmental Economics and Policy* , 7 (1): 23-46.
- [17] Golosov, M., J. Hassler, P. Krusell, A. Tsyvinski (2014). Optimal taxes on fossil fuel in general equilibrium, *Econometrica*, Vol. 82 (1), 41- 88.
- [18] Hasselmann, K., S. Hasselmann, R Giering, V Ocana, H v Storch (1997). Sensitivity of optimal CO₂ emissions paths using a simplified structural integrated assessment model (SIAM), *Climatic Change* 37: 345-387
- [19] Hassler, J., P. Krusell, and C. Olovsson (2012). Energy-saving technical change, NBER Working Paper 18456.
- [20] Heal, G., A. Millner (2014). Reflections: Uncertainty and Decision Making in Climate Change Economics *Rev Environ Econ Policy* (Winter 2014) 8 (1): 120-137.
- [21] Hoel, M., and T. Sterner, Discounting and relative prices (2007). *Climatic Change* 84: 265-280.

- [22] 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
- [23] 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. *Tellus* 55B(2):378-390.
- [24] 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.
- [25] IPCC (2007). Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M. Tignor and H.L. Miller (eds.) Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.
- [26] IPCC (2014). "Climate Change 2014: Impacts, Adaptation, and Vulnerability, Summary for Policy Makers", WGII AR5 Summary for Policymakers.
- [27] Jensen, S., and Traeger, C. (2014). Optimal climate change mitigation under long-term growth uncertainty: Stochastic integrated assessment and analytic findings, *European Economic Review*, Volume 69, July 2014, Pages 104-125.
- [28] Jensen, S., and Traeger, C. (2013). Optimally Climate Sensitive Policy under Uncertainty and Learning, Department of Agricultural & Resource Economics, UC Berkeley.
- [29] Joos, F. and Roth, R. and Fuglestad, J. S. and Peters, G. P. and Enting, I. G. and von Bloh, W. and Brovkin, V. and Burke, E. J. and Eby, M. and Edwards, N. R. and Friedrich, T. and Frölicher, T. L. and Halloran, P. R. and Holden, P. B. and Jones, C. and Kleinen, T. and Mackenzie, F. and Matsumoto, K. and Meinshausen, M. and Plattner, G.-K. and Reisinger, A. and Segschneider, J. and Shaffer, G. and Steinacher, M. and Strassmann, K. and Tanaka, K. and Timmermann, A. and Weaver, A. J. (2013), Carbon dioxide and climate impulse response functions for the computation of greenhouse gas metrics: a multi-model analysis, *Atmospheric Chemistry and Physics Discussions*, 13, 2793—2825.

- [30] Keller, G., and S. Rady, and M. Cripps (2005). Strategic experimentation with exponential bandits, *Econometrica*, Vol. 73, No. 1, 39-68.
- [31] Kelly, D.L., and C.D. Kolstad (1999). Bayesian learning, growth, and pollution, *Journal of economic dynamics and control* 23, 491-518.
- [32] Kelly, D.L., and Z. Tan (2015). Learning and Climate Feedbacks: Optimal Climate Insurance and Fat Tails, *Journal of Environmental Economics and Management* Volume 72 (July) Pages 98-122.
- [33] Kriegler, 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.
- [34] Lange, A., and N. Treich (2008). Uncertainty, learning and ambiguity in economic models on climate policy: some classical results and new directions, *Climatic Change* (2008) 89:7-21.
- [35] Leach, A. J. (2007). The climate change learning curve, *Journal of Economic Dynamics and Control* 31, 1728 -1752.
- [36] Lemoine, D. M. and Traeger, C. P. (2014). Watch Your Step: Optimal Policy in a Tipping Climate, *American Economic Journal: Economic Policy* 6(1): 137-166.
- [37] Liski, M., and F. Salanié (2015). Stabilisation and the control of catastrophes, mimeo.
- [38] Lontzek, T.S., Y. Cai, K.L. Judd, and T.M. Lenton (2015). Stochastic integrated assessment of climate tipping points indicates the need for strict climate policy, *Nature Climate Change* 5: 441-444.
- [39] Malueg, D., and S. O. Tsutsui (1997). Dynamic R&D Competition with Learning, *The RAND Journal of Economics* , Vol. 28, No. 4, pp. 751-772.
- [40] Maier-Reimer E. and K.Hasselmann (1987). Transport and storage of CO₂ in the ocean - an inorganic ocean-circulation carbon cycle model, *Climate Dynamics* 2, 63-90.
- [41] Mendelsohn R., A. Dinar and L. Williams (2006). The distributional impact of climate change on rich and poor countries, *Environment and Development Economics* 11: 159-178.

- [42] Mendelsohn R. K. Emanuel, S. Chonabayashi, and L. Bakkensen (2012). The impact of climate change on global tropical cyclone damage, *Nature Climate Change* 2: 205-209.
- [43] Millner, A. (2013). On welfare frameworks and catastrophic climate risks, *Journal of Environmental Economics and Management*, Volume 65, Issue 2, Pages 310- 325.
- [44] Nordhaus, W. D. (2007). A Review of The Stern Review on the Economics of Climate Change, *Journal of Economic Literature*, 45 (3), 686-702.
- [45] Nordhaus, W. D. (2008). A Question of Balance: Weighing the Options on Global Warming Policies. (Yale University Press, New Haven, CT).
- [46] Nordhaus, W. D. (2010). Economic Policy in the Face of Severe Tail Events, *Journal of Public Economic Theory*, 14 (2), 2012, pp. 197- 219.
- [47] Piketty, T., and G. Zucman (2014). Capital is Back: Wealth-Income Ratios in Rich Countries, 1700-2010, *The Quarterly Journal of Economics* 129 (3): 1255-1310.
- [48] Pindyck, R. S. (2013). Climate Change Policy: What Do the Models Tell Us?, *Journal of Economic Literature*, 51(3): 860-72.
- [49] Polasky, S., A.J. de Zeeuw and F. Wagener (2011). Optimal management with potential regime shifts, *Journal of Environmental Economics and Management*, 62, 2, 229-240.
- [50] Rezai, A. and R. van der Ploeg (2015), Robustness of a Simple Rule for the Social Cost of Carbon, CESifo Working Paper Series No. 4703.
- [51] Roe, G. H., and M.B. Baker (2007). Why Is Climate Sensitivity So Unpredictable? *Science* 26 October 2007: Vol. 318 no. 5850 pp. 629-632.
- [52] Sakamoto, H. (2014). Dynamic resource management under the risk of regime shifts, *Journal of Environmental Economics and Management* 68: 1-19.
- [53] Sargent, T. J. (1987). *Dynamic Macroeconomic Theory*. Harvard University Press.
- [54] Schlenker, W., and M. J. Roberts (2009). Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change, *Proc Natl Acad Sci U S A*. 2009 September 15; 106(37): 15594- 15598.

- [55] 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.
- [56] Stern, N. (2006). "The economics of climate change: the Stern review", Cambridge, UK: Cambridge University Press.
- [57] Tsur, Y., Zemel, A. (1995). Uncertainty and irreversibility in groundwater management. *Journal of Environmental Economics and Management* 29, 149-161.
- [58] Tsur, Y. and A. Zemel (1996). Accounting for global warming risks: resource management under event uncertainty, *Journal of Economic Dynamics & Control*, 20, pp. 1289-1305.
- [59] Traeger, C. P. (2014). Why uncertainty matters: discounting under intertemporal risk aversion and ambiguity, *Economic Theory*, 56, 627-664.
- [60] Traeger, C. P. (2015). *Analytic Integrated Assessment and Uncertainty* (September 30, 2015), SSRN.
- [61] Tsur, Y. and A. Zemel (1996). Accounting for global warming risks: resource management under event uncertainty, *Journal of Economic Dynamics and Control*, 20, 1289-1305.
- [62] van den Bijgaart, I. R. Gerlagh, and M. Liski (2016). A simple formula for the social cost of carbon, *Journal of Environmental Economics and Management* 77, 75-94.
- [63] van der Ploeg, R., and C. Withaagen (2012). Is there really a green paradox?, *Journal of Environmental Economics and Management* 64, 342-363.
- [64] van der Zwaan, B. C. C., R. Gerlagh, G. Klaassen, and L. Schrattenholzer (2002). Endogenous Technological Change in Climate Change Modelling. *Energy Economics* 24(1): 1-19.
- [65] Weitzman, M. (2009). On modelling and interpreting the economics of catastrophic climate change, *The Review of Economics and Statistics* 91(1), 1-19.
- [66] Weitzman, M. (2011). Additive Damages, Fat-Tailed Climate Dynamics, and Uncertain Discounting, Chapter in NBER book *The Economics of Climate Change: Adaptations Past and Present* (2011), Gary D. Libecap and Richard H. Steckel, editors (p. 23 - 46).

- [67] Weitzman, M. (2013). Tail-Hedge Discounting and the Social Cost of Carbon. *Journal of Economic Literature*, 51(3): 873-82.
- [68] Weyant, J.P, F. C. de la Chesnaye, and G. J. Blanford (2006). Overview of EMF-21: Multigas Mitigation and Climate Policy, *The Energy Journal*, Multi-Greenhouse Gas Mitigation and Climate Policy Special Issue.

ONLINE APPENDIX

Solution to the detailed energy-sector model

The online supplementary file contains a program for reproducing the graphs in the text, https://www.dropbox.com/sh/7meos655j14jh5p/_dlr8X_FHI. The labor allocation is numerically obtained as follows. The allocation can be solved period-by-period taking the (i) productivity parameters, (ii) total labor, (iii) savings g , and (iv) carbon policies h_t as given. We drop the time subscript in the variables:

1. We normalise prices for the final good to equalise marginal utility, so that factor prices can be interpreted as marginal welfare per factor endowment:

$$p = \frac{1}{c} = \frac{1}{(1-g)y}.$$

2. Final-good producers of y take capital k , wages w , and prices of energy q and output p as given. Since $y = k^\alpha [\min\{A_y l_y, A_e e\}]^{1-\alpha} \exp(-\Delta_{y,t} D_t)$, factor compensation for labour and energy together receives a share $(1-\alpha)$ of the value of output py :

$$wl_y + qe = (1-\alpha)py$$

where $e = e_f + e_n$.

3. Fossil-fuel energy production combines labor and fuels, with technology $e_{f,t} = \min\{A_{f,t} l_{f,t}, B_t z_t\}$. Fossil fuel use and labour employed, $z, l_f \geq 0$, are strictly positive if q covers the factor payments, including the carbon price τ

$$\left[q - \left(\frac{w}{A_f} + \frac{\tau}{B} \right) \right] \times l_f \leq 0.$$

The zero profit condition for fossil fuel energy allocates the value of fossil fuel energy to labour and emission payments; using the production identity we can express it in terms of labour employed,

$$qe_f = wl_f + \tau z = \left(w + \frac{\tau A_f}{B} \right) l_f.$$

4. Carbon-free energy inverse supply is given by the first-order condition

$$q = w \frac{\partial l_n}{\partial e_n} = \frac{w_t}{(A_n)^{\frac{\varphi}{\varphi+1}}} (l_n)^{\frac{1}{\varphi+1}}.$$

The value share of labour employed in the carbon-free energy sector equals $\varphi/(1+\varphi)$, so that the rent value is expressed in labour employed:

$$qe_n = \left(1 + \frac{1}{\varphi} \right) wl_n$$

We obtain four equations in four unknowns l_y, l_f, l_n, w :

$$A_y l_y = A_e (A_f l_f + \frac{\varphi + 1}{\varphi} (A_n l_n)^{\frac{\varphi}{\varphi+1}}) \quad (24)$$

$$w l + \frac{\tau A_f}{B} l_f + \frac{1}{\varphi} w l_n = \frac{1 - \alpha}{1 - g} \quad (25)$$

$$\frac{w}{A_f} + \frac{\tau}{B} \geq \frac{w}{(A_n)^{\frac{\varphi}{\varphi+1}}} (l_n)^{\frac{1}{\varphi+1}} \perp l_f \geq 0 \quad (26)$$

$$l_y + l_f + l_n = l \quad (27)$$

For (24) note that, for strictly positive input prices, $A_t(\cdot) = \min \{A_y l_y, A_e e\} \Rightarrow A_y l_y = A_e e$. In equation (25) we allocate the value of output that is not attributed to capital (the right-hand side) to the labour, carbon emissions, and land rent for the non-carbon energy (where we latter two terms are expressed in labour units). Equation (26) compares the production costs for fossil fuel energy with non-carbon energy, and the last equation is the labor market clearing equation. Note that the solution depends on the state of the economy only through total labor l and productivities A_y, A_e, A_f, A_n .

In the absence of a carbon policy, $\tau = 0$, we can solve the allocation in closed-form:

$$l_{n,t} = \frac{A_{n,t}^{\varphi}}{A_{f,t}^{\varphi+1}} \quad (28)$$

$$w_t = \frac{1 - \alpha}{1 - g} \frac{\varphi}{\varphi l_t + l_{n,t}} \quad (29)$$

$$l_{y,t} = \frac{A_{e,t}}{A_{y,t} + A_{e,t} A_{f,t}} [A_{f,t} (l_t - l_{n,t}) + \frac{\varphi + 1}{\varphi} (A_{n,t} l_{n,t})^{\frac{\varphi}{\varphi+1}}] \quad (30)$$

$$l_{f,t} = l_t - l_{y,t} - l_{n,t} \quad (31)$$

Here we include the time subscripts to emphasise the drivers of the solution. This business-as-usual allocation is used to calibrate the productivities. When $\tau > 0$, the solution is numerical, and available in the supplementary file.

Impact of utility function transformation and 100 % depreciation on the DICE outcome

Figure 6: We consider the numbers presented in Nordhaus (2008), Table 5.4. These are the optimal carbon prices in the central DICE run. We take the DICE model, reproduce these numbers, with the elasticity of marginal utility=2 (“DICE: base”). Then, we change the utility function to logarithmic and adjust the discount rate using the Ramsey rule to keep the effective discounting as in “DICE: base”. This gives the DICE run with elasticity of marginal utility=1 (“DICE: log utility”). Finally, we change DICE capital depreciation to 100 per cent (“DICE: log utility & full cap. dep.”). In the last case, DICE remains observationally equivalent with respect to savings and output, which is achieved by adjustments in capital share (from 30 per cent to 27 per cent), initial capital stock, and productivity growth. Together, the three lines show that the elasticity of marginal utility and capital depreciation have limited effect on the carbon price, if such choices are embedded in a broader calibration exercise. See also Appendix “Adjustments in the calibration: approximating less than 100 per cent capital depreciation” for detailed analysis of the adjustments to the calibration in the full capital depreciation case.

Comparison of DICE and the current paper policy-ramp predictions, smooth damages

Figure 7: This Figure continues from Figure 6. The DICE path depicted is obtained for log-utility (elasticity of marginal utility=1) and 2 % discount rate, as is assumed throughout this paper. Our model produces the GL path, using the sure-loss damages corresponding to those in DICE ($\Delta_y = 1$). We find that our numbers are ca. 10 per cent lower over period 2010-2100. The gap can be closed by adjusting the carbon cycle parameters (taken from the climate-science literature).

Bad news carbon price

We define next the bad news carbon price as the socially optimal price at time t if bad news arrives at time t . That is, $I_t = 1$ holds for the first time at t so that the carbon price at that moment is the one determined in Section 3.1. Fig. 8 depicts both the no news and bad news carbon price paths for the near and longer terms, for $\Delta_y = 4$ which was the main case in the text. Note that the bad news price path is “virtual” because it is drawn for an economy in which the output is not dynamically adjusted due to past damages to obtain the bad news price for any given t ; otherwise, time t would not be the first arrival date of damages.⁵⁰ The starting level is given by our calibration at 53 EUR/tCO_2 , 3.7 times larger than the no news price; the bad news price is about factor six larger by the end of time horizon in the Figure. The increase in the virtual price over time reflects purely the effect of income growth on carbon pricing: the marginal utility loss from current emissions reductions, $\frac{\partial y_t}{\partial z_t} u'_t$, weigh less and less over time while the future marginal utility loss of current emissions, h , remains the same.

Figure 8: “Bad news carbon price” is the optimal price at time t if impacts arrive at t . “No news carbon price” is the optimal price at t conditional on not observing impacts by t ($\mu_0 = 0.8$; $\lambda = 0.08$, $\Delta_y = 4$).

⁵⁰The immediate loss of output at time t is accounted for in the calculation of the tax but then again ignored when moving to $t + 1$ to obtain a consistent bad news price for $t + 1$.

Recursive formulation

Here we formally derive value functions recursively and the policy functions.⁵¹ For notational convenience, in contrast with the main text, we use superscript 0 in state $I_t = 0$, while superscript 1 refers to the state after observing a high damage. Thus, we connect notation here with that in the main text, as $h^1 = h_{\lambda=1}^l = h$, and $h^0 = h_{\lambda<1}^l = h^l$.

Formally, the full state vector is $\mathcal{S}_t = (k_t, s_t, I_t, \mu_t)$ where s_t collects the vector of climate state variables through the series of past emissions, and (I_t, μ_t) is the information state. The climate affects the continuations payoffs only through the weighted sum of past emissions, as expressed in the emissions-temperature (9).

Policies take the form $k_{t+1} = \mathcal{G}_t(k_t, s_t, I_t, \mu_t)$, $z_t = \mathcal{H}_t(k_t, s_t, I_t, \mu_t)$. For given policies $\mathcal{G}_t(\cdot)$ and $\mathcal{H}_t(\cdot)$, we can write welfare recursively as

$$\begin{aligned} W_t^0(k_t, s_t, \mu_t) &= u_t + (1 - \mu_t \lambda) \delta W_{t+1}^0(k_{t+1}, s_{t+1}, \mu_{t+1}) + \mu_t \lambda \delta W_{t+1}^1(k_{t+1}, s_{t+1}) \quad (32) \\ W_t^1(k_t, s_t) &= u_t + \delta W_{t+1}^1(k_{t+1}, s_{t+1}) \quad (33) \end{aligned}$$

where $W_{t+1}^1(k_t, s_t)$ is the value function after observing impacts, and $W_{t+1}^0(k_t, s_t, \mu_t)$ is the value function before learning. We note that $I_t = 1$ is an absorbing state; the optimal policy after observing a catastrophe is relatively straightforwardly determined, and by backwards induction, welfare and policies before learning are determined.

More specifically, we guess and verify that optimal policies can be described through constants (g, h^0, h^1) where $0 < g < 1$ is the share of the gross output invested,

$$k_{t+1} = g y_t, \quad (34)$$

and h^1 (in main text: h , as in (13)) is the climate policy variable after observing high damage, that is, a constant that measures the current utility-weighted marginal product of carbon, $\frac{\partial y_t}{\partial z_t} \frac{\partial u}{\partial c} = h^1$, and $\mu_t h^0$ (in main text: $\mu_t h^l$, as in Theorem 1) is the climate policy variable before learning.

The policies, through the functional assumptions, define the marginal product of the fossil fuel use, the carbon price, as

$$\frac{\partial y_t}{\partial z_t} = \mu_t h^0 (1 - g) y_t \text{ if } I_t = 0, \quad (35)$$

$$\frac{\partial y_t}{\partial z_t} = h^1 (1 - g) y_t \text{ if } I_t = 1, \quad (36)$$

⁵¹The proof builds on Gerlagh and Liski (2016) but is not the same since in that paper we do not consider beliefs and uncertainty.

Similarly as g measures the stringency of the savings policy, h^0 and h^1 measure the stringency of the climate policy. In particular, the carbon price, $\frac{\partial y_t}{\partial z_t}$, is monotonic in policy h^I , which allows an interchangeable use of these two concepts. Now, for any constants (g, h^0, h^1) such that (34), (35) and (36) are satisfied, we have in Lemma 1, just below, that the energy sector choices do not depend on the current state of the economy (k_t, s_t) but only on the policies as affected through the information set (I_t, μ_t) .

Lemma 1 *For all t :*

- (i) *Given policy $(g, \mu_t h^0)$ for $I_t = 0$ and (g, h^1) for $I_t = 1$ at time t , emissions $z_t = z_t^*$ at t implied by the policy are independent of the current state (k_t, s_t) , but depend only on the current technology at t as captured by $A_t(\cdot)$;*
- (ii) *Emissions z_t^* decrease monotonically in policy stringency $\mu_t h^0$ and h^1 .*

Proof. Substituting (5) into the first-order conditions for emissions (35) and (36), determines z_t^* as independent of k_t and s_t :

$$\begin{aligned} A'_t(z_t^*) &= \mu_t h^0 (1 - g) A(z_t^*) \text{ if } I_t = 0 \\ A'_t(z_t^*) &= h^1 (1 - g) A(z_t^*) \text{ if } I_t = 1 \end{aligned}$$

As $A'_t(z_t^*)$ (LHS) is decreasing and $A_t(z_t^*)$ (RHS) is increasing in z_t^* , it is immediate that z_t^* is unique, independent of (k_t, s_t) , and decreasing in the policy stringency. Q.E.D

■

The independence lemma helps to establish a separable representation of welfare:

Lemma 2 *It holds for every policy (g, h^0, h^1) that*

$$\begin{aligned} W_t^0(k_t, s_t, \mu_t) &= V(k_t) - \mu_t \Omega^0(s_t) + \tilde{A}_t^0(\mu_t) \\ W_t^1(k_t, s_t) &= V(k_t) - \Omega^1(s_t) + \tilde{A}_t^1 \end{aligned}$$

with parametric form

$$\begin{aligned} V(k_t) &= \xi \ln(k_t) \\ \Omega^0(s_t) &= \sum_{\tau=1}^{t-1} \zeta_\tau^0 z_{t-\tau}, \\ \Omega^1(s_t) &= \sum_{\tau=1}^{t-1} \zeta_\tau^1 z_{t-\tau}, \end{aligned}$$

where $\xi = \frac{\alpha}{1-\alpha\delta}$, $\zeta_1^0 = \zeta_1^1 - \Delta(1-\lambda) \sum_{i \in \mathcal{I}} \frac{a_i \pi \varepsilon}{[1-(1-\lambda)\delta(1-\eta_i)][1-(1-\lambda)\delta(1-\varepsilon)]}$, $\zeta_1^1 = \Delta \sum_{i \in \mathcal{I}} \frac{a_i \pi \varepsilon}{[1-\delta(1-\eta_i)][1-\delta(1-\varepsilon)]}$, $\Delta = \frac{1}{1-\alpha\delta}$ and $\tilde{A}_t^0(\mu_t)$ and $\tilde{A}_t^1$ are independent of k_t and s_t .

Proof. Induction hypothesis: assume (i) that future policies are given by a sequence of constants (g, h^0, h^1) such that (34), (35) and (36) are satisfied for all future $\tau \geq t$, and (ii) that Lemma 2 holds for $t + 1$. We can thus construct the value function for the current period, through (32) and (33) and check the validity of the Lemma.

Consider policies at t . From (35), $k_{t+1} = g_t y_t$. Emissions $z_t = z_t^{I,*}$ can be determined independently of the state variables k_t and s_t as shown in Lemma 1. Substituting the policies at t gives:

$$\begin{aligned}
W_t^0(k_t, s_t, \mu_t) &= [\ln(1 - g_t) + \ln(A_t) + \alpha \ln(k_t) - \lambda \mu_t D_t] \\
&\quad + \delta \xi [\ln(g_t) + \ln(A_t) + \alpha \ln(k_t) - \lambda \mu_t D_t] \\
&\quad - \delta [(1 - \lambda \mu_t) \mu_{t+1} \Omega^0(s_{t+1}) - \mu_t \lambda \Omega^1(s_{t+1})] \\
&\quad + \delta [(1 - \lambda \mu_t) \tilde{A}_{t+1}^0 + \lambda \mu_t \tilde{A}_{t+1}^1] \\
W_t^1(k_t, s_t) &= [\ln(1 - g_t) + \ln(A_t) + \alpha \ln(k_t) - D_t] \\
&\quad + \delta \xi [\ln(g_t) + \ln(A_t) + \alpha \ln(k_t) - D_t] \\
&\quad - \delta \Omega^1(s_{t+1}) \\
&\quad + \delta \tilde{A}_{t+1}
\end{aligned}$$

Collecting the coefficients that only depend on future policies g_τ and z_τ for $\tau > t$, and that do not depend on the next-period state variables k_t and s_t , we get the constant part of $W_t^I(\cdot)$:

$$\tilde{A}_t^0 = \ln(1 - g_t) + \delta \xi \ln(g_t) + (1 + \delta \xi) \ln(A_t) \quad (37)$$

$$- \delta [(1 - \lambda \mu_t) \mu_{t+1} \zeta_1^0 - \lambda \mu_t \zeta_1^1] z_t \quad (38)$$

$$+ \delta [(1 - \lambda \mu_t) \tilde{A}_{t+1}^0 + \lambda \mu_t \tilde{A}_{t+1}^1] \quad (39)$$

$$\tilde{A}_t^I = \ln(1 - g_t) + \delta \xi \ln(g_t) + (1 + \delta \xi) \ln(A_t) - \delta \zeta_1 z_t + \delta \tilde{A}_{t+1}^1. \quad (40)$$

Collecting the coefficients in front of $\ln(k_t)$ yields the part of $V(k_t)$ depending k_t with the recursive determination of ξ ,

$$\xi = \alpha(1 + \delta \xi).$$

so that $\xi = \frac{\alpha}{1 - \alpha \delta}$ follows.

Collecting the terms with s_t yields $\Omega^I(s_t)$ through

$$\begin{aligned}
\mu_t \Omega^0(s_t) &= -\lambda \mu_t (1 + \delta \xi) D_t + \delta (1 - \lambda) \mu_t \Omega^0(s_{t+1}) + \delta \lambda \mu_t \Omega^1(s_{t+1}) \\
\Omega^1(s_t) &= -(1 + \delta \xi) D_t + \delta \Omega^1(s_{t+1}),
\end{aligned}$$

where we used Bayesian updating of beliefs (8): $(1 - \lambda\mu_t)\mu_{t+1} = (1 - \lambda)\mu_t$, and where $z_t = z_t^*$ appearing in $s_{t+1} = (z_1, \dots, z_{t-1}, z_t)$ is independent of k_t and s_t so that we only need to consider the values for $z_1, \dots, z_t$ when evaluating $\Omega(s_t)$. The values for ζ_τ^I can be calculated by collecting the terms in which $z_{t-\tau}$ appear. Recall $\frac{dD_{t+\tau}}{dz_t} = \mathcal{R}(\tau)$ from (9):

$$\begin{aligned}\zeta_\tau^0 &= (1 + \delta\xi)\lambda \sum_{i \in \mathcal{I}} a_i \pi \varepsilon \frac{(1 - \eta_i)^\tau - (1 - \varepsilon)^\tau}{\varepsilon - \eta_i} + \delta(1 - \lambda)\zeta_{\tau+1}^0 + \delta\lambda\zeta_{\tau+1}^1 \\ \zeta_\tau^1 &= (1 + \delta\xi) \sum_{i \in \mathcal{I}} a_i \pi \varepsilon \frac{(1 - \eta_i)^\tau - (1 - \varepsilon)^\tau}{\varepsilon - \eta_i} + \delta\zeta_{\tau+1}^1\end{aligned}$$

We rewrite the first equation as

$$\zeta_\tau^0 - \zeta_\tau^1 = -(1 + \delta\xi)(1 - \lambda) \sum_{i \in \mathcal{I}} a_i \pi \varepsilon \frac{(1 - \eta_i)^\tau - (1 - \varepsilon)^\tau}{\varepsilon - \eta_i} + \delta(1 - \lambda)[\zeta_{\tau+1}^0 - \zeta_{\tau+1}^1]$$

Substitution of the recursive formula, for all subsequent τ , gives

$$\begin{aligned}\zeta_\tau^0 &= \zeta_\tau^1 - \frac{1}{1 - \alpha\delta} \sum_{i \in \mathcal{I}} \sum_{t=\tau}^{\infty} a_i \pi \varepsilon (1 - \lambda)^{t-\tau} \delta^{t-\tau} \frac{(1 - \eta_i)^t - (1 - \varepsilon)^t}{\varepsilon - \eta_i} \\ \zeta_\tau^1 &= \frac{1}{1 - \alpha\delta} \sum_{i \in \mathcal{I}} \sum_{t=\tau}^{\infty} a_i \pi \varepsilon \delta^{t-\tau} \frac{(1 - \eta_i)^t - (1 - \varepsilon)^t}{\varepsilon - \eta_i}.\end{aligned}$$

To derive the value of ζ_1^I as expressed in the lemma, we consider

$$\begin{aligned}\sum_{t=1}^{\infty} \delta^{t-1} \frac{(1 - \eta_i)^t - (1 - \varepsilon)^t}{\varepsilon - \eta_i} &= \frac{\sum_{t=1}^{\infty} [\delta(1 - \eta_i)]^t - \sum_{t=1}^{\infty} [\delta(1 - \varepsilon)]^t}{\delta(\varepsilon - \eta_i)} \\ &= \frac{\frac{\delta(1 - \eta_i)}{1 - \delta(1 - \eta_i)} - \frac{\delta(1 - \varepsilon)}{1 - \delta(1 - \varepsilon)}}{\delta(\varepsilon - \eta_i)} \\ &= \frac{1}{[1 - \delta(1 - \eta_i)][1 - \delta(1 - \varepsilon)]}\end{aligned}$$

The formulation for $\zeta_\tau^1 - \zeta_\tau^0$ follows by symmetry with $(1 - \lambda)\delta$ substituted for δ . This proves the value function representation lemma.

The expected future cost of the emission history is thus given by $\mu_t \Omega^0(s_t)$ and $\Omega^1(s_t)$, giving also the marginal cost of the current emissions as $\frac{\partial \mu_{t+1} \Omega^0(s_{t+1})}{\partial z_t} = \mu_{t+1} \zeta_1^0$, and $\frac{\partial \Omega^1(s_{t+1})}{\partial z_t} = \zeta_1^1$. It is then easily seen that policies (g, h^0, h^1) as defined in text are optimal, so that, indeed, the welfare function plus policies are consistent.

Proof of Proposition 4

Proof. The proof proceeds in two steps. First, we show that the carbon price, conditional on the event arrival, is log-convex in the discount factor δ . Second, we show that this property leads to the proposition.

Let us rewrite the marginal utility impact of emissions as

$$\begin{aligned} h_\lambda^l &= \tilde{h}(\delta) - \tilde{h}((1-\lambda)\delta) \\ \tilde{h}(x) &= x\Delta\pi \frac{\varepsilon}{1-x(1-\varepsilon)} \sum_{i \in \mathcal{I}} \frac{a_i}{1-x(1-\eta_i)} \end{aligned} \quad (41)$$

That is, the post-event carbon price in utility terms is given by $\tilde{h}(\delta)$. We now show log-convexity for the separate terms.

$$\begin{aligned} \frac{\partial}{\partial x} \frac{\partial \ln \left(\frac{x}{1-x(1-\varepsilon)} \right)}{\partial x} &= \frac{\partial}{\partial x} \frac{1-x(1-\varepsilon)}{x} \left(\frac{1}{1-x(1-\varepsilon)} + \frac{(1-\varepsilon)x}{(1-x(1-\varepsilon))^2} \right) \\ &= \frac{\partial}{\partial x} \frac{1-x(1-\varepsilon)}{x} \frac{1}{(1-x(1-\varepsilon))^2} \\ &= \frac{\partial}{\partial x} \frac{1}{x(1-x(1-\varepsilon))} \end{aligned}$$

so that

$$\frac{\partial}{\partial x} \frac{\partial \ln \left(\frac{x}{1-x(1-\varepsilon)} \right)}{\partial x} > 0 \Leftrightarrow x > \frac{1}{2} \frac{1}{1-\varepsilon} > \frac{1}{2}.$$

We then consider the second part of $\tilde{h}(\delta)$:

$$\begin{aligned} \frac{\partial}{\partial x} \frac{\partial \ln \left(\sum_{i \in \mathcal{I}} \frac{a_i}{1-x(1-\eta_i)} \right)}{\partial x} &= \frac{\partial}{\partial x} \frac{\sum_{i \in \mathcal{I}} \frac{a_i(1-\eta_i)}{(1-x(1-\eta_i))^2}}{\sum_{i \in \mathcal{I}} \frac{a_i}{1-x(1-\eta_i)}} \\ &= \frac{\sum_{i \in \mathcal{I}} \frac{\partial}{\partial x} \frac{a_i(1-\eta_i)}{(1-x(1-\eta_i))^2}}{\sum_{i \in \mathcal{I}} \frac{a_i}{1-x(1-\eta_i)}} - \frac{\sum_{i \in \mathcal{I}} \frac{a_i(1-\eta_i)}{(1-x(1-\eta_i))^2} \sum_{i \in \mathcal{I}} \frac{\partial}{\partial x} \frac{a_i}{1-x(1-\eta_i)}}{\left(\sum_{i \in \mathcal{I}} \frac{a_i}{1-x(1-\eta_i)} \right)^2} \\ &= \frac{\sum_{i \in \mathcal{I}} 2 \frac{a_i(1-\eta_i)^2}{(1-x(1-\eta_i))^3}}{\sum_{i \in \mathcal{I}} \frac{a_i}{1-x(1-\eta_i)}} - \frac{\left(\sum_{i \in \mathcal{I}} \frac{a_i(1-\eta_i)}{(1-x(1-\eta_i))^2} \right)^2}{\left(\sum_{i \in \mathcal{I}} \frac{a_i}{1-x(1-\eta_i)} \right)^2}. \end{aligned}$$

We multiply both terms by the quadratic denominator, and get

$$\begin{aligned}
& \frac{\partial}{\partial x} \frac{\partial \ln \left(\sum_{i \in \mathcal{I}} \frac{a_i}{1-x(1-\eta_i)} \right)}{\partial x} > 0 \iff \\
& \sum_{i \in \mathcal{I}} 2 \frac{a_i(1-\eta_i)^2}{(1-x(1-\eta_i))^3} \sum_{i \in \mathcal{I}} \frac{a_i}{1-x(1-\eta_i)} - \left(\sum_{i \in \mathcal{I}} \frac{a_i(1-\eta_i)}{(1-x(1-\eta_i))^2} \right)^2 > 0 \iff \\
& \sum_{i,j \in \mathcal{I}} 2 \frac{a_i(1-\eta_i)^2}{(1-x(1-\eta_i))^3} \frac{a_j}{1-x(1-\eta_j)} - \sum_{i,j \in \mathcal{I}} \frac{a_i(1-\eta_i)}{(1-x(1-\eta_i))^2} \frac{a_j(1-\eta_j)}{(1-x(1-\eta_j))^2} > 0 \iff \\
& \sum_{i,j \in \mathcal{I}} \omega_{ij} \left(2 \frac{(1-\eta_i)^2}{(1-x(1-\eta_i))^2} - \frac{(1-\eta_i)}{(1-x(1-\eta_i))} \frac{(1-\eta_j)}{(1-x(1-\eta_j))} \right) > 0 \iff \\
& \sum_{i,j \in \mathcal{I}} \omega_{ij} \left(\frac{(1-\eta_i)^2}{(1-x(1-\eta_i))^2} + \frac{(1-\eta_j)^2}{(1-x(1-\eta_j))^2} - \frac{(1-\eta_i)}{(1-x(1-\eta_i))} \frac{(1-\eta_j)}{(1-x(1-\eta_j))} \right) > 0 \iff \\
& \sum_{i,j \in \mathcal{I}} \frac{1}{2} \omega_{ij} \left[\frac{(1-\eta_i)^2}{(1-x(1-\eta_i))^2} + \frac{(1-\eta_j)^2}{(1-x(1-\eta_j))^2} + \left(\frac{(1-\eta_i)}{(1-x(1-\eta_i))} - \frac{(1-\eta_j)}{(1-x(1-\eta_j))} \right)^2 \right] > 0
\end{aligned}$$

where we moved from the second to the third line, recognizing that $\sum_{i \in \mathcal{I}} \sum_{i \in \mathcal{I}} = \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{I}} = \sum_{i,j \in \mathcal{I}}$, and in the fourth and subsequent lines, we substituted ω_{ij} for

$$\frac{a_i a_j}{(1-x(1-\eta_i))(1-x(1-\eta_j))}.$$

The last inequality is always satisfied. Thus, we conclude that $\tilde{h}(x)$ is log-convex for $x > 0.5$. We rewrite this as

$$\frac{\partial}{\partial x} \frac{\tilde{h}'(x)}{\tilde{h}(x)} > 0$$

and then compare $x = \delta$ with $x = (1-\lambda)\delta$:

$$\begin{aligned}
& \frac{(1-\lambda)\tilde{h}'((1-\lambda)\delta)}{\tilde{h}((1-\lambda)\delta)} < \frac{\tilde{h}'(\delta)}{\tilde{h}(\delta)} \Leftrightarrow \\
& -(1-\lambda)\tilde{h}(\delta)\tilde{h}'((1-\lambda)\delta) > -\tilde{h}((1-\lambda)\delta)\tilde{h}'(\delta) \Leftrightarrow \\
& \tilde{h}(\delta)\tilde{h}'(\delta) - (1-\lambda)\tilde{h}(\delta)\tilde{h}'((1-\lambda)\delta) > \tilde{h}(\delta)\tilde{h}'(\delta) - \tilde{h}((1-\lambda)\delta)\tilde{h}'(\delta) \Leftrightarrow \\
& \tilde{h}(\delta)(\tilde{h}'(\delta) - (1-\lambda)\tilde{h}'((1-\lambda)\delta)) > (\tilde{h}(\delta) - \tilde{h}((1-\lambda)\delta))\tilde{h}'(\delta) \Leftrightarrow \\
& \frac{\tilde{h}'(\delta) - (1-\lambda)\tilde{h}'((1-\lambda)\delta)}{\tilde{h}(\delta) - \tilde{h}((1-\lambda)\delta)} > \frac{\tilde{h}'(\delta)}{\tilde{h}(\delta)}.
\end{aligned}$$

The last line, is equivalent to the inequality stated in the proposition, when substituting the definition (41) for h_λ^l :

$$\frac{1}{h_\lambda^l} \frac{\partial h_\lambda^l}{\partial \delta} > \frac{1}{h_{\lambda=1}^l} \frac{\partial h_{\lambda=1}^l}{\partial \delta}$$

Q.E.D. ■

Adjustments in the calibration: approximating less than 100 per cent capital depreciation

Our analytical model assumes 100 per cent capital depreciation per period. Here we demonstrate the adjustments needed in the calibration if one is interested in accommodating a lower depreciation factor. Since the capital policy is separable from the climate policy in our setting, we can demonstrate the adjustments using the standard consumption choice model:

$$W_t = l_t \ln \frac{c_t}{l_t} + \delta W_{t+1} \quad (42)$$

$$y_t = A_t k_t^\alpha l_t^{1-\alpha} \quad (43)$$

$$k_{t+1} = (1 - \beta)k_t + i_t \quad (44)$$

$$y_t = c_t + i_t \quad (45)$$

where W_t denotes the value function, l_t is labor, c_t, i_t, k_t are consumption, investments and capital, A_t is total factor productivity, α is the capital share in output, β is the capital depreciation factor, and δ is the time discount factor. In balanced growth, we write the growth factors for $x = y, k, c, l, A$, as $\hat{x} = x_{t+1}/x_t$, and we note that in balanced growth we have $\hat{c} = \hat{i} = \hat{k} = \hat{y}$. We define the net rate of return on transferring goods between period t and $t + 1$ as $r = \hat{p}^{-1} - 1 = p_t/p_{t+1} - 1$, where $p_t = \partial W_0 / \partial c_t$.

Consider the set of balanced growth variables $\{y_0, k_0, i_0, c_0, \hat{y}, r\}$, where the subscript 0 refers to an observation in the year of reference. The true capital stock is unobservable and in macro-economic accounts, it is typically estimated based on observed historic investments and an assumed depreciation. We thus partition the set of variables in observed and unobserved variables and define vector spaces: $\mathcal{V}_{obs} = \{(y_0, i_0, c_0, \hat{y}, r)\}$; $\mathcal{V}_{unobs} = \{(k_0)\}$. Further, consider the set of parameters $\{\alpha, \beta, \delta, l_0, \hat{l}, A_0, \hat{A}\}$, which we partition into parameters that we choose based on exogenous scenarios or rules-of-thumb, $\mathcal{P}_{exo} = \{(\beta, l_0, \hat{l})\}$, and calibrated parameters, $\mathcal{P}_{cal} = \{(\alpha, \delta, A_0, \hat{A})\}$

The calibration is a mapping from observed variables to the calibrated parameters:

$$\mathcal{C} : \mathcal{V}_{obs} \rightarrow \mathcal{P}_{cal}$$

Below we present the calibration, and we specifically are interested in the calculation of A_0 (or A_t which is sometimes a more convenient notation) and α as dependent on the observed variables and exogenous parameters. Specifically, we want to see how a change

in β affects the calibrated value for α . The production function and capital dynamics identity (43)-(44) fix the output growth rate and the investment-capital ratio. Output growth comes from labour and productivity growth, and investments need to balance depreciation plus the build up of a growing capital stock:

$$\widehat{y} = \widehat{A}^{\frac{1}{1-\alpha}} \widehat{l} \quad (46)$$

$$\frac{i_t}{k_t} = \widehat{y} - 1 + \beta \quad (47)$$

The above first equation is used to calibrate productivity growth $\widehat{A}$. The first-order condition for the consumption-savings tradeoff gives the Ramsey rule:

$$1 + r = \frac{\widehat{y}}{\delta \widehat{l}}, \quad (48)$$

which we use to calibrate δ , and where finite welfare requires $\delta \widehat{l} < 1$ (that is, the net present value of output needs to be finite: $\widehat{y} < 1 + r$). The first-order condition for capital gives

$$\alpha \frac{y_t}{k_t} = r + \beta \quad (49)$$

which results in

$$\alpha = (r + \beta) \frac{k_t}{i_t} \frac{i_t}{y_t} = \frac{\frac{\widehat{y}}{\delta \widehat{l}} - 1 + \beta}{\widehat{y} - 1 + \beta} \frac{i_t}{y_t} \quad (50)$$

Note that the ratio in front of i_t/y_t on the far RHS exceeds one, the numerator exceeds the denominator, so that the RHS decreases in β . The equations inform us about the calibration adjustments that follow from $\beta < 1$. To keep the same observable output, consumption and investment levels, y_0, c_0, i_0 , growth rates, $\widehat{y}$, and returns on savings, r , a smaller capital depreciation ($\beta < 1$) must be balanced by a larger (unobserved) capital stock k_0 in (47), a larger capital-income share α in (50), and an adjusted productivity A_t in (43). The time preference parameter δ and productivity growth $\widehat{A}$ are not affected.