

Carbon prices for the next thousand years

Reyer Gerlagh and Matti Liski*

September 24, 2012

Abstract

Climate is a persistent asset, bar none: changes in climate-related stocks have consequences spanning over centuries or possibly millennia to the future. To reconcile the sensitivity of policies to such far-distant impacts and realism of the shorter-term decisions, we consider hyperbolic time-preferences in a climate-economy model. The climate-economy model is rich in details but can be solved in closed-form yielding Markov carbon prices dependent on climate system parameters, damage estimates, technology parameters, and both short- and long-term time preferences.

Preferences with declining discount rates have unexplored general-equilibrium effects: carbon prices exceed the pure carbon externality costs—the Pigouvian tax level—by multiple factors in our quantitative assessment. The model justifies high carbon taxes while preserving a realistic calibration, thus providing a solution for the dilemma centering the carbon tax-discount rate debate. The welfare ranking of the policy alternatives is unambiguous: enforcing the Pigouvian tax decreases a consistently-defined welfare measure.

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

*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 Geir Asheim, Larry Goulder, Baard Harstad, Jon Hassler, Michael Hoel, Terry Iverson, Larry Karp, Dave Kelly, Per Krusell, Rick van der Ploeg, Tony Smith, Sjak Smulders, Christian Traeger, Cees Withagen, participants at Cowles Foundation 2012, NBER Summer Institute 2012, and SURED 2012 meetings, and at seminars in Helsinki HECER, Stockholm IIES, Toulouse TSE, and Oxford OxCarre for many useful comments and discussions.

1 Introduction

“Climate” is an extremely persistent asset, with a complicated and long delay structure of impacts, involving atmospheric and ocean carbon dioxide diffusion, and land surface and ocean temperature adjustments (e.g., Maier-Reimer and Hasselmann 1987, Hooss et al. 2001). The consequences of current changes in climate-related stocks span over centuries or possibly millennia into the future. The persistence of climate change is a central feature in applied climate-economy models, the so called integrated assessment models (IAMs) put forward by Peck and Teisberg (1992), Nordhaus (1993), and Manne and Richels (1995). However, policy evaluations of the climate-economy models ignore the persistent far-distant socioeconomic climate impacts, when holding a view on discounting that respects revealed preferences on shorter-term decisions such as the savings behavior (for discussion, see, e.g., Nordhaus 2007, Weitzman 2007, Dasgupta 2008).

One may take it as given that policies should respect the shorter-term time preferences consistent with historical data (Nordhaus, 2007). But there is also evidence that preferences for the far-distant future should be treated differently. Weitzman (2001) surveyed 2,160 economists for their best estimate of the appropriate real discount rate to be used for evaluating environmental projects over a long time horizon, and used the data to argue that the policy maker should use a discount rate that declines over time — coming close to zero after 300 years. Rather than relying on experts, some studies invoke “similarity” as evidence for non-constant discounting: from the current perspective, generations living after 400 or, alternatively, after 450 years look the same. That being the case, no additional discounting arises from the added 50 years, while the same time delay commands large discounting in the near term. Rubinstein (2003) was among the first to argue that such similarity of alternatives can justify hyperbolic time discounting (see also Karp 2005).

While much has been said to motivate time-changing discounting for climate policy evaluations, the climate-economy general-equilibrium implications have gone unnoticed.¹ This paper considers the normative implications of such time preferences for climate policies: How should we design the current climate policies, internalizing all climate

¹The Stern Review (2006) triggered a heated debate on how the policies of the climate-economy models could be made more sensitive to the long-term climate outcomes — essentially, the literature has sought for reasons to use lower long-term discount rates. In general, non-constant discount rates can result, e.g., from: aggregation over heterogenous individuals (Gollier and Zeckhauser, 2005; Lengwiler, 2005, Jackson and Yariv, 2011); uncertainty (Gollier, 2002). We discuss the motivations directly consistent with our setting in detail below.

impacts of emissions, when the future decision makers design their own policies looking at the future from their vantage point? Policy makers living 400 years after us also make a distinction between the short and long runs, even though —from our point— they live in the long run.

In this setting, the distinction between short- and long-run discounting leads to a climate policy game between generations even when the current and all future policy makers internalize all climate impacts of emissions. The equilibrium time discount factor is then endogenous, as opposed to one that is assumed *a priori*, and this endogeneity of time preference allows partial de-linking of the equilibrium discount factor used for savings from the one used for evaluating the future impacts of current carbon emissions. Strikingly, the current optimal carbon price exceeds the Pigouvian tax level — unambiguously defined as the net present value of the future externality costs from current emissions — potentially by multiple factors. Apart from the time structure of preferences, the framework for quantitative analysis producing these results is a standard general-equilibrium growth framework with a climate module, following the Nordhaus’ approach and its recent gearing towards the macro traditions by Golosov, Hassler, Krusell, and Tsyvinsky (2011).

The key distortion introduced by time-changing time discounting is the lack of commitment to actions that we would like to implement in the future, as the future decision makers with their present-biases (from our point of view) control their own capital savings and emissions. But the future decision makers face the same dilemma — they value future savings and emission reductions, after their time, relatively more than the subsequent actual polluters. Therefore, also the future policy makers would value any devices that would allow commitment to long-run actions. The long persistence of climate impacts provides such a commitment device: the current climate policy decisions have a long lasting impact on the future agents. Also, future agents have no reason to undermine past climate investments as they value the climate capital for the same reason. The mechanism is similar to that delivering value for commitment devices in self-control problems (Laibson, 1997);² it explains why investments in climate protection is valued above the level implied by the pure Pigouvian externality costs of emissions.

Our point of departure is a Markov equilibrium where each generation sets its self-interested savings and climate targets understanding how the future generations respond to current choices. The model is analytically tractable having thus the same virtues as

²However, self-control at the individual level is not the interpretation of the “behavioral bias” in our economy, as we explain shortly.

the approach in Golosov et al. 2011, even though we include time-changing discounting and a more detailed delay-structure of the carbon-climate cycle — such delays have dramatic effects on the carbon price levels. As a result, we can derive the Markov equilibrium savings and carbon prices dependent on the multi-layer climate system parameters, damage estimates, production technology parameters, and both short and long-term time-preferences.

In the Markov equilibrium, the commitment value is larger the longer are the climate delays, and it can justify high carbon prices as suggested by the Stern Review (2006) while allowing realistic calibration of the shorter-term macro variables. Table 1 contains the gist of the quantitative assessment, detailed at the end of the paper. The technology parameters of the model are calibrated to 25 per cent gross savings, when preferences are consistent and the annual time discount rate is 2 per cent. This is consistent with the Nordhaus’ DICE 2007 baseline scenario (Nordhaus, 2007), giving 8.4 Euros per ton of CO₂ as the optimal carbon price in the year 2010 (i.e., 40 Dollars per ton C). When long-term preference parameters are chosen such that the long-term receives a higher weight (roughly consistent with Weitzman’s (2001) survey results), short-term preferences can be matched so that the model remains observationally equivalent to Nordhaus in terms of macroeconomic performance, savings in particular. But carbon prices increase: for very low long-term discounting, ultimately carbon prices approach those suggested by Stern.³ Note that the equilibrium outcome — high carbon prices and realistic savings — does not implement an efficient allocation for such preferences, but it presents the Markov equilibrium describing how the climate decisions are *de facto* made given the order of moves in the time line.

| | discount rate | | savings | carbon price |
|------------|---------------|-----------|---------|--------------|
| | short-term | long-term | | |
| “Nordhaus” | .02 | .02 | .25 | 8.4 |
| Markov | .026 | .001 | .25 | 116.4 |
| “Stern” | .001 | .001 | .30 | 151.8 |

Table 1: Equilibrium carbon prices in EUR/tCO₂ year 2010.

What is then the first-best in the climate policy game? We treat agents in different periods as distinct generations (as in Phelps and Pollak, 1968), and thus the multi-

³Under “Stern” the capital-share of output is fully saved (30 per cent); increasing the capital-share leads to unrealistic savings as discussed, e.g., in Weitzman 2007 and Dasgupta 2008.

generation Pareto optimality is a natural welfare concept (as, e.g., in Caplin and Leahy, 2004) for considering whether policy measures can improve welfare above that in the Markov equilibrium.⁴ A natural policy experiment is to impose the Pigouvian externality pricing of carbon as an institutional constraint — it can be thought of as a rule of behavior for an environmental agency scrutinizing the climate policies within each period. The result is striking: pricing carbon according to the common principle that each unit of emissions should pay the present-value marginal damages caused by that unit lowers welfare for all generations! That is, even if the present generation and all future generations could commit to follow this rule, the commitment would generate no social value, without additional policy measures. To explain this surprising result, we show the equivalence of the Pigouvian externality pricing and a standard definition of efficiency for consumption streams. But, we also show that such efficiency does not imply Pareto optimality for welfare levels when there is intergenerational altruism as implied by quasi-hyperbolic discounting (Saez-Marti and Weibull, 2005).⁵

The changing time-preference opens the door for our conceptual results; their quantitative significance follows from the unusual delays of the consequences of climate change. We develop and calibrate a tractable representation of the carbon cycle, with the peak impact lagging 60-70 years behind the date of emissions. The analytics allows us to decompose the contribution of the different layers of the climate system to the carbon price: ignoring the delay of impacts — as in Golosov et al. 2011 — misses the correct price levels by a factor of 2, even when preferences are consistent. It should be emphasized that getting the carbon price right is not merely an academic exercise; such prices are increasingly factored into the policy decisions, for example, into those that favor particular electricity generation technologies.⁶

We take the time-structure of preferences as given and focus on their general-equilibrium climate policy implications; however, multiple recent arguments can justify the deviation from geometric discounting. First, if we accept that the difficulty of distinguishing long-run outcomes describes well the climate-policy decision problem, then our decision procedure can imply a lower long-term discount factor than that for the short-term de-

⁴See Bernheim and Rangel (2009) for an alternative concept, and its relationship to the Pareto criterion. The Pareto criterion may not be reasonable when the focus is on the behavioral anomalies at the individual level. Our welfare analysis follows closely our working paper Gerlagh and Liski (2011)

⁵In a different context, Bernheim and Ray (1987) also show that, in the presence of altruism, efficiency does not imply Pareto optimality.

⁶See Muller et al. (2011) for the dramatic effect that carbon prices can have for the value-added evaluation of the US electricity sector.

cisions (see Rubinstein 2003 for the procedural argument). Second, climate investments are public decisions requiring aggregation over heterogeneous individual time-preferences, leading again to a non-stationary aggregate time-preference pattern, typically declining with the length of the horizon, for the group of agents considered (Zeckhauser and Gollier 2005; Jackson and Yariv 2011). We can also interpret Weitzman’s (2001) study based on the survey of experts’ opinions on discount rates as an aggregation of persistent views. Third, the long-term valuations must by definition look beyond the welfare of the immediate next generation; any pure altruism expressed towards the long-term beneficiaries implies changing utility-weighting over time (Phelps and Pollak 1968 & Saez-Marti and Weibull 2005).⁷

The paper is organized as follows. The next section introduces a simplified version of the model in three periods to pinpoint the biases in the carbon prices, policy proposals, as well as the main idea of the welfare analysis. The section has several subsections as we want to clarify all stages of the main plot in three periods in order to avoid confusion in the main model. Section 3 then introduces the full infinite-horizon climate-economy model, and provides the first look at the numbers on carbon pricing. Section 4 calibrates the full model with climate system parameters from the scientific evidence to generate the climate-economy outcomes over the next thousand years — due to the unique delays in the system non-trivial effects remain over such horizons. Section 5 discusses the relevance of the results for the climate policy debates. Section 6 concludes.

2 A three-period model

2.1 Technologies and preferences

Consider three generations, living in periods $t = 1, 2, 3$. In each period, consumers are represented by an aggregate agent having a concern also for future consumers’ utilities and welfare. The representative decision-maker in each period internalizes all future impacts of current actions, and thereby there is no climate-change commons problem but rather the focus is on the optimal design of policies under the preferences that we define shortly. In period one, production depends on the use of fossil fuels, with associated emissions. In period two, a substitute energy source has been developed, and consumption possibilities

⁷The first and third arguments for using lower discount rates in the long run are directly consistent with our formal model. For the second, we do not formally consider the aggregation of preferences, but Jackson and Yariv (2011) show that utilitarian aggregation leads necessarily to a present bias.

depend on the capital stock inherited. In period three, emissions from the first period have translated into a climate problem negatively affecting production. An allocation $(\mathbf{c}, \mathbf{k}, z) = (c_1, c_2, c_3, k_2, k_3, z) \in A \subseteq R_+^6$ (convex set) constitutes a consumption level for each generation c_t , the first-period use of fossil fuels z , which we also consider a proxy for the emissions of carbon dioxide emissions, and capital stocks k_2 and k_3 left for future agents (k_1 is given). Generations care about current and future utilities as follows

$$w_1 = u_1(c_1) + \beta[\delta u_2(c_2) + \delta^2 u_3(c_3)] \quad (1)$$

$$w_2 = u_2(c_2) + \beta[\delta u_3(c_3)] \quad (2)$$

$$w_3 = u_3(c_3), \quad (3)$$

where all utility functions u_t are assumed to be continuous and, in addition, strictly concave, differentiable, and satisfying $\lim_{c \rightarrow 0} u'_t = \infty$. Discount factor $\beta\delta^t$, with $\delta \in (0, 1)$, has the quasi-exponential form that contains the standard exponential function as special case $\beta = 1$. When $\beta < 1$, discounting is quasi-hyperbolic, implying that the first-period decision-maker's patience increases over time: one-period postponement of a utility gain is first discounted with $\beta\delta$ and then with δ . This increasing patience captures the idea of utility discount rates declining over time. Notice that w_t denotes the social preferences, used by each generation for its welfare evaluation of policies influencing future utilities. These preferences are specific for generation t , and in that sense, w_t is different from the generation-independent social welfare function (SWF) as discussed, e.g., in Goulder and Williams (2012) and Kaplow et al. (2010). Furthermore, the first generation's welfare attributes weights to future *utilities* but, obviously, these translate into weights on future *welfares*. The condition $\beta < 1$ is equivalent to pure altruism towards future decision makers (Saez-Marti and Weibull 2005):

$$w_1 = u_1(c_1) + a_2 w_2 + a_3 w_3 \quad (4)$$

$$a_2 = \beta\delta > 0, a_3 = \beta(1 - \beta)\delta^2 > 0,$$

where a_2, a_3 can be interpreted as welfare weights given by the first generation, implied by increasing patience over time. When $\beta = 1$, there is one-period pure altruism, and the typical recursive-dynastic representation of welfare follows.

In the first period, the consumption possibilities are determined by a strictly concave neoclassical production function $f_1(k_1, z)$, where k_1 is the capital stock, and z is the use of fossil fuels, or emissions of carbon dioxide, both having positive marginal products, $\frac{\partial f_1}{\partial k} = f_{1,k}, \frac{\partial f_1}{\partial z} = f_{1,z} > 0$. The first generation starts with a capital stock k_1 , and

produces output using z , which can be used to consume c_1 , or to invest in capital for the immediate next period k_2 :

$$c_1 + k_2 = f_1(k_1, z). \quad (5)$$

For convenience we abstract from fossil-fuel use in the second and third period, but the first-period fossil-fuel use impacts production negatively in the third period. The second agent starts with the capital stock k_2 , produces output using a strictly concave neoclassical production function $f_2(k_2)$, and can use its income to consume c_2 , or to invest in capital for the third period k_3 :

$$c_2 + k_3 = f_2(k_2). \quad (6)$$

The third consumer derives utility from its consumption, which equals production. Past emissions now enter negatively, as damages, in the production function, $f_{3,k} > 0$, $f_{3,z} < 0$:⁸

$$c_3 = f_3(k_3, z). \quad (7)$$

We assume that also this production function is strictly concave.

2.2 Equilibrium carbon price

Consider now the subgame-perfect equilibrium of the game where generations choose consumptions and emissions in the order of their appearance in the time line, given the preference structure (1)-(3) and choice sets defined through (5)-(7).

The third agent consumes all capital received and cannot influence past emissions. The second agent decides on the capital k_3 transferred to the third agent, given the capital inherited k_2 and the emissions z chosen by the first agent. We thus have a policy function $k_3 = g(k_2, z)$. The policy is defined by

$$\max_{k_3} u_2(c_2) + \beta \delta u_3(f_3(k_3)), \quad (8)$$

leading to equilibrium condition

$$u_2' = \beta \delta u_3' f_{3,k} \Rightarrow 1 = \frac{R_{2,3}}{MRS_{2,3}^{t=2}}, \quad (9)$$

where we introduce the notation $R_{i,j}$ for the rate of return on capital from period i to j , and $MRS_{i,j}^t$ for the absolute value of the marginal rate of substitution between consumptions in periods i and j for generation t .

⁸We follow standard practice in integrated assessment modelling to express all climate change damages in output losses. Including damages directly in utility does not change results.

The strict concavity of utility implies consumption smoothing, and thus if the second agent inherits marginally more capital k_2 , the resulting increase in output is not saved fully but rather split between the second and third generation:

Lemma 1 *Policy function g satisfies $0 < g_k < R_{1,2}$.*

Proof. Substitute the policy function $k_3 = g(k_2, z)$ in (9),

$$\beta\delta u'_3(f_3(g(k_2, z), z))f_{3,k}(g(k_2, z), z) = u'_2(f_2(k_2) - g(k_2, z)). \quad (10)$$

Full derivatives with respect to k_2 lead to

$$\begin{aligned} \beta\delta g_k(u''_3 f_{3,k} f_{3,k} + u'_3 f_{3,kk}) &= u''_2(f'_2 - g_k) \\ \Rightarrow g_k &= \frac{f'_2 u''_2}{\beta\delta u''_3 f_{3,k} f_{3,k} + \beta\delta u'_3 f_{3,kk} + u''_2} < f'_2 = R_{1,2}. \end{aligned} \quad (11)$$

as $u''_t, f_{3,kk} < 0$ and $f_{3,k}, u'_3 > 0$. ■

Understanding the second agent's policy, the first agent decides on consumption and fossil-fuel use to maximize its welfare

$$w_1 = u_1 + \beta\delta[u_2(f_2(k_2) - g(k_2, z)) + \delta u_3(f_3(g(k_2, z), z))].$$

The choice for leaving capital k_2 satisfies

$$\begin{aligned} u'_1 &= \beta\delta(f_{2,k} - g_k)u'_2 + \beta\delta^2 f_{3,k} g_k u'_3 \\ \Rightarrow MRS_{1,2}^{t=1} &= R_{1,2} + \left(\frac{1}{\beta} - 1\right)g_k. \end{aligned} \quad (12)$$

where we use (9). When $\beta = 1$, preferences are consistent, and the term in brackets vanishes as in standard envelope arguments for single decision makers; capital k is then valued according to the usual consumption-based asset pricing equation $MRS_{1,2}^{t=1} = R_{1,2}$. For $\beta < 1$, the second agent has a steeper indifference curve between consumptions in periods 2 and 3: the first-order effect in the bracketed term remains positive, leading to capital returns that no longer reflect the first generation's consumption trade-offs. Letting $MRS_{1,3}^{t=1} = MRS_{1,2}^{t=1} \times MRS_{2,3}^{t=1}$, we have

Lemma 2 *The compound capital return satisfies $MRS_{1,3}^{t=1} < R_{1,3}$ if and only if $\beta < 1$.*

Proof. Using (12), $MRS_{2,3}^{t=1} = \beta MRS_{2,3}^{t=2} = \beta R_{2,3}$, and Lemma 1:

$$\begin{aligned} MRS_{1,3}^{t=1} &= MRS_{1,2}^{t=1} \times MRS_{2,3}^{t=1} = \left[R_{1,2} + \left(\frac{1}{\beta} - 1 \right) g_k \right] \times MRS_{2,3}^{t=2} \\ &\Rightarrow MRS_{1,3}^{t=1} = \left[R_{1,2} + \left(\frac{1}{\beta} - 1 \right) g_k \right] \beta R_{2,3} \\ &< \left[R_{1,2} + \left(\frac{1}{\beta} - 1 \right) R_{1,2} \right] \beta R_{2,3} = R_{1,3}, \end{aligned}$$

where the inequality holds iff $\beta < 1$. ■

Capital returns are generally excessive from the first agent's point of view when $\beta < 1$, that is, the result holds without any restrictions on how emissions alter savings. But for the implications of the excessive capital returns on carbon pricing we must make assumptions on the effect of first-period emissions on the second-period policy, g_z ; these restrictions are common in the integrated assessment models, including ours and, e.g., Golosov et al. (2011), so we explicate them here. Taking the full derivatives of (10) with respect to z , we get

$$g_z = - \frac{\beta(u_3'' f_{3k} f_{3,z} + u_3' f_{3,kz})}{u_2'' + \beta u_3'' f_{3,k} f_{3,k} + \beta u_3' f_{3,kk}}. \quad (13)$$

The first term in the numerator captures the income effect of emissions and is positive. If the first generation emits more, the third generation has lower utility levels and the second generation will tend to save more, as the marginal utility of the third generation increases. The second term in the numerator captures the productivity effect and is negative. If the first generation emits more, productivity of capital in the third period will fall, and the return to investments in the second period will fall alongside. Assuming log utility, and that the production damage is multiplicative

$$u_t(c_t) = \ln(c_t) \quad (14)$$

$$f_3(k_3, z) = f_3(k_3) \omega(z), \quad (15)$$

where $\omega(z)$ is a strictly decreasing damage function, implies that the direct effect of emissions on savings vanishes, $g_z = 0$, as can be easily verified from (13).

Consider then the first generation's equilibrium carbon policy z :

$$u_1' f_{1,z} = \beta \delta g_z u_2' - \beta \delta^2 (f_{3,k} g_z + f_{3,z}) u_3'. \quad (16)$$

If we impose (14)-(15) and thus $g_z = 0$, the carbon policy implied by (16) is described by

$$MCP = \frac{MCD}{MRS_{1,3}^{t=1}} \quad (17)$$

where we let $MCP = f_{1,z}$ denote the marginal carbon product, and $MCD = -f_{3,z}$ denote the marginal carbon damages. If $\beta = 1$, then capital returns reflect consumption trade-offs, $MRS_{1,3}^{t=1} = R_{1,3}$, so that from (17) the carbon price becomes just equal to the damage, discounted with capital return:

$$MCP = \frac{MCD}{R_{1,3}}. \quad (18)$$

This is the general-equilibrium Pigouvian carbon price, under consistent preferences $\beta = 1$. However, if $\beta \neq 1$, in equilibrium, while (17) continues to hold as an internal cost-benefit rule for $t = 1$, Lemma 2 implies that the discounted damage no longer equals the carbon price:

$$MCP > \frac{MCD}{R_{1,3}} \text{ if and only if } \beta < 1. \quad (19)$$

In equilibrium, the first agent establishes a higher carbon price, compared to the Pigouvian level, if and only if $\beta < 1$, i.e., when the first agent gives a higher weight to the long-term utility than the second agent. The result has a very simple intuition. The first consumer would like to transfer more wealth to the third consumer, compared with the preferred wealth transfer of the second consumer: the high capital returns reflect this distortion (Lemma 2). The higher capital returns depress the present-value damages below the true valuation by the first consumer. The opposite deviation — carbon price below the Pigouvian price — occurs if $\beta > 1$.

Proposition 1 *Assume (14)-(15). If $\beta \neq 1$, the first-period carbon price does not satisfy the Pigouvian pricing rule, i.e., $MCP \neq \frac{MCD}{R_{1,3}}$. The carbon price exceeds the Pigouvian level if and only if $\beta < 1$.*

Proof. *Above.* ■

Note that in this general-equilibrium setting the market return $R_{1,3}$ depends both on the productivity of the technology, and on the savings generated by preferences. This is the reason why in standard climate-economy models, the (Pigouvian) carbon pricing is dictated by the same parameters that are used to calibrate the model to match capital returns and savings behavior. The result in (19) is the simplest possible illustration of the point in our paper: in the climate policy game, the equilibrium carbon price is not tied to the equilibrium capital return. This result opens a number of questions that will be subsequently analyzed. Since the result arises in a strategic interaction equilibrium, it cannot reflect a fully efficient outcome. Would enforcing the Pigouvian rule improve welfare?

2.3 Welfare, efficiency, and the Pigouvian rule

To evaluate the welfare implications of policies, we consider the Pareto frontier defined by the welfares of all generations; formally, this leads to the standard Bergsonian objective, as in Caplin and Leahy (2004). Consider an allocation $(\mathbf{c}, \mathbf{k}, z)$ that is Pareto optimal for welfare levels (w_1^*, w_2^*, w_3^*) defined in (1)-(3). If we maximize w_1 , subject to the constraints $w_2 \geq w_2^*$, and $w_3 \geq w_3^*$ and feasibility constraints (5)-(7), then we must find the same allocation, and non-negative Lagrange multipliers $(\alpha_2, \alpha_3) \in R_+^2$ for the welfare constraints. That is, the Pareto optimal allocation is also the solution of a program maximizing

$$W(\mathbf{c}, \mathbf{k}, z) = w_1 + \alpha_2 w_2 + \alpha_3 w_3 \quad (20)$$

subject to (5)-(7). The conclusion also holds the other way around: any solution to a maximization program with some $(\alpha_2, \alpha_3) \in R_+^2$ is Pareto optimal. Strict concavity of the production and utility functions ensures the uniqueness of the allocation. Therefore, we can associate with any Pareto optimal allocation a pair of non-negative welfare weights $(\alpha_2, \alpha_3) \in R_+^2$ that defines the welfare function $W(\mathbf{c}, \mathbf{k}, z)$. The welfare function is merely a tool for welfare analysis, and not intended to introduce dictatorial preferences over allocations: given an equilibrium allocation, we can verify if non-negative α_2, α_3 exist, and thus if a Bergsonian objective can be attached to the allocation.⁹

Rewrite the welfare aggregator in (20) in terms of utility, and compare it to a general utility aggregator $U(\cdot)$:

$$W(\mathbf{c}, \mathbf{k}, z) = u_1 + (\beta\delta + \alpha_2)u_2 + (\beta\delta^2 + \alpha_2\beta\delta + \alpha_3)u_3 \quad (21)$$

$$U(\mathbf{c}, \mathbf{k}, z) = u_1 + \alpha'_2 u_2 + \alpha'_3 u_3 \quad (22)$$

The utility aggregator (22) defines a broad concept of efficiency: since u_t is a strictly monotonic transformation of c_t , a consumption sequence (c_1^*, c_2^*, c_3^*) is efficient if and only if it maximizes $U(\cdot)$ for some positive utility weights $(\alpha'_2, \alpha'_3) \in R_+^2$.¹⁰ The efficient

⁹How are the weights in (20) related to the altruistic weights in (4) that necessarily follow from the (β, δ) -preferences? Letting $\alpha_2 = \alpha_3 = 0$ gives one point in the Pareto set, that is, the one that leads to the dictatorship of the present. But, when $\beta < 1$, generation $t = 1$ still gives altruistic weights to w_2 and w_3 as we can see from (4). Thus, the Pareto weighting (α_2, α_3) , a tool for welfare analysis, should not be confused with altruistic weights (a_2, a_3) that follow from the time-preference primitives. Note that, in their welfare analysis, Krusell et al. (2002) consider one welfare objective, corresponding to $\alpha_2 = \alpha_3 = 0$. We use the weights to describe the Pareto set, and thus we do not restrict the weights.

¹⁰Note that the definitions for efficiency and Pareto optimality are standard; see, for example, Bernheim and Ray (1987).

consumption streams include the Pareto optimal ones, when $(\alpha_2, \alpha_3) \geq 0$ exist such that $\alpha'_2 = \beta\delta + \alpha_2$ and $\alpha'_3 = \beta\delta^2 + \alpha_2\beta\delta + \alpha_3$. As we will see shortly, not all efficient consumption sequences are Pareto optimal. Once we have established efficiency for some allocation, testing whether both welfare weights implied, α_2, α_3 , are non-negative, suffices to test for Pareto optimality. But first, we will elaborate on the conditions for efficiency.

Given utility aggregator $U(\cdot)$, efficiency of (c_1, c_2, c_3) follows from the first-order conditions for $\{k_2, k_3, c_1, c_2, c_3\}$:

$$1 = \left[\frac{\alpha'_2 u'_2}{u'_1}\right] f_{2,k} = \left[\frac{\alpha'_3 u'_3}{u'_1}\right] f_{2,k} f_{3,k} \quad (23)$$

$$\Rightarrow 1 = \frac{R_{1,2}}{MRS_{1,2}} = \frac{R_{1,3}}{MRS_{1,3}}, \quad (24)$$

where we drop the time superscripts on MRS , as the inconsistency between preferences plays no role given objective (22). For emissions in the first period z , the first-order condition requires $u'_1 f_{1,z} + \alpha'_3 u'_3 f_{3,z} = 0$, or,

$$MCP = \frac{MCD}{MRS_{1,3}}, \quad (25)$$

which is equivalent to (17) but without the reference to the first generation's preferences. Efficiency leads to the Pigouvian pricing rule (as opposed to the equilibrium rule (19)):

$$MCP = \frac{MCD}{R_{1,3}}. \quad (26)$$

The Pigouvian principle is a necessary and sufficient test for efficiency.

Lemma 3 *An allocation with strictly positive consumption, capital, and fuel use is efficient if and only if the Pigouvian rule (26) is satisfied.*

Proof. Necessity of the Pigouvian rule has been established above. For sufficiency, we notice that given the allocation, we can construct positive weights (α'_2, α'_3) from (23). If (26) is satisfied as well, then all first-order conditions for the maximization of (22) are satisfied. ■

The equivalence with efficiency defines the Pigouvian rule in an impartial way — the definition is independent of the generation-specific welfare representations. The equivalence will be instrumental in our equilibrium analysis. If the Pigouvian rule is not satisfied, the equilibrium allocation is not efficient and it cannot be Pareto optimal (no Bergsonian objective exists). We have seen that in equilibrium the Pigouvian rule will not hold, so the conclusion for efficiency is immediate:

Corollary 1 *The equilibrium described in Proposition 1 is not efficient if $\beta \neq 1$.*

It is interesting to note that even though decision-makers internalize all future impacts of current actions, the equilibrium is observationally distinct from the planner's optimum for objective (22). This contrasts Barro (1999) where observational equivalence with an allocation chosen by a fictitious planner follows without restrictions on the actions of the decision-makers with hyperbolic preferences. Our result shows that when there is more than one capital-good the observational-equivalence does not hold in general.¹¹

Given Lemma 3, we have a straightforward policy instrument implementing efficiency: we can consider enforcing the Pigouvian rule as an institutional constraint on the equilibrium behavior — it can be thought of as an environmental agency scrutinizing the climate policies within each period. The environmental agency has no preferences, and it simply enforces the Pigouvian taxes, without restricting the choices of each generation in any other way.

Proposition 2 *Enforcing Pigouvian externality pricing leads to efficiency but not to Pareto optimality if preferences are quasi-hyperbolic, $\beta < 1$.*¹²

Proof. We have seen in Lemma 3 that the Pigouvian tax rule implements (22). From the first-order conditions for c_2 , c_3 and k_3 we have

$$\alpha'_2 u'_2 = \alpha'_3 u'_3 f_{3,k}. \quad (27)$$

As the second generation does not need fossil fuels, the policy function g does not change from that defined in (9). Combining (9) and (27) gives

$$\alpha'_3 = \beta\delta\alpha'_2.$$

Adding $\beta\delta^2 - (\beta\delta)^2 > 0$ to the right-hand-hand side gives

$$\alpha'_3 < \beta\delta^2 + (\alpha'_2 - \beta\delta)\beta\delta. \quad (28)$$

From (21) and (22), $(\alpha'_2, \alpha'_3) \geq 0 \Rightarrow (\alpha_2, \alpha_3) \geq 0$ if and only if

$$\alpha'_2 \geq \beta\delta \quad (29)$$

$$\alpha'_3 \geq \beta\delta^2 + (\alpha'_2 - \beta\delta)\beta\delta. \quad (30)$$

¹¹The capital asset nature of carbon becomes more explicit in the general model.

¹²If preferences are time-inconsistent in the other direction, $\beta > 1$, the proof is much more tedious. However, we prove both directions for the infinite horizon case.

Thus, (28) and (30) are in contradiction. ■

Enforcing Pigouvian externality pricing exploits all opportunities for efficiency improvements for consumption, although it does not achieve full Pareto optimality as just demonstrated. But could it lead to a Pareto improvement? In three periods, it is immediate that the answer is negative. The Pigouvian pricing will constrain only the first generation's choices.

Proposition 3 *The Pigouvian tax does not imply a Pareto improvement vis-a-vis the equilibrium without it: the welfare of the first generation decreases.*

The reason for this result is simple: the Pigouvian tax is only a constraint on the first generation, as it could have implemented the same allocation without the rule, or without consulting the later generations. Therefore, enforcing the Pigouvian tax must decrease welfare of the first generation if $\beta \neq 1$. When $\beta = 1$, imposing the Pigouvian tax has no effect on the equilibrium, so the requirement becomes redundant.

2.4 Discussion

The main plot should be clear after this three-period analysis. The carbon price that is optimal from the current generation's perspective is not dictated by the capital returns; see Proposition 1. The infinite-horizon model adds to the analysis in multiple ways. First, in three periods there is little room for gains from future efficiency improvements, as indicated by Proposition 3. With a longer sequence of generations, generation 1 could benefit from the later efficiency improvements by future generations' adherence to the Pigouvian tax. Indeed, as we will see at the end of the next section, in the infinite horizon model we can construct welfare-improving policies. Second, the delay structure of climate change impacts are important, and an infinite horizon model enables us to describe the climate dynamics more precisely. We can then analyze how the commitment value of climate policies depends on the persistence over time of the impacts of current choices. The infinite horizon model gives analytical results that prescribe the equilibrium carbon prices dependent on technology, climate, and preference parameters. Third, the infinite horizon model enables us to provide a quantitative assessment of the magnitudes for both Pigouvian and equilibrium carbon prices, using realistic parameters values.

3 An infinite horizon climate-economy model

3.1 Technologies and preferences: general setting

Consider a sequence of periods $t \in \{1, 2, \dots\}$. The economy's production possibilities $f_t(k_t, l_t, z_t, s_t)$ depend on capital k_t , labour l_t , current fossil-fuel use z_t , and the emission history (i.e., past fossil-fuel use),

$$s_t = (z_1, z_2, \dots, z_{t-2}, z_{t-1}).$$

History s_t enters production for two reasons. First, climate-change that follows from historical emissions changes production possibilities, as usual in integrated-assessment models. Second, the current fuel use is linked to historical fuel use through energy resources whose availability and the cost of use depends on the past usage. In the specific model that we detail below, we abstract from the latter type of history dependence, because the scarcity of conventional fossil-fuel resources is not a key factor for the long-run climate policies. The economy has one final good; closed-form solutions require that capital depreciates in one period, leading to the following budget accounting equation between period t and $t + 1$:

$$c_t + k_{t+1} = y_t = f_t(k_t, l_t, z_t, s_t), \quad (31)$$

where c_t is the total consumption, k_{t+1} is capital built for the next period, and y_t is gross output. In each period, the representative consumer makes the consumption, fuel use, and investment decisions. Let per-period utility be u_t and define generation t welfare generated by the choice sequence $\{c_t, z_t, k_t\}_{t=1}^{\infty}$ as

$$w_t = u_t + \beta \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} u_{\tau} \quad (32)$$

where we identify dynamically consistent preferences by $\beta = 1$. As in three periods, when $\beta < 1$, decision-makers use lower discount rates for long- than for short-term evaluations, and this increasing patience implies altruistic weights on future generations' welfare levels; see Saez-Marti and Weibull (2005) for the weights implied by quasi-hyperbolic discounting.¹³

¹³The formal analysis is not restricted to this quasi-hyperbolic setting. For interpretations, the quasi-hyperbolic case is the most natural to keep in mind, but we will state explicitly the formal results that require $\beta < 1$

3.2 The basis climate-economy model

Golosov et al. (2011) marks an important deviation from Nordhaus' approach (e.g., 1993) to integrated-assessment modeling: abatement does not enter as a separate decision but is implied by the energy input choices. We follow this approach but our modeling of the climate dynamics, in addition to preferences, departs substantially from both Golosov et al. and Nordhaus. We pull together the production structure as follows:

$$y_t = k_t^\alpha A_t(l_{y,t}, e_t)\omega(s_t) \quad (33)$$

$$e_t = E_t(z_t, l_{e,t}) \quad (34)$$

$$l_{y,t} + l_{e,t} = l_t \quad (35)$$

$$\omega(s_t) = \exp(-\Delta_y D_t), \quad (36)$$

$$D_t = \sum_{\tau=1}^{t-1} \theta_\tau z_{t-\tau} \quad (37)$$

Production. The gross production consists of: (i) the Cobb-Douglas capital contribution k_t^α with $0 < \alpha < 1$; (ii) function $A_t(l_{y,t}, e_t)$ for the energy-labour composite in the final-good production with $l_{y,t}$ denoting labor input and e_t the total energy use in the economy; (iii) total energy $e_t = E_t(z_t, l_{e,t})$ using fossil fuels z_t and labour $l_{e,t}$, and (iv) the damage part given by function $\omega(s_t)$ capturing the output loss of production depending on the history of emissions from the fossil-fuel use. The functional forms for the capital contribution and damages allow a Markov structure for policies, and thus the determination of the currently optimal policies as function of the state of the economy, say, at year 2010. Functions A_t and E_t are essential for capturing the two main options available for reducing emissions. First, emissions can decline through energy savings, obtained by substituting labor $l_{y,t}$ for total energy e_t . Second, emissions can decline through “de-carbonization” that involves substituting non-carbon inputs for carbon energy inputs z_t in energy production. In the quantitative analysis Section 4, where we specify A_t and E_t in detail, de-carbonization is obtained by allocating the total energy labor $l_{e,t}$ further between carbon and non-carbon energy sectors.¹⁴

Damages and carbon cycle. Equations (36)-(37) show that climate damages are interpreted as reduced output, dependent on the history of emissions through the state variable D_t , a non-linear proxy for the global mean temperature rise. The form in (37) assumes that anthropogenic emissions start at $t = 1$, taken to be the beginning of the industrial era. The exponential form for damages combined with linear dependence of

¹⁴Typically, the climate-economy adjustment paths will feature early emissions reductions through energy savings, whereas de-carbonization is necessary for achieving long-term reduction targets.

D_t on the past emissions is the same as in Golosov et al. (2011). But the specification of the parameters θ_τ in our model is very different: we develop a closed-form representation of the global carbon-climate cycle, allowing a transparent and detailed calibration that differs from that in DICE (Nordhaus, 2007), but is consistent with the main processes for CO_2 and temperature dynamics of DICE. This representation captures the delays with which emissions cause damages, having substantial implications for policies, which we show by contrasting our calibration results with both Nordhaus and Golosov et al. (2011). We explain now in detail how the weighting of past emissions θ_τ in the expression for D_t in (37) is obtained.

The carbon cycle is typically described as a diffusion process of carbon between various layers of the atmosphere, ocean and biosphere (Maier-Reimer and Hasselmann 1987). In reduced-form models, these layers are converted into a set of “boxes” (Hasselmann et al. 1997; see our Appendix “Calibration” for the conversion method). The core of the conversion from a description of layers to an equivalent formal representation of boxes is that, whereas carbon diffuses between the layers, emissions live only in the particular box they enter; interactions between the boxes are eliminated. Each box describes the atmospheric decay of a certain share of total emitted CO_2 . Label an atmospheric CO_2 box by i , and assume that share a_i of the total anthropogenic emissions in period t enter each box i such that $\sum_i a_i = 1/(1 + \mu)$, where μ is the share of emissions that is taken up, over short time, by the upper ocean layer.¹⁵ The initial values $S_{i,1}$ are the stocks of CO_2 in the boxes at the start of the first period. For the policy simulations, $S_{i,1}$ refers to the stock in 2015, following from the calibration of the carbon cycle model where we use data on emissions and concentrations for the past century. The depreciation rate of carbon in box i is η_i . Stock $S_{i,t}$ in box i , and total atmospheric CO_2 , S_t , develop for $t = 2, 3, \dots$ according to

$$\begin{aligned} S_{i,t} &= (1 - \eta_i)S_{i,t-1} + a_i z_{t-1}, \\ S_t &= \sum_i S_{i,t}, \end{aligned}$$

where z_t are the total anthropogenic emissions, expressed in Teratons of CO_2 (Tton CO_2). For our purposes, the calibration of a 3-box representation is sufficiently precise, as will be detailed in Section 4.¹⁶

¹⁵Strictly speaking, if the period length is long, some of the stock in box i may depreciate within the period, and $\sum_i a_i < 1/(1 + \mu)$. The condition $\sum_i a_i = 1/(1 + \mu)$ is a physical identity only for short periods of, say, less than a year. See the appendix for further clarifications.

¹⁶For example, a 3-box model almost perfectly fits the CO_2 dynamics in DICE (Nordhaus 2007). The illustration is available on request.

Similarly to the atmospheric CO_2 , we can describe the overall atmospheric temperature adjustment as a linear combination of several temperature adjustment processes, formally represented by “boxes” denoted by j . We follow Hooss et al (2001, table 2) and assume, for all boxes, a common asymptotic climate sensitivity $\varphi(S_t)$, but a different adjustment speed ε_j and weight b_j in the overall temperature adjustment:¹⁷

$$\begin{aligned} T_{j,t} &= T_{j,t-1} + \varepsilon_j(\varphi(S_t) - T_{j,t-1}) \\ T_t &= \sum_j b_j T_{j,t} \end{aligned}$$

where $\sum_j b_j = 1$ and $t = 2, 3, \dots$. Typically, the relationship between the asymptotic temperature sensitivity and the atmospheric CO_2 stock, $\varphi(S)$, is concave; the logarithmic form is frequently assumed. In steady state, we have $T = \varphi(S)$, but elsewhere temperature T_t changes depending on the atmospheric CO_2 stock. Damages, in turn, are a function of the temperature

$$D_t = \psi(T_t)$$

where $\psi(T_t)$ is convex. It has been noted in the literature that in the relevant domain of atmospheric CO_2 concentrations between 400 and 1000 ppmv,¹⁸ the composition of the typical convex damage and concave climate sensitivity functions returns an almost linear function through the origin:¹⁹

$$\psi(\varphi(S_t)) \approx \pi S_t$$

with $\pi > 0$, a constant sensitivity of damages to atmospheric CO_2 .²⁰ Using the approximation, we can rewrite the damage dynamics directly as dependent on stocks:

$$\begin{aligned} D_{j,t} &= D_{j,t-1} + \varepsilon_j(b_j \pi S_t - D_{j,t-1}) \\ D_t &= \sum_j D_{j,t}. \end{aligned}$$

where D_t can be taken to be a non-linear (e.g., quadratic) proxy for the temperature; inversely, the temperature can be derived from D_t through $\psi^{-1}(\cdot)$. Given the two layers

¹⁷We can formally relax the assumption that all boxes have the same asymptotic climate sensitivity, but the common parameter substantially simplifies the exposition below.

¹⁸ppmv=parts per million by volume.

¹⁹Indeed, the early calculations by Nordhaus (1991) based on local linearization, are surprisingly close to later calculations based on his DICE model with a fully-fledged carbon-cycle temperature module, apart from changes in parameter values based on new insights from the natural science literature.

²⁰Multiplying the constants Δ_y and π gives the damage sensitivity: the asymptotic percentage loss of output per TtCO₂ in the atmosphere. Inversely, $1/\Delta_y \pi$ is the amount of atmospheric CO_2 that leads to an asymptotic 63 per cent (e^{-1}) loss of output.

of climate variables — one for carbon stocks $S_{i,t}$, and the other for damages $D_{j,t}$ — it is a straightforward matter of verification that future damages depend on past emissions as follows:

$$S_{i,t} = (1 - \eta_i)^{t-1} S_{i,1} + \sum_{\tau=1}^{t-1} a_i (1 - \eta_i)^{\tau-1} z_{t-\tau} \quad (38)$$

$$D_{j,t} = (1 - \varepsilon_j)^{t-1} D_{j,1} + \sum_i b_j \pi \varepsilon_j \frac{(1 - \eta_i)^t - (1 - \eta_i)(1 - \varepsilon_j)^{t-1}}{\varepsilon_j - \eta_i} S_{i,1} + \sum_i \sum_{\tau=1}^{t-1} a_i b_j \pi \varepsilon_j \frac{(1 - \eta_i)^\tau - (1 - \varepsilon_j)^\tau}{\varepsilon_j - \eta_i} z_{t-\tau}, \quad (39)$$

where $S_{i,1}$ and $D_{j,1}$ are taken as given at $t = 1$, and then values for $t > 1$ are defined by the expressions. The model is applied to a situation where some climate change has taken place at the start of time $t = 1$, so we write the system dependent on $S_{i,1}, D_{j,1} > 0$ — however, interpreting $t = 1$ as the beginning of the industrial era, say 1850, we can set $S_{i,0} = D_{j,0} = 0$. Taking the sum $\sum_j D_{j,t}$ in (39), and collecting terms over (i, j) , allows us to express D_t as in (37).

The two layers of climate variables in (38)-(39) are an important part of the carbon-temperature-cycle model, as they introduce a time-lag between the current emissions and the peak in damages caused: an impulse of z has subsequent periodic damage effects peaking at a lag between $1/\eta_i$ and $1/\varepsilon_j$ periods.²¹

For the simple 1-box representation, a typical estimate for the atmospheric CO_2 depreciation η_1 is one per cent per year, while for the temperature adjustment ε_1 it is two per cent per year; the associated peak in temperature-response is after about 70 years. In other words, peak temperatures lag 70 years behind emissions, providing a rule of thumb for discounting damages: discounting future damages at a rate of r per cent per year results in a discount factor of about 2^{-r} after 70 years. An annual 1, 2 or 3 per cent discount rate results in a discount factor of damages of $1/2$, $1/4$, or $1/8$, respectively.

Golosov et al.’s (2011) specification can be understood as one where the temperature adjustment is immediate: $\varepsilon_j = 1$ so that an emission impulse leads to an immediate temperature shock that slowly decays. Our reduced-form model above can be calibrated to approximate the DICE model (Nordhaus 2007). Figure 1 shows the life path of damages (percentage of total output) caused by inserting one Teraton of Carbon (TtonC) in the first period, and then contrasting the impact with a counterfactual path without the carbon impulse.²² The specification following Golosov et al. produces an immediate

²¹The term within the summation in expression (39), $\frac{(1 - \eta_i)^\tau - (1 - \varepsilon_j)^\tau}{\varepsilon_j - \eta_i}$, peaks at a period between $1/\eta_i$ and $1/\varepsilon_j$.

²²See Appendix “Comparison of climate response functions” for the details of the experiment.

peak but a fat tail, while the DICE model shows an emissions-damage peak after 60 years with a thinner tail. Our model, that we calibrate with data from the natural sciences literature, produces a combination of the effects: a peak in the emission-damage response function after about 60 years and a fat tail; about 16 per cent of emissions do not depreciate within the horizon of thousand years.

We postpone the detailed calibration of the carbon cycle to Section 4.

Figure 1 here: Emissions-Damage responses

Periodic utility. We assume that the utility function is logarithmic, and through a separable linear term we also include the possibility of intangible damages associated with climate change:

$$u_t = \ln\left(\frac{c_t}{l_t}\right) - \Delta_u D_t. \quad (40)$$

The utility loss $\Delta_u D_t$ is not necessary for the substance matter of this paper, but it proves useful to explicate how it enters the carbon price formulas. In calibration, we let $\Delta_u = 0$ to maintain an easy comparison with the previous studies.²³ It is important to note that we write utility as an average in our analysis. Alternatively, we can write aggregate utility within a period by multiplying utility with population size, $u_t = l_t \ln(c_t/l_t) - l_t \Delta_u D_t$. The latter approach is feasible but it leads to considerable complications in the formulas below. Scaling the objective with labor rules out stationary strategies — they become dependent on future population dynamics —, and also impedes a clear interpretation of inconsistencies in discounting. While the formulas in the Lemmas depend on the use of an average utility variable, the substance of the Propositions is not altered.²⁴

Strategies. Our starting point is the symmetric and stationary Markov equilibrium. Symmetry means that all generations use the same policy functions — even though there is technological change and population growth, the form of the objective in (40) ensures that there will be an equilibrium where the same policy rule will be used for all t . The Markov restriction means that the policy does not condition on the history of past behavior: strategies are identical in states where the continuation payoffs are identical (see Maskin and Tirole, 2001).²⁵ In equilibrium, the policy will take the form

²³See Tol (2009) for a review of the existing damage estimates. While the estimates for intangible damages are mostly missing, our carbon pricing formulas can help to transform output losses into equivalent intangible losses to gauge the relative magnitudes of such losses that can be associated with a given carbon price level.

²⁴The expressions for this case are available on request

²⁵We will construct a natural Markov equilibrium where policies have the same functional form as

$k_{t+1} = g_t(k_t, \Theta_t)$, $z_t = h_t(k_t, \Theta_t)$, where Θ_t collects the vector of climate state variables. However, since the climate affects the continuations payoffs only through the weighted sum of past emissions, as expressed in (37), we will replace Θ_t by history s_t below, treating it as a state variable.

Structure of equilibria. Given policies $g_t(k_t, s_t)$ and $h_t(k_t, s_t)$, we can write welfare in (32) as follows

$$\begin{aligned} w_t &= u_t + \beta \delta W_{t+1}(k_{t+1}, s_{t+1}), \\ W_t(k_t, s_t) &= u_t + \delta W_{t+1}(k_{t+1}, s_{t+1}) \end{aligned}$$

where $W_{t+1}(k_{t+1}, s_{t+1})$ is the (auxiliary) value function. All equilibria considered in this paper will be of the form where a constant share $0 < g < 1$ of the gross output is invested,

$$k_{t+1} = gy_t, \tag{41}$$

whereas the climate policy defines emissions through a constant h that defines the carbon policies through

$$f_{t,z} = h(1 - g)y_t, \tag{42}$$

where $f_{t,z} = MCP_t$ is the marginal product of fossil fuel use, the carbon price.²⁶ Equilibrium policies will be characterized simply by a pair of constants (g, h) . That a constant fraction of output is saved should not be surprising, given the log utility and Cobb-Douglas contribution of capital in production.²⁷ Condition (42) implies that the marginal carbon price per consumption is a constant, $h = f_{t,z}/c_t$ where $c_t = (1 - g)y_t$. This may seem surprising given the complicated delay structure (37), and changing productivities in (33)-(37), and preference inconsistencies.²⁸

Postponing the discussion on welfare and the verification that the policies actually take the above form, it proves useful to state the properties of the value function implied by (g, h) policies (the proofs, unless helpful in the text, are in the Appendix).

Lemma 4 (*separability*) *Given the model (33)–(40), assuming that future policies $g_\tau(\cdot)$ and $h_\tau(\cdot)$ for $\tau = t + 1, t + 2, \dots$ satisfy (41) and (42), then the value function is separable*

when $\beta = 1$. For multiplicity of equilibria in this setting, see Krusell and Smith (2003) and Karp (2007).

²⁶Later, we also consider non-stationary equilibria, which are thus not Markovian, but both on-the-path and off-the-path strategies have the form in (41)-(42). Thus, the value expressions derived here can be used to analyze these equilibria as well.

²⁷See, e.g., Barro 1999, for the analysis of the one-capital good case.

²⁸Golosov et al. find emission policies that have the same features; our policies exploit the same functional assumptions, despite the added complexity.

in capital and historical emissions

$$W_{t+1}(k_{t+1}, s_{t+1}) = V_{t+1}(k_{t+1}) - \Omega(s_{t+1}).$$

with parametric form

$$\begin{aligned} V_{t+1}(k_{t+1}) &= \xi \ln(k_{t+1}) + \tilde{A}_{t+1} \\ \Omega(s_{t+1}) &= \sum_{\tau=1}^{t-1} \zeta_{\tau} z_{t+1-\tau}. \end{aligned}$$

where $\xi = \frac{\alpha}{1-\alpha\delta}$, $\zeta_1 = (\frac{\Delta_y}{1-\alpha\delta} + \Delta_u) \sum_i \sum_j \frac{a_i b_j \pi \varepsilon_j}{[1-\delta(1-\eta_i)][1-\delta(1-\varepsilon_j)]}$, and \tilde{A}_{t+1} is independent of k_{t+1} and s_{t+1} .

The result shows that we can obtain the value of savings k_{t+1} and the costs from fossil-fuel use z_t separately.

3.3 The Markov equilibrium carbon price

Consider first savings. Each generation takes the future policies, captured by constants (g, h) in (41)-(42), as given and chooses its current savings to satisfy

$$u'_t = \beta \delta V'_{t+1}(k_{t+1}),$$

where function $V(\cdot)$ from Lemma 4 captures the continuation value implied by the equilibrium policy.

Lemma 5 (*savings*) *The equilibrium investment share $g = k_{t+1}/y_t$ is*

$$g^* = \frac{\alpha\beta\delta}{1 + \alpha\delta(\beta - 1)}. \quad (43)$$

The proof of the Lemma is a straightforward verification exercise following from the first-order condition. The savings depend only on the capital share α and preference parameters; note that when preferences are consistent $\beta = 1$, we obtain $g = \delta\alpha$, as is expected.

Consider then the equilibrium carbon price $f_{t,z}$, that is, the marginal product of the fossil-fuel use z_t . The first-order condition is

$$u'_t f_{t,z} = \beta \delta \frac{\partial \Omega(s_{t+1})}{\partial z_t},$$

where function $\Omega(\cdot)$ gives the future costs implied by the equilibrium policy. Given Lemma 4, the equilibrium carbon price and the fossil-fuel use can be obtained:

Lemma 6 *Equilibrium emissions $z_t = z_t^*$ depend only on the current technology at period t as captured through $A_t(\cdot)$ and $E_t(\cdot)$. The carbon price satisfies (42) where $h^* = \beta\delta\zeta_1$:*

$$MCP_t = f_{t,z} = \beta\delta\zeta_1(1 - g)y_t \quad (44)$$

When y_t is known, say $y_{t=2010}$, the carbon policy for $t = 2010$ can be obtained from (44), by reducing fossil-fuel use to the point where the marginal product of z equals the consumption-weighted externality cost of carbon, as expressed by the right-hand-side of (44).

The Markov equilibrium carbon price, as indicated by (44) and Lemma 4, depends on the delay structure in the carbon cycle captured by parameters η_i and ε_j for each box. Carbon prices increase with the damage sensitivity ($\partial\zeta_1/\partial\pi > 0$), slower carbon depreciation ($\partial\zeta_1/\partial\eta_i < 0$), and faster temperature adjustment ($\partial\zeta_1/\partial\varepsilon_j > 0$). Higher short- and long-term discount rates both decrease the carbon price ($\partial\zeta_1/\partial\beta > 0$; $\partial\zeta_1/\partial\delta > 0$).

3.4 Efficiency and the Pigouvian tax rule

As in three periods, when $\Delta_u = 0$, we define a feasible program $\{c_t, z_t, k_t\}_{t=1}^\infty$ to be efficient if there does not exist another feasible program such that $c'_t \geq c_t$ for all $t > 0$ and with strict inequality for at least some $t = \tau$. If intangible climate change damages are positive, $\Delta_u > 0$, efficiency is defined as a maximal utility stream. In both cases, an efficient allocation maximizes a utility aggregator

$$\sum_{t=1}^\infty \alpha'_t u_t$$

where $\{\alpha'_t\}_{t=1}^\infty$ is some sequence of utility weights; we rule out dynamically inefficient allocations by requiring a bounded mass for α_t . Given the structure of our equilibria with constant savings, we can focus directly on geometric weights, $\gamma = \alpha'_{t+1}/\alpha'_t$, with $0 < \gamma < 1$, that in equilibrium is obtained from

$$u'_t = \gamma u'_{t+1} R_{t,t+1}$$

where $R_{t,t+1}$ is the capital return between t and $t + 1$. Thus,

$$\gamma = \frac{u'_t}{u'_{t+1} R_{t,t+1}} = \frac{c_{t+1}}{c_t R_{t,t+1}} = \frac{c_{t+1}}{c_t} \frac{k_{t+1}}{\alpha y_{t+1}} = \frac{g}{\alpha}. \quad (45)$$

In the Markov equilibrium where $g = g^*$, we have

$$\gamma^* = \frac{\beta\delta}{1 + \alpha\delta(\beta - 1)}. \quad (46)$$

This is the imputed geometric utility discount factor that is consistent with the efficiency of the equilibrium consumption stream. Given the equilibrium utility discount factor γ^* , we can find the carbon externality costs caused by the equilibrium actions.²⁹

Lemma 7 *The net present value of marginal damages of emissions, the Pigouvian tax, τ_t^γ is proportional to gross output,*

$$\tau_t^\gamma = h^\gamma(1-g)y_t \quad (47)$$

$$h^\gamma = \sum_i \sum_j \frac{(\frac{\Delta_y}{1-g} + \Delta_u)\gamma\pi a_i b_j \varepsilon_j}{[1-\gamma(1-\eta_i)][1-\gamma(1-\varepsilon_j)]} \quad (48)$$

Imposing the Pigouvian tax policy on the equilibrium is, then, the requirement that marginal product of emissions should equal the expression in (47), that is, $f_{t,z} = \tau_t^\gamma$. Note that this procedure is valid since the equilibrium savings, and thus γ^* , is unaffected by the carbon pricing policy (due to separability of policies from Lemma 4). We can now compare the two carbon prices:

Proposition 4 *For quasi-hyperbolic preferences, $\beta < 1$, the equilibrium carbon price strictly exceeds the Pigouvian carbon price if climate change delays are sufficiently long. Formally, the ratio of the equilibrium carbon price and the Pigouvian carbon price, $f_{t,z}/\tau_t^\gamma$, is continuous in parameters $\beta, \delta, \eta_i, \varepsilon_j, a_i, b_j$, and γ . Evaluating at $\beta < 1, \eta_i = \varepsilon_j = 0$, and γ^* by (46), gives*

$$\frac{f_{t,z}}{\tau_t^{\gamma^*}} > 1$$

If preferences are quasi-hyperbolic and the climate system is sufficiently persistent, then the current generation sees the climate asset as a commitment device to transfer wealth to far-future generations, and therefore it values the external climate costs above the Pigouvian level. For $\beta > 1$, the inverse holds.

As in three periods, requiring that emissions should follow the rule $f_{t,z} = \tau_t^\gamma$ has a clear justification: it implements efficiency. Only when Pigouvian carbon prices are imposed, is it not possible to increase consumption at any t without decreasing consumption at some $t' \neq t$.

Proposition 5 *An equilibrium with utility discount factor γ (not necessarily γ^*) is efficient if and only if the Pigouvian tax policy in Lemma 7 is implemented.*

²⁹Factor γ can also deviate from γ^* , if the equilibrium is non-Markovian. Since we come back to such policies shortly, we label the Pigouvian tax by the superscript γ , so that the formula provides the Pigouvian carbon price for any equilibrium with constant savings g .

Proof. Given utility discount factor γ , we can obtain the Planner's allocation as the one that follows in Markov equilibrium when $\beta\delta = \delta = \gamma$. The Markov equilibrium implies the first-order conditions for capital investments, and these conditions are sufficient for efficient savings in this economy. The remaining efficiency condition is for emissions. Efficiency implies and is implied by the Pigouvian carbon price. ■

Yet, efficiency does not imply Pareto optimality. In the three-period model, the implied utility weights α'_2 and α'_3 had to satisfy some conditions. For the infinite horizon model, the condition for efficiency turns out to be a simple condition on the utility discount factor (proof in Appendix but see also Gerlagh and Liski, 2011):³⁰

Lemma 8 *An efficient allocation with utility discount factor γ is Pareto optimal if and only if $\gamma \geq \max\{\beta\delta, \delta\}$.*

If time preferences are consistent, $\beta = 1$, then the equilibrium utility discount factor equals both short and long-term preferences, $\gamma = \beta\delta = \delta$, and the allocation is efficient and Pareto optimal. If time preferences are inconsistent, $\beta \neq 1$, then the equilibrium utility discount factor is bounded by the two contrafactuals: by straightforward substitution, we can see that for $\beta \neq 1$, the Markov-equilibrium discount factor satisfies

$$\min\{\beta\delta, \delta\} < \gamma^* < \max\{\beta\delta, \delta\} < 1. \quad (49)$$

The reasoning for this result is straightforward. Consider $\beta < 1$. The current generation cares more for total future welfare, and thus saves more, than a planner would do if the planner had consistent preferences with discount factor $\gamma = \beta\delta$. But then, the current agent cares less, and saves less, compared to a representative planner who would have consistent preferences with $\gamma = \delta$. Clearly, the equilibrium savings must be somewhere between the extremes.

Proposition 6 *If $\beta \neq 1$, the equilibrium with Pigouvian carbon prices is not Pareto optimal.*

While the Pigouvian carbon price does not restore full efficiency, it might be argued that the productive inefficiency removed produces at least a Pareto improvement. This we will consider in Section 3.6, but first, we want to assess the significance of the gap between the equilibrium carbon prices that are based on current welfare preferences, and the Pigouvian carbon prices based on the efficiency criterion.

³⁰Following Bernheim and Ray (1987), we could define quasi-Pareto optimality, which also takes into account the welfare of past generations. In our economy, this does not change the results, and we leave the concept of quasi-Pareto optimality out of the paper.

3.5 First look at numbers

For the Markov and Pigouvian carbon prices in (44) and (47), we do not have to solve the full model: the income level, savings, and the carbon cycle parameters allow obtaining the carbon price levels. The calibration of the carbon cycle is explicated in the Appendix.³¹ We note here only that the model is decadal (10-year periods),³² and that our box presentation of the carbon cycle, when compared to DICE 2007, results in a slightly lower response of damages to emissions for the first 80 years, and substantially higher damages after 300 years.³³ That is, our parameters suggest a slightly slower and more persistent climate response, as also illustrated in Figure 1.³⁴

We calibrate the damage parameter so that they are equivalent to 2.7 per cent of output at a temperature rise of 3 Kelvin, as in Nordhaus (2001), we obtain $\Delta_y = 0.003$; we set $\Delta_u = 0$. We calibrate the Cobb-Douglas capital elasticity so that annual pure time discounting of 2 per cent leads to a 25 per cent gross savings rate. These parameter choices result in a Pigouvian carbon price of 8.4 Euro/tCO₂, equivalent to 40 USD/tC, for 2010.³⁵ This number is very close to the level found by Nordhaus.³⁶

We can decompose the carbon price (44) into three contributing parts. First, consider the one-time costs if damages were immediate (*ID*) but only for one period,³⁷

$$ID = \left(\frac{\Delta_y}{1 - \alpha\delta} + \Delta_u \right) \pi(1 - g)y_t,$$

where $1 - \alpha\delta$ is replaced by $1 - \alpha\gamma = 1 - g$ when we evaluate this effect for the Pigouvian tax rule. This value is multiplied by a factor to correct for the persistence of climate

³¹Appendix "Calibration: carbon cycle"

³²Bill Nordhaus suggested to us that the period length could be longer, e.g., 20-30 years to better reflect the idea that the long-term discounting starts after one period for each generation. We have these results available on request.

³³see Appendix "Comparisons of damage responses".

³⁴The main reason for the deviation is that DICE assumes an almost full CO₂ storage capacity for the deep oceans, while large-scale ocean circulation models point to a reduced deep-ocean overturning running parallel with climate change (Maier-Reimer and Hasselman 1987). The positive feedback from temperature rise to atmospheric CO₂ through the ocean release is essential to explain the large variability observed in ice cores in atmospheric CO₂ concentrations.

³⁵Note that 1 tCO₂ = 3.67 tC, and 1 Euro is about 1.3 USD.

³⁶Minor differences are caused by a correction for the price index, and a somewhat more persistent damage structure in our reduced model

³⁷The adjustment of the output loss Δ_y by $(1 - \alpha\delta)^{-1}$ is to account for the decrease in the future capital stock caused by a current drop in output.

change, the persistence factor (PF),

$$PF = \sum_i \frac{a_i}{[1 - \delta(1 - \eta_i)]},$$

which we then multiply by a factor to correct for the delay in the temperature adjustment, the delay factor (DF),

$$DF = \sum_j \frac{\beta \delta b_j \varepsilon_j}{1 - \delta(1 - \varepsilon_j)}.$$

Table 2 below presents the decomposition of the Pigouvian tax, as well as the Markov tax when the short-term annual discount rate is 2.55 and the long-term is rate .5 per cent; we discuss these choices shortly.

| | ID | PF | DF | Carbon price |
|--------|------|------|-----|--------------|
| Pigou | 7.61 | 2.44 | .45 | 8.4 |
| Markov | 8.05 | 5.79 | .63 | 29.5 |

Table 2: Decomposition of Carbon price, MCP [Euro/tCO₂]. ID =immediate costs, PF =persistence factor, DF =delay factor, $MCP = ID \times PF \times DF$. Parameter values in text.

Leaving out the time lag between CO₂ concentrations and temperature amounts to replacing the column DF by $\beta\delta$: abstracting from the delay in temperature adjustments, as in Golosov et al. (2011), increases the Pigouvian carbon price by almost factor 2.

From Table 2, the Markov carbon prices stand at 29.5 Euro/tCO₂, three and half times the Pigouvian level (and well above the current carbon prices at the EU emissions trading system). How did we choose the short- and long-run discount rates underlying the Markov price? Weitzman's (2001) survey led to discount rates declining from 4 per cent for the immediate future (1-5 years) to 3 per cent for the near future (6-25 years), to 2 per cent for medium future (26-75 years), to 1 per cent for distant future (76-300), and then close to zero for far-distant future. Roughly consistent with Weitzman and our 10-year length of one period, we use the short-term discount rate close to 3 per cent, and the long-term rate at or below 1 per cent. This still leaves degrees of freedom in choosing the two rates. We present three cases that all match the savings rate of 25 per cent and thus the macroeconomic performance in Nordhaus (2007): we choose β and δ to maintain the equilibrium utility discount factor at $\gamma = 0.817$ (2 per cent annual discount rate). Moreover, since the equilibrium utility discount rate remains at 2 per cent, the Nordhaus'

case is the Pigouvian benchmark; using the 2 per cent rate to determine the present-value damages implies Pigouvian taxes as proposed by Nordhaus.³⁸ See now Table 3.

| | annual discount rate | | | ID | PF | DF | Carbon price |
|------------|----------------------|-----------|-------------|------|-------|-----|--------------|
| | short-term | long-term | equilibrium | | | | |
| “Nordhaus” | .02 | .02 | .02 | 7.61 | 2.44 | .45 | 8.4 |
| Markov | .0235 | .01 | .02 | 7.89 | 3.70 | .55 | 16.1 |
| Markov | .0255 | .005 | .02 | 8.05 | 5.79 | .63 | 29.5 |
| Markov | .0271 | .001 | .02 | 8.19 | 19.55 | .73 | 116.4 |
| “Stern” | .001 | .001 | .001 | 8.19 | 19.55 | .95 | 151.8 |

Table 3: Disconnecting carbon pricing from equilibrium discounting. Parameter values as in Table 2

The first row reproduces the Pigouvian case from Table 2 assuming consistent preferences when the annual utility discount rate is set at 2 per cent: this row presents the carbon price under the same assumptions as in Nordhaus (2007), but using our carbon cycle model (which approximates his case very well, as discussed). Keeping the equilibrium time-preference rate at 2, thus maintaining the savings rate at a constant level (reported also in Table 1 of the Introduction), we move to the Markov equilibrium by departing the short- and long-term discount rates in first and second columns. We obtain a radical increase in the carbon price as the long-term discounting decreases, while savings remain unchanged from one set of preferences to the next. The highest carbon tax, 116.4 EUR/tCO₂, corresponds to the case where the long-run discounting is as proposed by Stern (2006); this case also best matches the Weitzman’s values. For reference, we report the Stern case where the long-term discounting holds throughout, the carbon price takes a value of 151.8 EUR/tCO₂, and gross savings cover about 30 per cent of income. Thus, the Markov equilibrium closes considerably the gap between Stern’s and Nordhaus’s carbon prices, without having unrealistic by-products for the macroeconomy.³⁹

³⁸Using the terminology of Barro (1999): Nordhaus is observationally equivalent to our β, δ - equilibrium when the Pigouvian rule is imposed on the equilibrium. That is, the equilibrium outcome is the one that a fictitious planner would choose, assuming that the planner uses the equilibrium utility discount factor γ . We come back to this issue in Section 3.6.

³⁹The deviation between the Markov (thus Nordhaus) and Stern savings can be made extreme by sufficiently increasing the capital share of the output that gives the upper bound for the fraction of y_t saved; close to all income is saved under Stern preferences as this share approaches unity (Weitzman, 2007). However, with reasonable parameters such extreme savings do not occur, as in Table 3.

The empirical evidence is not helpful in gauging the size of intangible damages (parameter Δ_u) associated with a given temperature rise, but using the model structure we can transform the intangible damages into consumption losses and illustrate the damage sensitivity of carbon prices. Note that if output loss caused by damages is x per cent, then the equivalent consumption loss is $x/(1 - g)$ per cent. Above, we had $\Delta_y = .003$ that implied $x = 2.7$ per cent output loss when temperature rises to 3 Kelvin; setting the intangible damage parameter at $\Delta_u = .004$ gives the same loss but now working through the utility only.

Consider then three potential damage scenarios expressed in terms of the percentage output loss, $x = \{1.3, 2.7, 10\}$. That is, in all cases both tangible and intangible losses can contribute to the total loss. For example, in the high-damage case, $x = 10$, setting $\Delta_y = .0056$ and $\Delta_u = .0074$ gives a total loss that is equivalent to 10 per cent output loss. Table 4 shows that the Markov equilibrium tax for low damages are between the median and high damage tax of Nordhaus (i.e, Pigouvian). The results unambiguously show that both the damage estimates and the time preference structure are equally important determinants of carbon prices.

| Carbon Prices | low damages | median damages | high damages |
|--------------------|-------------|----------------|--------------|
| “Nordhaus” | 4.0 | 8.4 | 30.9 |
| Markov equilibrium | 14.2 | 29.5 | 109.4 |
| “Stern” | 73.1 | 151.8 | 562.1 |

Table 4: Carbon prices [Euro/tonCO₂] dependence on structure of time preferences and damage estimates

3.6 Welfare

We have seen that the Pigouvian rule implies efficiency, but not Pareto optimality. We still need to assess whether welfare increases or decreases when Pigouvian carbon prices are implemented. To study this, and to identify policies that improve welfare, we consider how future policies affect current welfare. Note that from the perspective of the current generation, future savings and emission levels are optimal if they are consistent with the long-term time preference δ , that is, if $g = \alpha\delta$ and $h^{\gamma=\delta}$ where h^γ is defined in Lemma 7; then future agents would behave as if they were consistent with present long-term preferences. This thought-experiment gives a clear benchmark against which we can test

how various policy proposals affect current welfare through future policies.

Lemma 9 For $\beta \neq 1$ and $\tau > t$,

$$\begin{aligned} \frac{\partial w_t}{\partial g_\tau} &> 0 \text{ iff } g_\tau < \alpha\delta \\ \frac{\partial w_t}{\partial h_\tau} &> 0 \text{ iff } h_\tau < h^\delta. \end{aligned}$$

Since the equilibrium policies depart from those optimal for the long-run preference δ , any policy that manages to take the decision variables closer to the long-run optimal levels increases current welfare. It turns out that imposing the stand-alone Pigouvian tax principle implies a correction in the wrong direction.

Proposition 7 For slow climate change, implementing Pigouvian carbon prices from period t onwards implies a welfare loss for generation t vis-a-vis the Markov equilibrium.

Proof. By Lemma 4, the change to the Pigouvian carbon price does not affect policy g ; thus, we can focus on the change in current welfare w_t due to the effect of future carbon prices. Let $\beta < 1$ so that $\beta\delta < \gamma < \delta$, and let climate change be a slow process such that $\tau_t^\delta > f_{t,z} > \tau_t^\gamma$; see Proposition 4. Imposing the Pigouvian carbon price will then decrease the future carbon price, taking it further away from τ_t^δ , decreasing current welfare as shown in Lemma 9. The same mechanism applies for $\beta > 1$, when we have $\tau_t^\delta < f_{t,z} < \tau_t^\gamma$. Add the finding that imposing Pigouvian carbon prices on current policies must reduce welfare compared to the unrestricted Markov policy, and it follows that present welfare decreases through both channels. ■

The remarkable feature of the above proposition is that a Pigouvian carbon price policy strictly decreases welfare, not as a second-order effect, but as a first-order effect; this extends the three-period result of Proposition 3 to infinite horizon. However, Lemma 9 suggests that we can achieve more in the infinite horizon setting. We can look for welfare-improvements through self-enforcing policy rules; that is, there is surplus to be created by increasing the policy choices (g, h) for all generations while maintaining their incentive constraints, coming from the threat to switch back to the Markov equilibrium where no rules regarding future behavior apply. Consider a policy pair $(\widehat{g}, \widehat{h})$ that the current generation would like to propose for all generations, including itself. For clarity, consider the quasi-hyperbolic case $\beta < 1$ for the rest of this section. Note from Lemma 9 that generation t would like to propose for all future generations the decision rules $g^{PO} = \delta\alpha$ and $h^{PO} = h^\delta$; these policies would implement efficiency and Pareto optimality, as

the resulting carbon price is Pigouvian, and the utility discount factor satisfies Lemma 8. But, achieving such a future policy requires that these policies are followed also at t , which is ruled out by the current incentive constraints: consumption and carbon pricing given by (g^{PO}, h^{PO}) does not maximize w_t . But, also from Lemma 9, the current generation is willing to give up part of its consumption, by increasing g and h , beyond their Markov equilibrium level, anticipating that all subsequent decision-makers will follow suit when facing the same decision. The best self-enforcing policy pair supported by the Markov equilibrium is between the Pareto Optimal and the Markov policy:

Proposition 8 *For $\beta < 1$, there exist a policy pair $g^{PO} > \hat{g} > g^*$ and $h^{PO} > \hat{h} > h^*$, such that this policy rule maximizes welfare of the first generation, and no future generation can benefit from switching back to the Markov equilibrium, also when applied for each policy rule separately.*

We have seen that the Markov equilibrium allows us to disconnect utility discounting and carbon pricing. Proposition 8 implies that we can tighten the carbon policy ($\hat{h} > h^*$) while keeping savings at the Markov level ($\hat{g} = g^*$). The proposition justifies, from the welfare point of view, carbon prices that are not only higher than the Pigou tax but also higher than the Markov price, even if for institutional reasons, one cannot adjust the savings. Recall that the discussion following Stern (2006) on the appropriate level of the carbon price is centered around the choice of the social discount factor. To justify the Stern's level for carbon prices, one may invoke ethical arguments or fundamental uncertainties (Weitzman, 2007). Our argument is completely different. When the short and long-term time-preferences differ, the self-enforcing policies can support carbon prices above Markov levels, closing the gap between the Stern and Nordhaus estimates even further. In our quantitative analysis below, we illustrate that the quantitative magnitude of such a policy is significant.⁴⁰

⁴⁰The model can also justify Stern through high equilibrium utility discount factors: self-enforcing policies that target both savings and emissions imply simultaneously higher equilibrium discount factors, $\hat{g}/\alpha > g^*/\alpha$, and lower emissions, $\hat{h} > h^*$. Obviously, this maximizes the welfare potential of the policies considered, but can lead to savings rates close to α , which we calibrated in our numerical model below to 0.306. Such savings may be hard to justify, in particular if the capital share α is large, which is the common critique on Stern. We concur with the practical relevance of this criticism, but note that the endogenous link between equilibrium discounting and welfare shows that the Stern proposal is intelligible as an equilibrium outcome. We come back to the numbers in the quantitative analysis.

4 Quantitative analysis

While we can obtain the current climate policies for the current state of the economy, the evaluation of future policies depends on how we expect that technologies in the economy will develop. Thus, to obtain the future climate-economy adjustment paths we need to further specify technologies, and make additional quantitative choices. In particular, we need a detailed structure for the energy sector.

4.1 Production and energy

Consider a production function as in (33) but further specified to

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

where the overall labor-energy composite $A_t(l_{y,t}, e_t)$ takes a CES form with extremely low elasticity of substitution between labor in the final-good sector $l_{y,t}$ and energy e_t , i.e., we take it as Leontief. Productivities $A_{y,t}$ and $A_{e,t}$ are calibrated and thus exogenous. The Leontief assumption avoids unrealistically deep immediate cuts of emissions; see also Hassler, Krusell and Olovsson (2011).⁴¹ Energy e_t also uses labor: the core allocation problem on which we add detail here is how to allocate a given total labor l_t at time t between final output $l_{y,t}$, fossil-fuel energy, $l_{f,t}$, and non-carbon energy, $l_{n,t}$. Thus the energy and climate policy steers the labor allocation $(l_{y,t}, l_{f,t}, l_{n,t})_{t \geq 0}$ and thereby the quantities of fossil-fuel, $e_{f,t}$, and non-carbon energy, $e_{n,t}$. Both energy sources are intermediates, summing up to the total energy input:

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

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

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

where $A_{f,t}$ and B_t describe productivities. The fuel resource is not a fixed factor and commands no resource rent; by this assumption, our focus is on the “coal phase”, as coined by Golosov et al. (2011), where the fuel resource relevant for long-term climate

⁴¹We have the quantitative analysis also for a Cobb-Douglas form for $A_t(l_{y,t}, e_t)$. These results are available on request. See also the next footnote.

policies is in principle unlimited. In contrast, the non-fossil fuel energy production is land-intensive and subject to diminishing returns and land rents (Fischer and Newell, 2008):

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

where $\varphi > 0$ describes the elasticity of supply from this sector, given the labor cost. To find the competitive equilibrium, we first find the policies (g, h) as outlined in the sections above. Subsequently we determine the labour allocation $(l_{y,t}, l_{f,t}, l_{n,t})$, per period.⁴² We provide the equilibrium conditions for the labor-market allocation in the Appendix “Equilibrium labor market allocation”. We also include the Appendix “Calibration”, separately for climate and economic parameters.⁴³

Our calibration, detailed in the Appendix, progresses as follows. When there is no carbon policy, $h = 0$, the labor market allocation can be solved in closed form; thus, we can invert the model to map from quantities $(l, y, e_f, e_n)_{t \geq 0}$ to productivities $(A_y, A_e, A_f, A_n)_{t \geq 0}$.⁴⁴ We match the business-as-usual (BAU) quantities $(y, e_f, e_n)_{t \geq 0}$ with the A1F1 SRES scenario from the IPCC (2000). 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]. We calibrate to 25 per cent gross savings in all scenarios. Under consistent preferences with 2 per cent annual pure time discount, the capital share, under 25 per cent savings, becomes $\alpha = g/\rho = .306$. For the Markov equilibrium, we take $\beta\delta = .7724$ and $\delta = .9511$, corresponding to 2.55 and .5 per cent annual discounting, respectively. Thus, the scenarios correspond to the Markov carbon price at 29.5 for year 2010 in Tables that we have presented in Section 3.5. These choices preserve $g = .25$. From the BAU calibration we proceed to policy experiments: Markov, Pigouvian, and more advanced policies. We turn next to these results.

⁴²The climate policy is essentially a labor market policy. The low elasticity between the final-good and energy in $A_t(l_{y,t}, e_t)$, combined with the decreasing returns to scale for non-fossil fuel energy, prevents unrealistic short-term reallocations of labor following emissions penalties. Note that, in principle, the Leontief does not rule out growth in the long run: the economy can be scaled up without limits with labor, if $h = 0$. Thus, it is the land-intensive carbon-free production that prevents the scale-up in the presence of carbon policies.

⁴³The supplementary material, available online, contains a file that pulls together all parameters, as well as a program files that can be used for replicating the results.

⁴⁴We express all energy in carbon units; to obtain this, we set $B_t = 1$ and employ three distinct energy productivities (A_e, A_f, A_n) .

4.2 Climate-economy adjustment paths

We use now the full model for four scenarios. The simulations are run through the years 2010-3000; the model is initiated at the period labeled ‘2010’, representing 2006-2015.

We consider the following scenarios. First, in the business as usual (BAU) scenario, we assume no climate policy and set $h = 0$, while savings g are at the Markov equilibrium level. Second, in the Markov equilibrium generations choose their equilibrium savings and climate policies at each income level as we have described in our theory (Section 3.3). Third, in the Pigouvian equilibrium, we impose the Pigouvian tax as detailed in Section 3.4. For interpretation, it is useful to notice that because, under our assumptions, savings and carbon policies are separable, moving between scenarios one to three does not change the savings rate; only the carbon policy is affected. Moreover, as we calibrate the model to match savings under consistent preferences with 2 per cent annual discounting, the Pigouvian outcome is observationally equivalent to the Nordhaus DICE 2007 cost-benefit scenario. Finally, we consider an “Advanced policy” scenario where the carbon prices are chosen to maximize each generation’s welfare in a self-enforcing manner: as in Proposition 8 we find the highest carbon tax, characterized by constant \hat{h} , such that each generation has incentives to follow this policy rather than switching the continuation path to the Markov equilibrium. We consider the advanced policy only for the carbon tax and not for the savings.⁴⁵

Figure 2 shows the carbon prices for the three policy scenarios. The Pigouvian tax rule has the lowest carbon prices, from 8.3 EUR/tCO₂ in 2010, increasing with output to 43 EUR/tCO₂ in 2100.⁴⁶ The Markov equilibrium implements a considerably higher carbon price, rising from 29.4 to 106 EUR/tCO₂ during the same time period. Finally, the advanced policy further increases the carbon price to 34 EUR/tCO₂ in 2010, arriving at 106 EUR/tCO₂ in 2100. Between the scenarios, welfare is, of course, the lowest for all generations in the BAU scenario and increases as we move to the Pigouvian scenario, and then to the Markov equilibrium, and, finally, the highest welfare is reached by all generations in the Advanced Policy scenario.

Figure 2 here: Carbon prices over 2010-2100

Figure 3 shows the emissions associated with these scenarios.⁴⁷ The key take-away

⁴⁵There are two reasons for this focus. First, describing the self-enforcing policies in the two dimensions is computationally challenging. Second, such policies imply high savings and thus violate our basic premise that the shorter-term decisions should be consistent with actual historical behavior.

⁴⁶Multiply the values by factor 4.7 to arrive at the USD/tC values.

⁴⁷We calculate carbon prices for 2010, but assume that they will affect emissions from the second

from this figure is that Pigouvian taxes cannot prevent emissions from increasing substantially beyond current levels, for the full two coming centuries. In contrast, the carbon prices in the Markov equilibrium trigger a complete transition towards carbon-free energy sources by the end of 2100, and the Advanced Policy has the most drastic implications for energy supply and use. The increasing sequence of carbon prices translates into lower cumulative emissions over the period 2006-2105. They range from 6600 GtCO₂ in BAU, via 5300 GtCO₂ for the Pigouvian scenario, and 1770 GtCO₂ for the Markov scenario, to 1410 GtCO₂ for the Advanced Policy scenario. We note that while the Pigou scenario matches the Nordhaus DICE 2007 emissions well, for the Markov and Advanced scenarios the model predicts a change in the energy system that is unrealistically fast, given that the existing capital structure pre-determines energy infrastructure options in the near future.⁴⁸

Figure 3 here: CO2 emissions 2000-2200

Figure 4 shows the temperature rise associated with these scenarios. Due to the extreme persistence of climate, we show the temperatures up to the year 3000. There is large uncertainty about the climate sensitivity parameter π .⁴⁹ We use the best scientific estimates for π , implying that taking no policy action can lead to temperatures rising above 10 Kelvin, changing earth into an environment that hardly resembles anything humans have encountered in history. Applying the Pigouvian carbon price rule curbs the temperature rise to a maximum temperature rise of 3.7 K, reached around 2220. Even after thousand years, global temperatures are hardly below 3K.⁵⁰ The Markov equilibrium where emissions over the coming century are cut more than 70 per cent vis-a-vis the BAU scenario almost keeps the temperature rise below 2 K. Finally the Advanced Policy scenario keeps the temperature slightly lower, but even in this drastic scenario, for the coming thousand years, temperatures will remain above levels seen for the past 400,000 years.

* * *Figure 4 here: Temperature rise 2000-3000 * * *

Figure 5 shows the effect of the policies on consumption per capita. The Pigouvian period onwards only, that is, from the period labeled ‘2020’.

⁴⁸Because of the added complexity, only a few integrated assessment models describe a vintage-based capital structure (e.g., van der Zwaan et al. 2002).

⁴⁹See Kelly and Kolstad (1999) for seminal work on modeling the learning of climate sensitivity.

⁵⁰Such a temperature path could already lead to the collapse of the Greenland ice sheet. See Krieglera et al. (2009), and Alley et al. (2010) , and the following quote in the latter: “The evidence suggests nearly total ice-sheet loss may result from warming of [...] perhaps as little as 2°C or more than 7°C.”

carbon prices decrease consumption at about 0.1 percentage, relative to the Business as Usual for the coming century, and prevents a downturn at the second half of the millennium when temperature rises would cause severe damage. The Markov policy decreases consumption at about one percentage in addition to the Pigouvian tax scenario, for the coming century. In return, consumption increases from 2150 onwards, by up to 3 per cent between 2200 and 2300. The advanced policy takes it one step further, but the difference with the Markov equilibrium is small.

We emphasize that the results describing BAU and the Pigouvian scenarios are consistent with the experiences from numerous integrated assessment models for the medium-term future.⁵¹ However, not all integrated climate-economy assessment models feature the same persistence of climate change, implying that the far-distant impacts are relatively pronounced in our model. This seems justified because the main consequences of climate policies of the 21st century will be felt beyond 2200, and because our parameters are based on rigorous scientific insights.⁵²

Figure 5 here: Consumption per capita 2000-3000

5 Discussion

In an editorial comment Nordhaus (1997) discusses a benchmark climate change policy, comparable to our Pigouvian tax rule, and its various alternatives giving more weight to the concerns for climate change. The first alternative applies a lower than market discount rate uniformly to both capital and climate investments. The second alternative uses differential discounting: a market discount rate is used for capital, while a lower discount rate is used for climate investments; the approach is based on Hasselmann et al. (1997). The third approach sets a longer-term climate target, based on the variables believed to matter most for ultimate climate change damages. The target could be a limit on the global temperature rise by the year 2500. The fourth approach sets intermediate climate change targets, e.g., for emissions or atmospheric CO_2 concentrations. In his assessment of the Pigouvian tax rule and the alternative policies, Nordhaus finds that the Pigouvian tax rule outperforms all alternatives. That is, all alternatives result in

⁵¹Running the baseline DICE 2007 using our .02 discount rate throughout, leads to a climate-economy outcome that is very close our Pigouvian scenario during the coming two centuries, not only with respect to the initial tax but also to the time paths described in this section.

⁵²It is not possible to calculate the same scenarios in other integrated assessment models, as they are not well suited to the analysis of Markov type of equilibria.

strict welfare losses compared to the Pigouvian approach. The long-term climate target has the lowest welfare losses given the gains from climate stabilization; using intermediate climate variables as targets comes next; differential discounting comes third; and using lower discount rates for both capital and climate change policy is the worst policy.

While our results are certainly supportive of setting a price on emissions, our assessment of the policies is markedly different. Our Proposition 8 establishes that when intergenerational preferences are not time-consistent, the welfare potential is maximized by self-enforcing commitment to a lower than shorter-term market discount rates for both capital investments and climate change policy. If self-enforcing policies are deemed unrealistic, the Markov equilibrium is a simple alternative. In contrast with the Pigouvian approach, it sets differential discount rates for climate and private capital; the Markov policies perform hugely better than the Pigouvian policies applying a uniform (market-return) discounting for all assets.⁵³

To illustrate the gains from differential discounting (in the Markov equilibrium) when compared to the uniform discounting at the level of capital returns (in the Pigouvian equilibrium), we construct a “wrinkle experiment”, following Nordhaus (2007). In his criticism on Stern, Nordhaus (2007) proposes the following thought experiment. Consider a project that can yield a small but perpetual gain in a distant future, such as 0.1 per cent gain in income 200 years from now onwards. How much the current decision-maker is willing to give up to achieve such a future gain? If we discount future gains at 0.1 per cent per year, the net present value of the project amounts to 80 per cent of one year of income (Nordhaus 2007). We certainly do not want to give up so much for such a small future gain, even if perpetual. Our wrinkle experiment compares the Pigou policy with the Markov equilibrium scenario: viewing such a comparison as a project, it brings the peak global average warming from 3.7 K to about 2 K, and a large part of this gap (1.3K) is almost perpetual. The question we must ask is whether the current generation is willing to give up 1-2 per cent of its income to implement that project. Under our parameter choices, the answer is affirmative to this project decision. The associated “consumption wrinkle” is depicted in Figure 6. A subsequent project, in the same Figure, compares the Advanced Policy with the Markov scenario, which would shave off another 0.1K from global warming, from 2200 onwards, perpetually. The costs of this challenging policy

⁵³The climate target policies suffer from another type of problem: they are infeasible, as envisioned by Nordhaus, without a supporting intertemporal global commitment device. The target is arbitrary, and once we realize that, why would the next generation set the same target as the current generation? But the moment you can change the target over time, the concept becomes internally inconsistent.

would be an additional 0.1 per cent of our income for the coming century.⁵⁴ Under our parameter choices, the answer is again affirmative to this far-reaching project as well.

The parameters may err on both sides. We may have considered too low discount rates for the long-term structure of time preferences. On the other hand, we may underestimate the damages of climate change severely; a constant damage parameter Δ_y over centuries is a strong assumption. Despite the scope for errors, the conclusions from our wrinkle experiment seem robust: the experiment favors the stricter climate policies under the Markov equilibrium over the Pigouvian carbon price with a large margin.

Figure 6 here: Changes in consumption between scenarios 2000-3000

The most important consequences of climate change might not be related to the loss of income as measured through industrial activities and services (market impacts influencing prices), but future generations may measure their loss in terms of non-market impacts including losses of large ecosystems such as rain-forests and coral reefs that can collapse when they can not adapt quickly enough to the rapidly changing global environment. Though global warming receives most attention, ocean acidification is increasingly considered a complementary consequence of anthropogenic CO_2 emissions, and a major threat to marine biodiversity (Orr et al. 2005). As market-based effects are easier to evaluate, the set of available estimates tends to be biased towards such impacts.

Our model shows that, to carry out relevant policy assessments, we are in need of three sources of information. First, there should be a more structured approach to estimating climate impacts, including non-market damages. Second, we need to better understand the carbon-climate cycle. Third, given the assessment of possible consequences, we have to assess the willingness of people to give up part of their income to prevent part of these consequences from happening, respecting the willingness to pay as a decision separate from savings.

6 Concluding remarks

September 2011, the U.S. Environmental Protection Agency (EPA) sponsored a workshop to seek advice on how the benefits and costs of regulations should be discounted for

⁵⁴To calculate the costs of the climate policy on its own merits, we constructed an auxiliary scenario where we fixed the savings policy to the BAU levels, and determined the optimal carbon tax policy h , that would be self-enforcing.

projects with long horizons; that is, for projects that affect future generations. The EPA invited 12 academic economists to address the following overall question: “What principles should be used to determine the rates at which to discount the costs and benefits of regulatory programs when costs and benefits extend over very long horizons?” In the background document, the EPA prepared the panelists for the question as follows: “Social discounting” in the context of policies with very long time horizons involving multiple generations, such as those addressing climate change, is complicated by at least three factors: (1) the “investment horizon” is significantly longer than what is reflected in observed interest rates that are used to guide private discounting decisions; (2) future generations without a voice in the current policy process are affected; and (3) compared to shorter time horizons, intergenerational investments involve greater uncertainty. Understanding these issues and developing methodologies to address them is of great importance given the potentially large impact they have on estimates of the total benefits of policies that impact multiple generations.”

In this paper, we have developed a methodology for addressing the over-arching question posed above, and also provided a quantitative evaluation of optimal carbon prices using the methods developed. The analysis builds on two premises that seem consistent with practical program evaluation principles: (i) shorter-term decisions should respect the revealed attitudes towards choices over time; and (ii) the far-distant future should be treated differently and discounted with a lower rate. The climate-change problem stipulates a general-equilibrium approach where the returns on assets become dependent on the policies adopted. In this setting, distortions cannot be avoided: even when climate externalities are internalized, i.e., when policy makers fully take into account the inflicted impacts on future generations, the consumption trade-offs over time of the future decision makers are incongruent with those of the current generation. These distortions imply that the capital returns do not reflect how the future damages from climate impacts should be discounted; when time discounting declines with time, the current decision makers value investments in the climate asset more than indicated by the capital returns. As a result, the current optimal price on carbon emissions exceeds the Pigouvian level. When welfare is measured such that the original preferences of the generations are respected (no dictatorial objectives), imposing the Pigouvian rule decreases welfare for all generations. We have quantitatively demonstrated the significance of the difference between equilibrium and Pigouvian policies in all quantitative dimensions of the climate problem using state of the art approaches to the economy as well as to the climate system.

References

- [1] Alley, Andrews, Brigham-Grette, et al. (2010), “History of the Greenland Ice Sheet: paleoclimatic insights”, in *Quaternary Science Reviews* 29: 1728-1756.
- [2] Barro R.J. (1999), Ramsey meets Laibson in the neoclassical growth model, *Quarterly Journal of Economics* 114: 1125-1152.
- [3] Bernheim, D.B., and A. Rangel (2009), *Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics*, The Quarterly Journal of Economics, MIT Press, vol. 124(1), pages 51-104.
- [4] Bernheim, D.B., and D. Ray (1987), Economic growth with intergenerational altruism, *The Review of Economic Studies*, Vol. 54, No. 2: 227-241.
- [5] Boden, T.A., G. Marland, and R.J. Andres. 2011. Global, Regional, and National Fossil-Fuel CO₂ Emissions. Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, U.S. Department of Energy, Oak Ridge, Tenn., U.S.A.
- [6] Caldeira K. and M. Akai (eds) 2005, Ocean storage, Ch 6 in IPCC special report on carbon dioxide capture and storage, edited by Metz B., O. Davidson, H. de Coninck, M. Loos, and L. Meyer, Cambridge University Press.
- [7] Caplin A, and Leahy J (2004), The social discount rate, *Journal of Political Economy* 112: 1257-1268.
- [8] Dasgupta, P., Discounting Climate Change, *Journal of Risk and Uncertainty*, 2008, 37(2-3), 141-169.
- [9] Fischer C. and R. Newell, 2008, Environmental and technology policies for climate mitigation, *J of Environmental and Economic Management* 55:142-162.
- [10] Gerlagh, R., and M. Liski, Public investment as commitment, CESifo Working Paper Series No. 3330, available at hse-econ.fi/liski/papers/pubinv.pdf
- [11] Gollier, C., and R. Zeckhauser (2005), Aggregation of Heterogeneous Time Preferences, *Journal of Political Economy*, 2005, vol. 113, issue 4, ppp. 878-896.
- [12] Golosov, M., J. Hassler, P. Krusell, A. Tsyvinski (2011), Optimal taxes on fossil fuel in general equilibrium, NBER WP 17348.

- [13] Goulder, L.H, and R.C. Williams III (2012), The choice of discount rate for climate policy change policy evaluation, NBER Working Paper No. 18301.
- [14] 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
- [15] Hassler, J., P. Krusell, and C. Olovsson (2011), Energy-saving technical change, mimeo.
- [16] 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
- [17] 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.
- [18] 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.
- [19] Jackson, O. M., and L. Yariv, Collective Dynamic Choice: The Necessity of Time Inconsistency, WP-2011
- [20] Jevrejeva, S., J.C. Moore, A. Grinsted (2011), Sea level projections to AD2500 with a new generation of climate change scenarios, *Global and Planetary Change* 80-81:14-20.
- [21] Kaplow L., E. Moyer, and D. A. Weisbach (2010) òThe Social Evaluation of Intergenerational Policies and Its Application to Integrated Assessment Models of Climate Change,ó The B.E. Journal of Economic Analysis and Policy: Vol. 10: Iss. 2 (Symposium), Article 7.
- [22] Karp L. (2005), Global warming and hyperbolic discounting, *Journal of Public Economics* 89: 261-282.
- [23] Karp L. (2007), Non-constant discounting in continuous time, *Journal of Economic Theory* 132: 557-568.

- [24] Kelly, D.L., and C.D. Kolstad (1999), Bayesian learning, growth, and pollution, *Journal of economic dynamics and control* 23, 491-518.
- [25] Krieglera, Hall, Helda, Dawson, and Schellnhuber (2009), “Imprecise probability assessment of tipping points in the climate system”, *Proceedings of the National Academy of Sciences*, vol 106: 5041-5046.
- [26] Krusell P., Kuruscu B, Smith A.A. (2002), Equilibrium welfare and government policy with quasi-geometric discounting, *Journal of Economic Theory* 105: 42-72.
- [27] Krusell P. and A.A. Smith (2003), Consumption-savings decisions with quasi-geometric discounting, *Econometrica* 71: 365-375.
- [28] Laibson D. (1997), Golden eggs and hyperbolic discounting, *Quarterly Journal of Economics* 112: 443-477.
- [29] 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.
- [30] Manne A.S., R.Mendelsohn, and R.Richels (1995), MERGE, A model for evaluating regional and global effects of GHG reduction policies, *Energy Policy* 23, 17-34.
- [31] Nordhaus, W. D.(1993), Optimal greenhouse-gas reductions and tax policy in the “DICE” model, *The American Economic Review*, 83 (2), pp. 313-317.
- [32] Nordhaus, W.D. (1997), Discounting in economics and climate change, an editorial, in *Climatic Change* 37: 315-328.
- [33] Nordhaus, W. D.(2007), A Review of The Stern Review on the Economics of Climate Change, *Journal of Economic Literature*, 45 (3), 686-702.
- [34] Maskin, E., and J. Tirole. 2001. “Markov Perfect Equilibrium I: Observable Actions”, *Journal of Economic Theory* 100, 191-219.
- [35] Muller, N.Z, R. Mendelsohn, and W. Nordhaus (2011), Environmental accounting for pollution in the United States economy, *American Economic Review*, 101 (5): 1649-75

- [36] Orr, J.C., V.J. Fabry, O. Aumont, et al. (2005), Anthropogenic ocean acidification over the twenty-first century and its impacts on calcifying organisms, *Nature* 437: 681-686.
- [37] Peck S.C. and T.J. Teisberg (1992), CETA: a model for carbon emissions trajectory assessment, *Energy Journal* 13, 55-77.
- [38] Phelps E.S. and R.A. Pollak (1968), On second-best national saving and game-equilibrium growth, *Review of Economic Studies* 35(2): 185-199.
- [39] Rubinstein, A. (2003), Economics and psychology? The case of hyperbolic discounting, *International Economic Review* 44: 1207-1216.
- [40] Saez-Marti M. and J.W. Weibull (2005), Discounting and altruism to future decision makers, *Journal of Economic Theory* 122: 254-266
- [41] 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.
- [42] Stern, N. (2006), "The economics of climate change: the Stern review", Cambridge, UK: Cambridge University Press.
- [43] Tol, R. (2009), The economic effects of climate change, *Journal of economic perspectives*, 23(2): 29-51.
- [44] Weitzman, M. (2001). Gamma Discounting, *American Economic Review*, Vol. 91, No. 1, pp. 260-271.
- [45] Weitzman, M. (2007). A Review of The Stern Review on the Economics of Climate Change, *Journal of Economic Literature*, 45 (3), 703-724.

Appendix

Lemma 4

The proof is by induction. Assume that (41) and (42) hold for all future periods $t+1, t+2, \dots$, and that the lemma holds for $t+2$. We can thus construct the value function for the next period, as

$$W_{t+1}(k_{t+1}, s_{t+1}) = u_{t+1} + \delta W_{t+2}(k_{t+2}, s_{t+2}).$$

Substitution of the investment decision at time $t + 1$, $k_{t+2} = gy_{t+1}$ and emissions $z_{t+1} = z_{t+1}^*$, gives

$$W_{t+1}(k_{t+1}, s_{t+1}) = [\ln(1 - g_{t+1}) + \ln(A_{t+1}) + \alpha \ln(k_{t+1}) + \ln(\omega(s_{t+1}))] - \Delta_u D_{t+1} \\ + \delta \tilde{A}_{t+2} + \delta \xi [\ln(g_{t+1}) + \ln(A_{t+1}) + \alpha \ln(k_{t+1}) + \ln(\omega(s_{t+1}))] + \delta \Omega(s_{t+2})$$

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+1} and s_{t+1} , we get the constant part of $V_{t+1}(k_{t+1})$:

$$\tilde{A}_{t+1} = \ln(1 - g_{t+1}) + \delta \xi \ln(g_{t+1}) + (1 + \delta \xi) \ln(A_{t+1}) - \delta \zeta_1 z_{t+1} + \delta \tilde{A}_{t+2}. \quad (50)$$

Collecting the coefficients in front of $\ln(k_{t+1})$ yields the part of $V_{t+1}(k_{t+1})$ depending k_{t+1} 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+1} yields $\Omega(s_{t+1})$ through

$$\Omega(s_{t+1}) = \ln(\omega(s_{t+1}))(1 + \delta \xi) - \Delta_u D_{t+1} + \delta \Omega(s_{t+2}).$$

where $z_{t+1} = z_{t+1}^*$ appearing in $s_{t+2} = (z_1, \dots, z_t, z_{t+1})$ is independent of k_{t+1} and s_{t+1} (by Lemma 6 that holds by the induction hypothesis) so that we only need to consider the values for z_1, \dots, z_t when evaluating $\Omega(s_{t+1})$. The values for ζ_τ can be calculated by collecting the terms in which $z_{t+1-\tau}$ appear. Recall that $\ln(\omega(s_{t+1})) = -\Delta_y D_{t+1}$ so that

$$\zeta_\tau = ((1 + \delta \xi) \Delta_y + \Delta_u) \sum_{(i,j)} a_i b_j \pi \varepsilon_j \frac{(1 - \eta_i)^\tau - (1 - \varepsilon_j)^\tau}{\varepsilon_j - \eta_i} + \delta \zeta_{\tau+1}$$

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

$$\zeta_\tau = \left(\frac{\Delta_y}{1 - \alpha \delta} + \Delta_u \right) \sum_{(i,j)} \sum_{t=\tau}^{\infty} a_i b_j \pi \varepsilon_j \delta^{t-\tau} \frac{(1 - \eta_i)^t - (1 - \varepsilon_j)^t}{\varepsilon_j - \eta_i}$$

To derive the value of ζ_1 , we consider

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

Finally, we notice that a careful examination shows that the final equation still holds when $\eta_i = \varepsilon_i$, even though we then cannot follow the same derivation. Q.E.D.

Lemma 6

The first-order conditions for fossil-fuel use z_t , and the labor allocations over the final goods $l_{y,t}$ and the energy sectors $l_{e,t}$ give:

$$u'_t \frac{\partial y_t}{\partial z_t} = \beta \delta \frac{\partial \Omega_{t+1}}{\partial s_{t+1}} \frac{\partial s_{t+1}}{\partial z_t} \Rightarrow \frac{1}{1-g} \frac{1}{A_t} \frac{\partial A_t}{\partial e_t} \frac{\partial E_t}{\partial z_t} = \beta \delta \zeta_1 \quad (51)$$

$$\frac{\partial A_t}{\partial l_{y,t}} = \frac{\partial A_t}{\partial e_t} \frac{\partial E_t}{\partial l_{e,t}} \quad (52)$$

The second part of the Lemma follows immediately from (51):

$$\frac{\partial y_t}{\partial z_t} = f_{t,z} = \beta \delta \zeta_1 (1-g) y_t.$$

The second equation equates the marginal product of labor in the final good sector with the indirect marginal product of labor in energy production. We have thus four equations, energy production (34), labour market clearance (35), and the above two first-order conditions, that jointly determine four variables: $z_t, l_{y,t}, l_{e,t}, e_t$, only dependent on technology at time t through $A_t(l_{y,t}, e_t)$ and $E_t(z_t, l_{e,t})$. Thus, $z_t = z_t^*$ can be determined independent of the state variables k_t and s_t . Q.E.D.

Lemma 7

To determine the Pigouvian tax as the net present value of marginal damages, note that from the emissions-damage response function we have

$$\begin{aligned} \frac{du_{t+\tau}}{dz_t} &= \left(\frac{\Delta_y}{1-g} + \Delta_u \right) \frac{dD_{t+\tau}}{dz_t} \\ &= \left(\frac{\Delta_y}{1-g} + \Delta_u \right) \sum_i a_i b_j \pi \varepsilon_j \frac{(1-\eta_i)^\tau - (1-\varepsilon_j)^\tau}{\varepsilon_j - \eta_i} \end{aligned}$$

Furthermore, we note that the marginal rate of substitution for utility, between two periods, in equilibrium, is γ , so that the net present value of future damages associated with one extra unit of emissions, in current utility terms, h^{Pig} , is given by

$$\begin{aligned} h^{Pig} &= \sum_{\tau=1}^{\infty} \gamma^\tau \frac{du_{t+\tau}}{dz_t} \\ &= \left(\frac{\Delta_y}{1-g} + \Delta_u \right) \sum_i \frac{a_i b_j \pi \varepsilon_j}{\varepsilon_j - \eta_i} \sum_{\tau=1}^{\infty} \gamma^\tau (1-\eta_i)^\tau - \gamma^\tau (1-\varepsilon_j)^\tau \\ &= \left(\frac{\Delta_y}{1-g} + \Delta_u \right) \sum_i a_i b_j \pi \varepsilon_j \gamma \sum_{\tau=0}^{\infty} \gamma^\tau (1-\eta_i)^\tau - \gamma^\tau (1-\varepsilon_j)^\tau \\ &= \left(\frac{\Delta_y}{1-g} + \Delta_u \right) \sum_i \frac{\gamma \pi a_i b_j \varepsilon_j}{[1-\gamma(1-\eta_i)][1-\gamma(1-\varepsilon_j)]} \\ &= h^\gamma \end{aligned}$$

Q.E.D.

Lemma 8

This proof follows Gerlagh and Liski (2011). Pareto optimal program is a feasible allocation $\{c_t, z_t, k_t\}_{t=1}^{\infty}$ maximizing objective

$$W(\cdot) = \sum_{t=1}^{\infty} \alpha_t w_t$$

for some non-negative welfare weights α_t having a bounded mass $\sum_{t=1}^{\infty} \alpha_t < \infty$. The lemma considers allocations represented through geometric utility weights $\alpha'_t = \gamma^{t-1}$. If $\gamma < \beta\delta$ or $\gamma < \delta$, it is not possible to construct a sequence of non-negative welfare weights α_t consistent with a Planner maximizing under the sequence of utility weights α'_t . Suppose the contrary, that welfare weights $\alpha_t \geq 0$ consistent with α'_t exist. Then, using the definition of welfare, we see that for some $\tau > t$, the relationship between the two is $\alpha'_1 = \alpha_1$ and $\alpha'_\tau = \alpha_\tau + \sum_{t=1}^{\tau-1} \alpha_t \beta \delta^{\tau-t}$. Expanding the latter gives

$$\alpha'_\tau = \alpha_1 \beta \delta^{\tau-1} + \alpha_2 \beta \delta^{\tau-2} + \dots + \alpha_{\tau-1} \beta \delta + \alpha_\tau. \quad (53)$$

If $\gamma < \delta$ and $\alpha_1 > 0$, we see that the equation cannot hold with $\alpha_\tau \geq 0$ for sufficiently large τ : $\alpha'_\tau - \alpha_1 \beta \delta^{\tau-1} < 0$ for some finite $\tau > 0$.

If $\gamma < \beta\delta$, include only the last two terms on the right in (53) to obtain

$$\alpha'_{\tau+1} \geq \beta\delta\alpha_\tau + \alpha_{\tau+1}.$$

Substitute $\alpha'_{\tau+1} = \gamma\alpha'_\tau$,

$$\gamma\alpha'_\tau - \beta\delta\alpha_\tau \geq \alpha_{\tau+1}.$$

Since $\gamma < \beta\delta$, this cannot hold with $\alpha_{\tau+1} \geq 0$ for any τ .

Consider now $\gamma \geq \max\{\beta\delta, \delta\}$. We show that now one can construct the non-negative welfare weights. We construct an algorithm for finding the weights. Let $\tilde{\alpha}_1 = \{\alpha_\tau^1\}_{\tau \geq 1}$, $\tilde{\alpha}_2 = \{\alpha_\tau^2\}_{\tau \geq 2}$, and so on. Define

$$\begin{aligned} \alpha_\tau^1 &= \gamma^{\tau-1}, \tau \geq 1 \\ \alpha_\tau^2 &= \alpha_\tau^1 - \alpha_1^1 \delta^{\tau-2}, \tau \geq 2 \\ &\dots \\ \alpha_\tau^{t+1} &= \alpha_\tau^t - \alpha_t^t \delta^{\tau-t-1}, \tau \geq t. \end{aligned}$$

The value of α_τ^t measures the weight remaining for generation τ after all altruistic weights of generations 1 to $t-1$ have been subtracted. Note that the equilibrium implies utility

weights $\tilde{\alpha}_1$, and $\{\alpha_t^t\}_{t \geq 1}$ is the sequence of welfare weights consistent with $\tilde{\alpha}_1$. The main intermediate result that we need, in order to prove that the sequence of welfare weights $\{\alpha_t^t\}_{t \geq 1}$ is non-negative, is that for all $\tau \geq t$:

$$\frac{\alpha_{\tau+1}^t}{\alpha_\tau^t} > \max\{\beta\delta, \delta\}. \quad (54)$$

By construction, this condition is satisfied for $t = 1$. It implies that next sequence $\tilde{\alpha}_2$, induced by the algorithm, is non-negative, as

$$\alpha_\tau^2 = \gamma^{\tau-1} - \delta^{\tau-t-1} > \alpha_\tau^1 \{(\max\{\beta\delta, \delta\})^{\tau-1} - \beta\delta^{\tau-t}\} > 0, \tau \geq 2.$$

By induction, if the condition holds for $\tilde{\alpha}_t$, the sequence $\tilde{\alpha}_{t+1}$ is non-negative:

$$\alpha_\tau^{t+1} > \alpha_\tau^t \{(\max\{\beta\delta, \delta\})^{\tau-t} - \beta\delta^{\tau-t}\} > 0, \tau \geq t.$$

Thus, we are done if we can show that condition (54) holds. Notice that

$$\alpha_{\tau+1}^{t+1} = \alpha_{\tau+1}^t - \alpha_t^t \beta \delta^{\tau-t+1} > \max\{\beta\delta, \delta\} \alpha_\tau^t - \alpha_t^t \beta \delta^{\tau-t+1} \geq \delta \{\alpha_\tau^t - \alpha_t^t \beta \delta^{\tau-t}\} = \delta \{\alpha_\tau^{t+1}\}.$$

If $\beta < 1$, this proves that $\alpha_{\tau+1}^{t+1} > \delta \{\alpha_\tau^{t+1}\} > \beta \delta \{\alpha_\tau^{t+1}\}$. On the other hand, if $\beta > 1$, we have

$$\alpha_{\tau+1}^{t+1} = \alpha_{\tau+1}^t - \alpha_t^t \beta \delta^{\tau-t+1} > \max\{\beta\delta, \delta\} \alpha_\tau^t - \alpha_t^t \beta \delta^{\tau-t+1} \geq \beta \delta \{\alpha_\tau^t - \alpha_t^t \beta \delta^{\tau-t}\} = \beta \delta \{\alpha_\tau^{t+1}\},$$

which completes the proof. Q.E.D.

Lemma 9

Consider a given policy path $(g_\tau, z_\tau)_{\tau \geq t}$. We then look at variations in policies at time τ , consider the effect on welfare at time t . All effects are captured by W_{t+1} that we analyze in Lemma 4. The analysis in the proof of Lemma 4 implies: the value function at time t is separable in states and the parameters ξ and ζ do not depend on future polices (g_τ, z_τ) , but term \tilde{A}_t does. Technically, we need to show that, for some $\tau > t$,

1. \tilde{A}_t increases in g_τ for $g_\tau < \alpha\delta$,
2. \tilde{A}_t decreases in z_τ for $z_\tau > z_\tau^\delta$,

where z_τ^δ is the emission level that is consistent with the policy variable h^δ and z_τ is the emission level consistent with some $h < h^\delta$. From the proof of Lemma 4, consider

(50). Term \tilde{A}_t increases with \tilde{A}_τ for some $\tau > t$. Moreover, \tilde{A}_τ is strictly concave in g_τ , and maximal when g_τ maximizes $\ln(1 - g_\tau) + \delta\xi \ln(g_\tau)$, that is, for $g_\tau = \frac{\delta\xi}{1+\delta\xi} = \alpha\delta$. This proves item 1. Furthermore, notice that A_τ depends on z_τ , that \tilde{A}_τ is strictly concave in z_τ and maximal when z_τ maximizes $(1 + \delta\xi) \ln(A_\tau(z_\tau)) - \delta\zeta_1 z_\tau$, that is, for $\frac{d \ln A_t}{A_\tau dz_\tau} = \delta\zeta_1(1 - \alpha\delta)$. This is the value of z_τ consistent with h^δ . This proves item 2. We now shown the “if” part of Lemma 9. The “only if” follows from the strict concavity of \tilde{A}_τ with respect to (g_τ, z_τ) . Q.E.D.

Proposition 4

We consider the ratio between the carbon price and the Pigouvian carbon price for very long climate change delays, $\eta_i = \varepsilon_j = 0$, and quasi-hyperbolic preferences, $\beta < 1$:

$$\begin{aligned} \frac{f_{t,z}}{\tau_t^\gamma} &= \frac{(1 - \gamma)^2 \beta \delta}{(1 - \delta)^2 \gamma} \\ &= \frac{\left(1 - \frac{\beta \delta}{1 - \alpha \delta + \alpha \beta \delta}\right)^2}{(1 - \delta)^2} (1 - \alpha \delta + \alpha \beta \delta) \\ &= \frac{(1 - \delta(\alpha + (1 - \alpha)\beta))^2}{(1 - \delta)^2(1 + \alpha\delta(\beta - 1))} \\ &> 1 \end{aligned}$$

The first equality follows from substitution of $\eta_i = \varepsilon_j = 0$ in the equation for the equilibrium carbon price and Pigouvian carbon price. The second equality substitutes the value for γ . The final inequality follows as for $\beta < 1$, we have that $\alpha + (1 - \alpha)\beta < 1$, and thus the numerator exceeds $1 - \delta$, while $\beta < 1$ also ensures that the second term in the denominator falls short of 1. Q.E.D.

Equilibrium labor market allocation in Section 4

This Appendix details the labor market allocation for the functional forms introduced in Section 4; this allocation is then numerically solved to obtain the overall climate-economy adjustment path in Section 4.2. As explained in the text, the allocation can be solved period-by-period taking the (i) productivity parameters, (ii) total labor, (iii) savings g , and (iv) carbon policies $h = \beta\delta\zeta_1$ as given. Items (i)-(ii) change over time, implying reallocations of labor. But these reallocations satisfy equations (55)-(58) below. We drop the time subscript in the variables.

1. Consumers choose the fraction g of output y that is available for consumption. We normalize prices for the final good to equalize 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 the prices of energy q and output p as given. Since $y = k^\alpha [\min\{A_y l_y, A_e e\}]^{1-\alpha} \omega(s)$, factor compensation for labour and energy together receives a share $(1-\alpha)$ of the value of output py :

$$w l_y + q e = (1-\alpha) p y$$

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 $\beta \delta \zeta_1$

$$\left[q - \left(\frac{w}{A_f} + \frac{\rho \zeta_1}{B} \right) \right] \times l_f \leq 0.$$

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

$$q e_f = w l_f + \rho \zeta_1 z = \left(w + \frac{\rho \zeta_1 A_f}{B} \right) l_f$$

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

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

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:

$$q e_n = \left(1 + \frac{1}{\varphi} \right) w l_n$$

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

$$A_y l_y = A_e \left(A_f l_f + \frac{\varphi+1}{\varphi} (A_n l_n)^{\frac{\varphi}{\varphi+1}} \right) \quad (55)$$

$$w l + \frac{\rho \zeta_1 A_f}{B} l_f + \frac{1}{\varphi} w l_n = \frac{1-\alpha}{1-g} \quad (56)$$

$$\frac{w}{A_f} + \frac{\rho \zeta_1}{B} \geq \frac{w}{(A_n)^{\frac{\varphi}{\varphi+1}}} (l_n)^{\frac{1}{\varphi+1}} \perp l_f \geq 0 \quad (57)$$

$$l_y + l_f + l_n = l \quad (58)$$

Equation (55) follows since, for strictly positive input prices, $A_t(\cdot) = \min \{A_y l_y, A_e e\} \Rightarrow A_y l_y = A_e e$. Equation (56) allocates 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 (57) 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, $\zeta_1 = 0$, we can solve the allocation in closed-form:

$$l_{n,t} = \frac{A_{n,t}^\varphi}{A_{f,t}^{\varphi+1}} \quad (59)$$

$$w_t = \frac{1 - \alpha}{1 - g} \frac{\varphi}{\varphi l_t + l_{n,t}} \quad (60)$$

$$l_{y,t} = \frac{A_{e,t}}{A_{y,t} + A_{e,t} A_{f,t}} [A_{f,t}(l_t - l_{n,t}) + \frac{\varphi + 1}{\varphi} (A_{n,t} l_{n,t})^{\frac{\varphi}{\varphi+1}}] \quad (61)$$

$$l_{f,t} = l_t - l_{y,t} - l_{n,t} \quad (62)$$

Here we include the time subscripts to emphasize the drivers of the solution. This business-as-usual allocation is used to calibrate the productivities as explained in the text.

Calibration

In this section, we pull together all carbon cycle and economic parameters used in the quantitative assessment. The online supplementary material contains program files for replicating the results.

Calibration: Climate parameters

We describe a 3-layer/3-box representation of the global biogeochemical cycles. We denote the layers by the vector L_t , where $L_{1,t}$ represents the atmosphere plus the upper ocean layer, $L_{2,t}$ represents the biomass, and the third layer represents the deep ocean. Below, the layer system is transformed to a system of boxes that we denote by vector S_t . We denote the individual layers by j and the boxes by i .

The layers contain physical carbon stocks measured in $[TtCO_2]$. The atmosphere and upper ocean layer completely mix within the model's decadal period, but as the CO₂ stored in the upper ocean layer does not add to the greenhouse effect, we correct the

CO_2 in the first layer to find the atmospheric CO_2 stock

$$Atm_t = \frac{L_{1,t}}{1 + \mu}.$$

where μ is the amount of CO_2 stored in the upper ocean layer, relative to the amount in the atmosphere. In each period, share b_j of total emissions z_t enters layer j . The diffusion and between the layers is described through matrix \mathbf{M} that has real and distinct eigenvalues $\lambda_1, \lambda_2, \lambda_3$. No CO_2 leaves the system, so that row elements of \mathbf{M} sum to one, $\mathbf{M}(1, 1, 1)' = (1, 1, 1)'$, and the elements of b sum to 1. Dynamics satisfy

$$L_{t+1} = \mathbf{M}L_t + bz_t.$$

Using the eigen decomposition theorem of linear algebra, we can define the linear transformation of co-ordinates $H_t = \mathbf{Q}^{-1}L_t$ where $\mathbf{Q} = [v_1 \ v_2 \ v_3]$ is matrix of eigenvectors v_λ such that

$$\mathbf{Q}^{-1}\mathbf{M}\mathbf{Q} = \mathbf{\Lambda} = \mathbf{diag}[\lambda_1, \lambda_2, \lambda_3].$$

We obtain

$$\begin{aligned} H_{t+1} &= \mathbf{Q}^{-1}L_{t+1} = \mathbf{Q}^{-1}\mathbf{M}\mathbf{Q}H_t + \mathbf{Q}^{-1}bz_t \\ &= \mathbf{\Lambda}H_t + \mathbf{Q}^{-1}bz_t, \end{aligned}$$

which enables us to write the (uncoupled) dynamics per box i of the vector H_t as

$$H_{i,t+1} = \lambda_i H_{i,t} + c_i z_t$$

where λ_i are the eigenvalues and we defined the vector $c = \mathbf{Q}^{-1}b$. Then, we can solve for the development of the atmospheric CO_2 as

$$Atm_t = \frac{\sum_i q_{1,i} H_{i,t}}{1 + \mu}$$

where $q_{1,i}$ corresponds to the first row of \mathbf{Q} . Redefine $S_{i,t} = \frac{q_{1,i}}{1+\mu} H_{i,t}$, $a = \frac{q_{1,i}}{1+\mu} \mathbf{Q}^{-1}b$, and $\eta_i = 1 - \lambda_i$, to obtain

$$\begin{aligned} S_{i,t+1} &= (1 - \eta_i)S_{i,t} + a_i z_t \\ Atm_t &= \sum_i S_{i,t} \end{aligned}$$

Notice that we know one eigenvalue $\lambda = 1$, from which it follows that we have one box i with no depreciation, $\eta_i = 0$. If the model is run in almost continuous time, that is, with short periods so that most of the emissions enter the atmosphere, $b_1 = 1$, it follows that $\sum_i a_i = 1/(1 + \mu)$. Otherwise, we have $\sum_i a_i < 1/(1 + \mu)$.

We take data from Houghton (2003) and Boden et al. (2011) for carbon emissions in 1751–2008. We calibrate the layer model parameters \mathbf{M} , b , μ , to minimize the error between the atmospheric concentration prediction from the 3-layer model and the Mauna Loa observations, while maintaining CO_2 stocks in the various layers and flows between the layers consistent with scientific evidence as reported in Fig 7.3 from the IPCC fourth assessment report from Working Group I (Solomon et. al. 2007). There are 4 parameters we calibrate:⁵⁵ (1) the CO_2 absorption capacity of the “atmosphere plus upper ocean”, (2) the CO_2 absorption capacity of the biomass layer relative to the atmosphere, while we fix the relative size of the deep ocean layer at 4 times the atmosphere, based on the IPCC special report on CCS, Fig 6.3 (Caldeira and Akai, 2005). We furthermore calibrate (3) the speed of CO_2 exchange between the atmosphere and biomass and (4) between the atmosphere and the deep ocean. Subsequently, we transform this annual 3-layer model into a decadal layer model adjusting the exchange rates within a period between the layers and the shares of emissions that enter the layers within the period of emissions. Then, we transform the decadal 3-layer model into the decadal 3-box model, as described above. The transformed box model has no clear physical meaning other than this: box 0 measures the amount of atmospheric carbon that never depreciates; box 1 contains the atmospheric carbon with a depreciation of about 7 per cent in a decade; while carbon in box 2 depreciates 50 per cent per decade.⁵⁶ About 20 per cent of emissions enter either the upper ocean layer, biomass, or the deep ocean within the period of emissions. In the reduced-form model, they do not enter the atmospheric carbon stock, so that the shares a_i sum to 0.8. Our procedure provides an explicit mapping between the physical carbon cycle and the reduced-form model for atmospheric carbon with varying depreciation rates; the Excel file available as supplementary material contains these steps and allows easy experimentation with the model parameters. The resulting boxes, their emission shares, and depreciation factors are:

$$\begin{aligned} S_{t=2005} &= (.304, .250, .256) \\ a &= (.163, .184, .449) \\ \eta &= (0, .074, .470). \end{aligned}$$

⁵⁵We set $b = (1, 0, 0)$ for a short period. The matrix \mathbf{M} has 9 elements. The condition that the rows sum to one removes 3 parameters. We assume no diffusion between the biosphere and the deep ocean, removing 2 other parameters. We fix the steady state share of the deep ocean at 4 times the atmospheric share. This leaves us with 3 elements of \mathbf{M} to be calibrated, plus μ .

⁵⁶As explained above, the decay rates in the final reduced-form model come from the eigenvalues of the original model.

These coefficients are used to calculate the stock dynamics as given by (38) when the path of future emissions is given by the economy model. To transform stocks into damages as given by (39), we assume a 1-box damage model and choose the parameters as follows:

$$b = 1, \varepsilon = .183, \pi = 4.23$$

We thus have one box ($b = 1$). We interpret D_t in (39) as the Global Mean Temperature (GMT) squared ($GMT_t = \sqrt{D_t}$); $\varepsilon = .183$ in the decadal model implies a temperature adjustment speed of 2 per cent per year. Choice $\pi = 4.23 [K^2/GtCO_2]$ implies a climate sensitivity of 3 *Kelvin* per 2.129 $TtCO_2$. These choices are within the ranges of scientific evidence (Solomon et al. 2007).

We seek to calibrate the damage parameters to match the case presented in Nordhaus (2007) as a benchmark. Assuming damages equivalent to 2.7 per cent of output at a temperature rise of 3 Kelvin, as in Nordhaus (2001), we obtain $\Delta_y = 0.003$; we set $\Delta_u = 0$, unless otherwise stated.

Calibration: Economic parameters

In Section 4 (see also Appendix on the labor allocation solution) we explained the business-as-usual calibration, that is, the mapping from quantities path $(l, y, e_f, e_n)_{t \geq 0}$ to productivities $(A_y, A_e, A_f, A_n)_{t \geq 0}$. Thus, we need to make choices that drive the quantities $(l, y, e_f, e_n)_{t \geq 0}$. Population is assumed to follow a logistic growth curve:

$$l_{t+1} = [1 + \gamma_L(1 - \frac{l_t}{l_{\max}})]l_t$$

with parameters given by the World Bank forecasts. Population in 2010 (L) is set at 6.9 [billion], while the maximum population growth rate γ_L 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].

Consider then the determinants of initial output $y_{t=2010}$. We take Gross Global Product as 600 Trillion Euro [$Teuro$] for the first decade, 2006-2015 (World Bank, using PPP). Fossil-fuel energy input e_f , measured in CO_2 , is .318 [$TtCO_2$] for 2006-2015 (SRES IPCC 2000); we set e_n at 10 per cent of e_f . These are the raw quantities entering production function at $t = 2010$; now, we use the model structure to calibrate the remaining variables at $t = 2010$. There is only one energy sector parameter to be set: the elasticity of carbon-free supply; see Section 4. We set this parameter to $\varphi = 2$.

We calibrate the preference parameters to yield 25 per cent savings ($g = .25$), which together with $u(c) = \ln(c)$ and $y_{t=2010} = 600$, gives the price of output as $p = 1/((1 - g)600)$. The relative price of energy is taken to be $q/p = 50$ [*euro/tCO₂*]. In Appendix for the labor allocation, we derived the labor allocation as a function of productivities. We have now information on the endogenous variables: the energy outputs (e_n, e_f), total output y and the energy price q/p , so we can solve both the labor allocation (l_y, l_f, l_n) and the productivities (A_y, A_e, A_f, A_n) at $t = 2010$.

To progress to the next decade $t = 2020$, we take capital k given by savings, keep the energy price $q/p = 50$ [*euro/tCO₂*], and match y and (e_n, e_f) to the A1F1 SRES scenario from the IPCC (2000). This way the calibration procedure for productivities can be repeated for all future decades.

Finally, capital elasticity α follows from the assumed time-preference structure β and δ , and observed historic gross savings g . As a base-case, we consider net savings of 25% ($g = .25$), and a 2 per cent annual pure rate of time preference ($\beta = 1, \delta = 0.817$), resulting in $\alpha = g/\rho = 0.306$. For the Markov equilibrium, we take $\beta\delta = .7724$ and $\delta = .9511$. These choices preserve $g = .25$.

Sea level rise in damages

Climate change does not stop at temperature changes. After temperature rises, the sea level will rise as well, and it may do so more or less proportionally to temperatures (Jevrejeva et al. 2011). When damages are proportional to output and quadratic in the level of sea-level rise, we can use SLR_t for the damages associated with sea level rise, and write for the dynamics of damages

$$SLR_t = SLR_{t-1} + \varepsilon_{SLR}(\pi_{SLR}D_t - SLR_t)$$

It is a tedious but straightforward matter to derive the resulting dynamics as

$$SLR_t = \sum_i \sum_j \sum_\tau a_i b_j \pi \pi_{SLR} \varepsilon_j \varepsilon_{SLR} \times \frac{(\varepsilon_j - \varepsilon_{SLR})(1 - \eta_i)^{\tau+1} + (\varepsilon_{SLR} - \eta_i)(1 - \varepsilon_j)^{\tau+1} + (\eta_i - \varepsilon_j)(1 - \varepsilon_{SLR})^{\tau+1}}{(\varepsilon_j - \eta_i)(\varepsilon_j - \varepsilon_{SLR})(\varepsilon_{SLR} - \eta_i)} z_{t-\tau}$$

where we forego the terms associated with the initial conditions. Let Δ_{SLR} be the costs relative to output of 1m sea level rise. After some tedious substitutions, the Markov

equilibrium carbon price is then calculated as

$$\rho\zeta_{SLR} = \sum_i \sum_j a_i b_j \pi \pi_{SLR} \varepsilon_j \varepsilon_{SLR} \rho \Delta_{SLR} \times \frac{(\varepsilon_j^2 - \eta_i^2) \varepsilon_{SLR} + (\varepsilon_{SLR}^2 - \varepsilon_j^2) \eta_i + (\delta_i^2 - \varepsilon_{SLR}^2) \varepsilon_j}{(\varepsilon_j - \delta_i)(\varepsilon_j - \varepsilon_{SLR})(\varepsilon_{SLR} - \eta_i)(1 - \theta[1 - \varepsilon_{SLR}])(1 - \theta[1 - \varepsilon_j])(1 - \theta[1 - \eta_i])}$$

For sea level rise, both the sensitivity π_{SLR} and the speed of adjustment ε_{SLR} are both very uncertain, but the estimates for both parameters are strongly and negatively correlated: a higher sensitivity must be matched with a lower adjustment speed, to match the historically observed records, $\varepsilon_{SLR} \pi_{SLR} \in [0.02, 0.05]$ (per decade). Estimates for the sensitivity range from 0.2 to 5 meter sea level rise per W/m² forcing increase (Jevrejeva 2011), where typically 1 W/m² leads to a temperature rise just below 1 Kelvin, so that our parameters range would be $\pi_{SLR} \in [0.2, 4]$. A choice of $\pi_{SLR} = 1$ [m/K] and $\varepsilon_{SLR} = 0.04$ would represent a reasonable assumption. Yet, the resulting carbon prices will not deviate too much from the carbon prices presented in the main text. The literature does not provide estimates for damages associated with sea level rise that substantially exceed those for temperature rise, that is $\Delta_{SLR} \pi_{SLR} \ll \Delta_y$, so that, given the extended lag in sea level rise, the increase in the level of carbon prices associated will be small in relative terms.

Comparison of climate response functions

We compare our response function for damages, as percentage of output, resulting from emissions, with those in Nordhaus (2007) and Golosov et al. (2011). The GAMS source code for the DICE model provides a large variety of scenarios with different policies such as temperature stabilization, concentration stabilization, emission stabilization, the Kyoto protocol, a cost-benefit optimal scenario, and delay scenarios. For each of these scenarios we calculated the damage response function by simulating an alternative scenario with equal emissions, apart from a the first period when we decreased emissions by 1GtC. Comparison of the damages, in terms relative of output, then defines the response function for that specific scenario. It turns out that the response functions are very close, and we took the average over all scenarios. To interpret the response function in Nordhaus (2007), we notice that the average DICE carbon cycle and damage response can very accurately be described by our reduced form using the parameters $a = (0.575, 0.395, 0.029)$, $\eta = (0.310, 0.034, 0)$, which give a perfect fit for the carbon cycle of DICE2007, and $\varepsilon = 0.183$, $\pi = 4.09$ for the temperature delay. That is, the

carbon-cycle in DICE (Nordhaus 2007) is characterized by a very large long-term uptake of CO₂ in the oceans. The reduced model in Golosov et al. is represented by $a = (0.2, 0.486, 0.314)$, $\eta = (0, 0.206, 1)$, which implies a similar carbon cycle model to ours, but Golosov et al. have no temperature delay structure, $\varepsilon = 1$. Figure 1 presents the emissions damage responses.

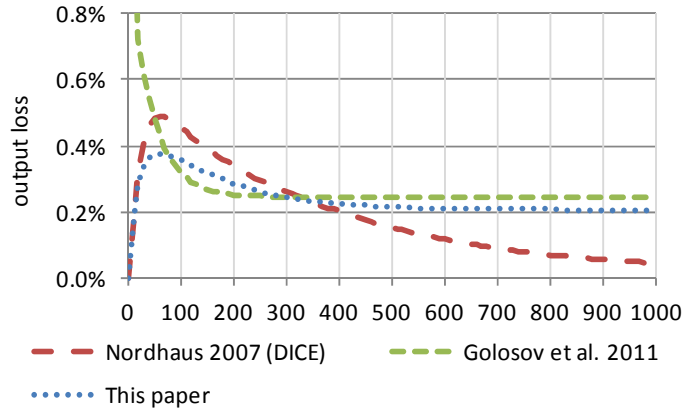


Figure 1. Emissions-Response function in 3 models. Damages as share of output for a 1 TtCO₂ impulse.

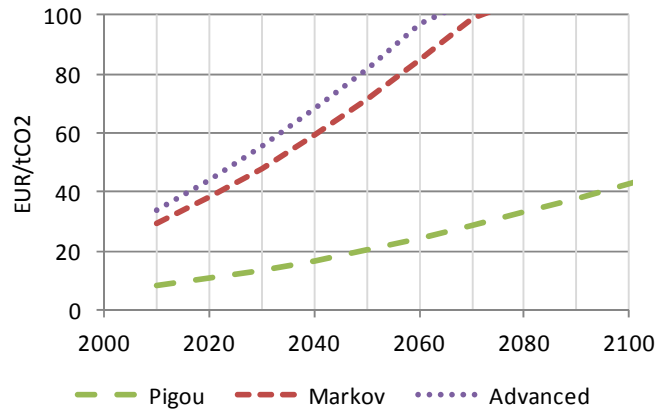


Figure 2. Carbon prices in 3 scenarios

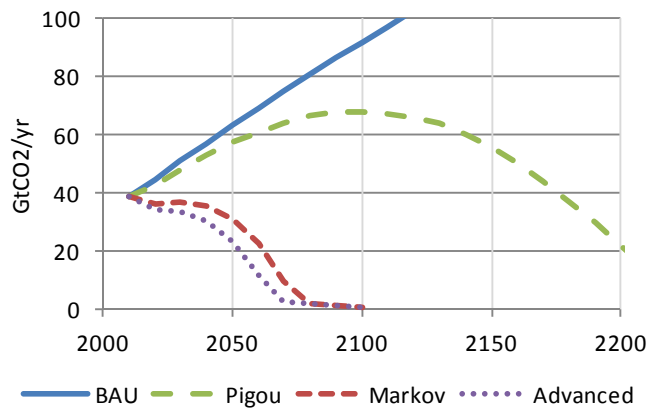


Figure 3. CO₂ emissions, per year, in 4 scenarios




Figure 4. Temperature rise in 4 scenarios

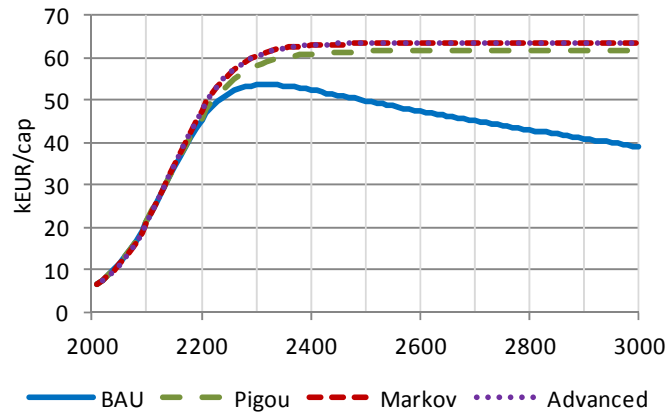


Figure 5. Per capita consumption levels in 4 scenarios

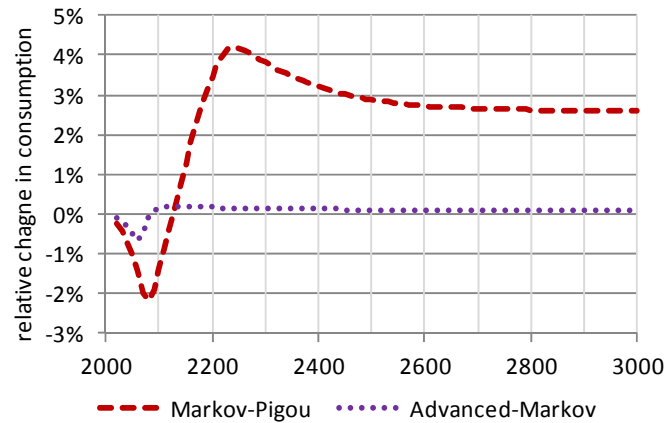


Figure 6. Change in consumption, Markov vs. Pigou and Advanced vs. Markov