Copyedited by: ES MANUSCRIPT CATEGORY: Article

Review of Economic Studies (2020) **0**, 1–36 doi:10.1093/restud/rdaa048

Advance access publication 21 August 2020

Carbon Taxes and Climate Commitment with Non-constant Time Preference

TERRENCE IVERSON

Department of Economics, Colorado State University

and

LARRY KARP

Department of Agricultural and Resource Economics, University of California, Berkeley

First version received May 2017*; Editorial decision January* 2020*; Accepted June* 2020 (*Eds.*)

We study the Markov perfect equilibrium in a dynamic game where agents have non-constant time preference, decentralized households determine aggregate savings, and a planner chooses climate policy. The article is the first to solve this problem with general discounting and general functional forms. With time-inconsistent preferences, a commitment device that allows a planner to choose climate policy for multiple periods is potentially very valuable. Nevertheless, our quantitative results show that while a permanent commitment device would be very valuable, the ability to commit policy for "only" 100 years adds less than 2% to the value of climate policy without commitment. We solve a log-linear version of the model analytically, generating a formula for the optimal carbon tax that includes the formula in Golosov et al. (2014, *Econometrica*, **82**, 41–88) as a special case. More importantly, we develop new algorithms to solve the general game numerically. Convex damages lead to strategic interactions across generations of planners that lower the optimal carbon tax by 45% relative to the scenario without strategic interactions.

Key words: Hyperbolic discounting, Time consistency, Climate policy

JEL Codes: D61, D62, E61, Q51, Q54.

1. INTRODUCTION

The pure rate of time preference is an important determinant of the discount rate and thus important for climate policy. The assumption of constant time preference was challenged from its inception on normative grounds (Samuelson, 1937), and it has recently been challenged on empirical grounds. Constant time preference leads to time-consistent preferences (Koopmans, 1960), which simplifies the calculation of equilibrium behaviour. This simplicity is especially helpful when solving Integrated Assessment Models, which combine climate policy and private savings. To incorporate non-constant time preference (NCTP) into an Integrated Assessment Model, we need to compute a subgame perfect equilibrium among generations (Strotz, 1955),

The editor in charge of this paper was Michele Tertilt.

while also imposing a fixed-point condition within each period to ensure consistency between individual and aggregate savings rules (Krusell et al., 2002). Our article is the first to solve this problem in a model with arbitrary discounting and general functional forms.

The choice between constant and non-constant time preference can have a large effect on the optimal carbon tax. Agents with NCTP may invest in capital at observed rates and also incur significant costs to avoid climate damages in the distant future (Gerlagh and Liski, 2017). The conflict among generations arising from NCTP gives current policy-makers an incentive to develop a commitment device to constrain future choices, just as with private agents (Strotz, 1955; Laibson, 1997). We assess the value of commitment and examine the effect of NCTP on climate policy and savings. Our normative analysis examines the optimal climate policy for a global planner, subject to the restrictions imposed by a limited ability to commit to future policy.

Arguments for NCTP suggest a path of time preference rates that declines over many periods. Most macroeconomic applications approximate this declining path with quasi-hyperbolic (or β, δ) discounting. Here, utility n periods in the future is discounted by factor $\beta\delta^n$ (Phelps and Pollak, 1968; Harris and Laibson, 2001; Krusell and Smith, 2003). Exponential discounting corresponds to the case $\beta = 1$; quasi-hyperbolic discounting corresponds to the case $\beta < 1$. The latter implies present bias. Apart from special cases, constant discounting, quasi-hyperbolic discounting, and general discounting are not observationally equivalent. The different types of discounting can have both a qualitative and a large quantitative effect on policy.[1]

If the actual discount rate declines over many periods, the approximation error due to the restriction to quasi-hyperbolic discounting grows in the decision horizon; the error is likely large for climate policy, where the relevant horizon spans centuries. In addition, the quasi-hyperbolic approximation misrepresents the nature of intergenerational decision conflict stemming from the difference between the discount rate future agents will use and the rates the current agent would like them to use. With quasi-hyperbolic discounting, this conflict is constant after one period, while with continuously decreasing discount rates it grows with the distance between generations.

There are two ways to think about the dynamic game when private agents have NCTP and the climate planner internalizes their preferences. The first is a two-instrument problem where a planner in each period controls both an emissions tax and an investment tax. The second is a single-instrument problem in which planners control an emissions tax only. We emphasize the second approach, considering the two-instrument setting only to provide a comparison. For a log-linear special case (Definition 3), we show that welfare is actually higher without the investment tax. Because policymakers are unlikely to make the effort to coordinate on a global investment tax when doing so lowers welfare, we view the setting with only an emissions tax as more relevant for policy.[2] The single-instrument setting entails a second-best optimal tax problem in which a planner each period takes as given both the optimal response of future planners and the equilibrium savings response of decentralized households. The optimal tax is the Markov (perfect) equilibrium of the corresponding dynamic game.

The single-instrument problem presents previously unsolved technical difficulties that stem from the need to solve a sequence of fixed-point problems. As noted in Krusell et al. (2002),

1. Harstad (2020, p. 4) illustrates the qualitative impact in an investment game. He notes: "…the result that upstream technologies should be subsidized more does *not* hold under quasi-hyperbolic discounting—which is therefore a poor approximation for hyperbolic discounting." Iverson (2012) documents the quantitative impact of the quasi-hyperbolic restriction.

2. Even if other functional forms were to reverse the welfare ranking, the difficulty of achieving global coordination of investment taxes, combined with the lack of observed efforts to do so, provides further justification for focusing on the problem without an investment tax.

the need to solve a fixed-point problem within each period does not arise when the planner also controls investment. The fixed-point problems exacerbate the known numerical instability associated with approximation methods in non-stationary dynamic programs (Cai and Judd, 2014). By exploiting analytic structure of the problem, we develop a novel numerical approach that solves the problem in a stable and accurate way. The methods developed can be applied with a general production function, a general climate module, and a general discount function, provided the single-period utility function is isoelastic.

Our principle policy question concerns the value of commitment in climate regulation when households have time-inconsistent preferences. With declining time preference, future generations prefer less climate regulation than the initial generation would like. Because a commitment device resolves this conflict, it is potentially valuable. Quantitatively, however, we find that although a permanent commitment device would be very valuable, the ability to commit over realistic horizons is worth little. In the log-linear model, a 50-year commitment device is worth approximately one-seventh the value of a 200-year commitment device. The large difference arises because the degree of decision conflict between planners in different periods increases in the number of periods between them. Relatedly, the value of commitment over realistic horizons is tiny compared to the value of climate policy without commitment. In the baseline calibration with more general functional forms, climate policy without commitment is worth 4.6% decadal Gross World Product, while a 100-year commitment adds only an additional 0.06%.

The findings have two important policy implications. First, because any commitment mechanism is likely to be expensive, and because these mechanisms produce little welfare gain, society should not attempt them. For example, we might consider making substantial and premature sunk investments in a new technology or infrastructure to induce our successors to complete the project; our investment reduces their completion cost. More premature investments have greater potential to influence future generations but are more expensive. Our results caution against using such investments to induce commitment. Second, we might be concerned that the value to future generations of contemporary altruism, reflected in a declining pure rate of time preference, could be largely undone by the inability to commit. We find that this theoretical possibility has little practical importance.

The article makes a further contribution by examining the effect on optimal climate policy and competitive equilibrium savings of changes in the depreciation rate, the elasticity of intertemporal substitution, the damage function, and the growth rate of Total Factor Productivity (TFP). If damages are strongly convex in the stock of carbon, climate policies are dynamic strategic substitutes, meaning a decrease in current emissions induces future generations to emit more. This strategic interaction causes the optimal carbon tax to be lower, all else equal. In the baseline calibration, the optimal carbon tax in the Markov equilibrium of the intergenerational game is 45% lower than it would be if strategic interactions among generations did not arise. When damages are roughly linear in the stock of carbon, as in Nordhaus (2013) and Golosov et al. (2014), strategic incentives across generations are insignificant even when the model contains other sources of nonlinearity.

Finally, we analytically solve a generalization of Gerlagh and Liski's (2017) log-linear model (Section 3). The log-linear structure decouples decisions across generations and between savings and climate policy. These decoupling properties give rise to the model's analytic tractability, but they limit its ability to shed light on intergenerational conflict and on the links between savings and climate policy. Our most important contributions stem from our ability to solve the general model. Nevertheless, because much of the prior literature focuses on the log-linear case and because it provides a simple setting in which to develop intuition, we examine that case first. Section 4 considers the general setting.

1.1. *Literature review*

We first explain our contributions relative to Krusell et al. (2002), Golosov et al. (2014), and Gerlagh and Liski (2017). We then review other climate applications with NCTP and summarize empirical evidence on declining time preference.

Krusell et al. (2002) study savings decisions in the Markov equilibrium of a log-linear economy with quasi-hyperbolic discounting. They show that the planner's ability to control current investment, *e.g.*, by using an investment tax, lowers equilibrium investment. The planner's recognition that additional savings lowers the return to capital reduces her incentive to save relative to that of price-taking agents. Those agents already save too little, so allowing the current planner to control investment, but not to commit to future policies, reduces welfare because it reduces savings. Hiraguchi (2014) extends this welfare comparison to the general discounting case; we further generalize it to the non-loglinear case and to a setting with climate damage and climate policy (Corollary 3).

Golosov et al. (2014) extend Brock and Mirman's (1972) log-linear setting to include a climate module, leading to an explicit formula for the optimal carbon tax that can be easily calibrated. They assume a constant pure rate of time preference. Gerlagh and Liski (2017) combine elements of both Krusell et al. (2002) and Golosov et al. (2014), incorporating quasi-hyperbolic discounting and extending the climate module in Golosov et al. (2014) to include delay between carbon emissions and damages. Our article follows Gerlagh and Liski (2017) in studying climate policy in a neoclassical economy with capital accumulation and time-inconsistent preferences. But the similarity largely ends there. Gerlagh and Liski (2017) consider only the log-linear case with quasi-hyperbolic discounting, while we study the case with general discounting and general functional forms.[3] They emphasize the policy problem where planners control both carbon emissions and investment, while we emphasize the case where planners control climate policy only. The two papers also adopt different timing conventions in the stage game, which lead to important differences in the carbon tax, unrelated to the differences in functional forms (Section 3). Finally, the papers examine different policy questions. Gerlagh and Liski (2017) emphasize the fact that current emissions have a delayed but persistent effect on future welfare. A reduction in carbon emissions provides transfers from the current generation to the distant future, without the intermediation of intervening generations. They refer to the delayed and persistent effect of policy as a commitment device. In contrast, we study the value of a literal commitment device, one that enables current policymakers to choose future carbon taxes.

Our article also builds on an earlier literature that studies climate policy with NCTP. Karp (2005, 2017) and Karp and Tsur (2011) use NCTP in analytic climate models, and Fujii and Karp (2008) use numerical methods, but these abstract from the savings decision and include only a rough approximation of climate dynamics. Iverson et al. (2015) use the formulae in Proposition 1 to show that long-term features of the carbon cycle are unimportant for near-term climate policy under constant time preference but important under hyperbolic discounting, even when both models are calibrated to match historical savings rates. Rezai and Van der Ploeg (2017) study a partial equilibrium model under the assumption that the initial planner can commit to the sequence of future policies.

Most of the empirical and macro-theory studies of NCTP take the decision-maker as a consumer or a firm. However, one of the earliest applications involves a dynasty consisting of many generations, each of which lives for a single period (Phelps and Pollak, 1968). Climate policy involves both intra- and intergenerational transfers. Current abatement can benefit both a

3. Our version of the log-linear model with arbitrary non-constant time preference is as analytically tractable as Gerlagh and Liski's (2017) extension of the Golosov et al. (2014) climate module.

TABLE 1
Recent estimates of the β, δ model. $\beta < 1$ implies present bias.

Article	Decision involving	β	δ
Shui and Ausebel (2005)	Credit cards	0.8	0.99
Paserman (2008)	Job search	(0.4, 0.9)	≈ 1
Fang and Wang (2015)	Labour supply and welfare	0.34	0.88
Mahajan and Tarozzi (2012)	Technology adoption	0.56	0.79
Fang and Wang (2015)	Mammography	(0.51, 0.79)	(0.68, 0.94)
Laibson et al. (2017)	Consumption over lifecycle	0.504	0.987

currently living agent later in her life and people not yet born. Our model is consistent with both an infinitely lived agent with NCTP or, as in Ekeland and Lazrak (2010), with an overlapping generations model where altruistic agents have constant mortality risk.

Frederick et al. (2002) conclude that empirical evidence overwhelmingly supports hyperbolic discounting over exponential discounting. Table 1 summarizes recent empirical estimates. The literature assumes quasi-hyperbolic preferences and usually provides estimates of present bias ($\beta < 1$). Augenblick et al. (2015) find present bias in consumption, concluding that people value commitment. Survey evidence suggests that people discount the distant future less heavily than the near future (Cropper et al., 1994; Layton and Levine, 2003; Drupp et al., 2014). Montiel Olea and Strzalecki (2014) provide an axiomatic foundation for quasi-hyperbolic discounting.[4]

Brocas et al. (2004) and Bryan et al. (2010) survey the literature on commitment devices. Examples of these include mandatory pension plans with limited accessibility, mandatory minimum years worked before retirement, and externally enforced deadlines. Barro (1999) shows that short-run savings falls, and long-run savings rises, with the length of a commitment device.

2. MODEL

After describing preferences, the climate, and technology, we discuss the equilibrium concept. Time is discrete and runs from 0 to infinity. There is no uncertainty. A unit continuum of identical households discount future utility with arbitrary time weights $\{\lambda_m\}_{m=0}^{\infty}$, where $\lambda_0 = 1$. The weights are unrestricted, but we emphasize the case of decreasing time preference rates. Household welfare (assumed finite) in t is the present value of discounted utility,

$$\sum_{\tau=0}^{\infty} \lambda_\tau u(c_{t+\tau}),$$

where utility is isoelastic in consumption:

$$u(c) = \begin{cases} \frac{c^{1-\eta}-1}{1-\eta}, & \text{if } \eta > 0, \eta \neq 1 \\ \ln(c), & \text{if } \eta = 1. \end{cases} \tag{2.1}$$

Agent i at t owns capital $k_{i,t}$, and aggregate capital is

$$K_t = \int_0^1 k_{i,t} \mathrm{d}i.$$

4. A growing literature studies reasons for non-constant *consumption* discount rates arising from the correlation between the returns to climate investments and a market portfolio (Traeger, 2014; Giglio et al., 2015; Dietz et al., 2018). Our deterministic model cannot include this feature.

We use a representative agent model and drop the agent index when the meaning is clear.

Aggregate carbon emissions, E_t, affect the evolution of a vector of climate states, $\mathbf{S}_t$, which could include carbon stocks and temperatures in different reservoirs. The climate states evolve according to

$$\mathbf{S}_{t+1} = \mathbf{f}(\mathbf{S}_t, E_t). \tag{2.2}$$

Output of the final good is

$$Y_t = F_t(K_t, E_t, \mathbf{S}_t). \tag{2.3}$$

Symbols in bold denote vectors. $F_t(\cdot)$ is concave, increasing in capital and emissions and decreasing in the climate state, with constant returns to scale in K, E, and labour (normalized to 1); the subscript t accommodates exogenous technological change. Climate damages arise from output losses, affecting consumption; we abstract from damage channels directly impacting utility, *e.g.*, a loss in amenity value.[5] Current output equals aggregate consumption, C_t, plus investment,

$$Y_t = C_t + K_{t+1} - (1-\delta)K_t, \tag{2.4}$$

where δ is the depreciation rate for capital.

Most carbon emissions result from burning fossil fuel, whose stocks may be exhaustible. We ignore stock constraints, making the model appropriate for coal, but probably not for oil. An earlier version of the article discussed modifications to include an exhaustible resource.

2.1. *Equilibrium savings and climate policy*

The climate policy problem is situated within a competitive economy where the decentralized decisions of price-taking households determine equilibrium savings, price-taking firms maximize profit, and both households and firms take as given the aggregate impact of climate policy on current and future prices. We begin with the household problem, then consider the planner's problem. To simplify our description of the dynamic game, we leave firms' input choices in the background until Section 5.1.

Under the assumption that revenue from a carbon tax (or cap and trade system) is returned in a lump sum to households, the income unrelated to the agent's ownership of capital equals the agent's wage plus the tax revenue (or quota rents). We define the sum of these payments as "transitory income" and denote it by w_t. Given constant returns to scale, w_t equals the value of output minus payments to capital.

In period t, agent i earns the return on capital r_t. With capital stock k_t, the agent earns income from capital $r_t k_t$ and receives transitory income w_t. We study the Markov equilibrium of the savings game where agent i chooses current savings, taking as given her future savings policies. Agents have zero mass and rational expectations, so they take as given the trajectories of aggregate capital, aggregate emissions, and climate states—which depend on equilibrium climate policy.

Suppressing the agent index i, the representative agent in t takes as given its successors' savings rules, $g_s(k_s, K_s, \mathbf{S}_s; E_s)$, $s > t$, and the current and future aggregate savings rules, $G_s(K_s, \mathbf{S}_s; E_s)$, $s \geq t$. The current emissions level, E_s, affects the level of savings via its effect on current income. The agent chooses k_{t+1} to maximize

$$u\left(r_t k_t + w_t + (1-\delta)k_t - k_{t+1}\right) + \tilde{V}\left(k_{t+1}, G_t(K_t, \mathbf{S}_t; E_s), \mathbf{S}_{t+1}, t+1\right), \tag{2.5}$$

5. The Supplementary Appendix describes an extension to include population growth and, under appropriate separability conditions, damages directly affecting utility.

where

$$\tilde{V}\left(k_{t+1}, K_{t+1}, \mathbf{S}_{t+1}, t+1\right) \equiv \sum_{j=1}^{T} \lambda_j u\left(r_{t+j} k_{t+j} + w_{t+j} + (1-\delta) k_{t+j} - g_{t+j}\left(k_{t+j}, K_{t+j}, \mathbf{S}_{t+j}\right)\right). \quad (2.6)$$

The optimum defines $g_t(k_t, K_t, \mathbf{S}_t; E_s)$, leading to agent welfare

$$V(k_t, K_t, \mathbf{S}_t, t) \equiv u(r_t k_t + w_t + (1-\delta) k_t - g_t(k_t, K_t, \mathbf{S}_t)) + \tilde{V}\left(g_t(k_t, K_t, \mathbf{S}_t), G_t(K_t, \mathbf{S}_t; E_s), \mathbf{S}_{t+1}, t+1\right). \quad (2.7)$$

In equilibrium, individual and aggregate capital must be equal. We obtain equilibrium welfare, denoted $V^e(\cdot)$, by imposing the condition $k = K$. Because the dependence of equilibrium welfare on current and future climate policy is important for describing the climate problem, we make it explicit:

$$V^e\left(K_t, \mathbf{S}_t, t; \{E_s\}_{s=t}^{\infty}\right) \equiv V(K_t, K_t, \mathbf{S}_t, t).$$

We turn next to the determination of climate policy. In each period, a planner chooses policy to maximize equilibrium welfare for the contemporaneous representative household. Because household preferences are time inconsistent, we model climate policy as a Markov equilibrium of a dynamic game. Planners take the equilibrium aggregate savings rule, $G_t(K_t, \mathbf{S}_t; E_t)$ and future climate policy rules as given. There may be many Markov equilibria, but we follow Krusell et al. (2002) in focusing on the limit equilibrium: the limit, as the horizon goes to infinity, of the equilibria to a sequence of finite horizon problems.[6]

To study the role of commitment, we introduce the following technology.

Definition 1 *(Commitment device) A planner in t with a j-period commitment device chooses climate policy, E_s, for periods $s = t, \ldots, t+j-1$, provided that the emissions level during these periods was not already fixed by a commitment device in an earlier period.*

We denote the Markov equilibrium emissions policy in period s by $E_s^*(K_s, \mathbf{S}_s)$. To simplify notation, we suppress the dependence of optimal policy on the commitment window, j. The planner in t with j-period commitment anticipates that planners outside the commitment window will choose $\{E_s\}_{s=t+j}^{\infty} = \left\{E_s^*(K_s, \mathbf{S}_s)\right\}_{s=t+j}^{\infty}$. The planner thus chooses $\{E_s\}_{s=t}^{t+j-1}$ to maximize welfare for the period-t representative agent. The planner solves

$$\max_{\{E_s\}_{s=t}^{t+j-1}} V^e\left(K_t, \mathbf{S}_t, t; \{E_s\}_{s=t}^{t+j-1}, \left\{E_s^*(K_s, \mathbf{S}_s)\right\}_{s=t+j}^{\infty}\right). \quad (2.8)$$

Combining the capital market equilibrium and the climate policy equilibrium, a Markov perfect equilibrium for the model is defined as follows.

Definition 2 *A Markov perfect equilibrium satisfies the following for all t:*

(i) Prices, determined competitively, satisfy

$$\text{(rental rate) } r_t = r_t(K_t, E_t, \mathbf{S}_t) = \frac{\partial F_t(K_t, E_t, \mathbf{S}_t)}{\partial K_t},$$

6. For both the savings and climate policy decisions, we prove uniqueness in the log-linear model, and we find no numerical evidence to indicate non-uniqueness in the general model.

$$\textit{(transitory income)}\ w_t = w_t(K_t, E_t, \mathbf{S}_t) = F_t(K_t, E_t, \mathbf{S}_t) - \frac{\partial F_t(K_t, E_t, \mathbf{S}_t)}{\partial K_t} K_t.$$

(ii) The physical constraints in (2.2)*,* (2.3)*, and* (2.4) *hold.*

(iii) Individual savings maximizes equilibrium welfare, defined in (2.5) *and* (2.6)*.*

(iv) Individual and aggregate savings are consistent: $g_t(K_t, K_t, \mathbf{S}_t; E_t) = G_t(K_t, \mathbf{S}_t; E_t)$*.*

(v) The climate policy function solves (2.8)*.*

The solution to the model consists of time-dated equilibrium aggregate savings rules and emissions policies. We describe the emissions policy using the carbon tax, $\tau \equiv \frac{\partial F_t(K_t, E_t, \mathbf{S}_t)}{\partial E}|_{E_t = E^*}$, that supports the equilibrium emissions level E^*.

3. ANALYTICAL SOLUTION FOR THE LOG-LINEAR MODEL

Before considering the general case, we derive results in a model with a closed-form savings rate and carbon tax. Golosov et al. (2014) use the same assumptions in a model with constant time preference, and Gerlagh and Liski (2017) use them with quasi-hyperbolic discounting. In addition to using a general discounting function, our treatment of the log-linear economy differs from Gerlagh and Liski (2017) because the two papers study different planning problems. In the scenario where the planner has a single instrument (enabling her to control current emissions but not investment), we assume that the planner takes as given the current investment *rule*; Gerlagh and Liski assume that the planner takes as given the current investment *level*. Our assumption leads to "intra-temporal decoupling" (Corollary 1) between the climate and investment problems. This decoupling implies that the carbon tax is the same whether or not the planner controls investment by means of an investment tax. The equivalence implies that welfare is higher when the planner does not use an investment tax (Corollary 3).[7]

In the remainder of the section, we state the functional assumptions of the log-linear model, then characterize the equilibrium—first, for the case without commitment, then with commitment. Finally, we study the value of commitment.

3.1. *Functional assumptions*

Golosov et al. (2014) extend the Brock–Mirman growth model (Brock and Mirman, 1973)—log utility, 100% depreciation of physical capital, and Cobb Douglas production—by assuming that climate damages are multiplicative with a particular "linear-exponential" functional form. Because the resulting equilibrium is linear in the climate states and linear in the log of capital, we refer to it as the log-linear model.[8]

Definition 3 *The log-linear model*

(i) The utility function is logarithmic: $u(C) = \log(C)$*.*

(ii) Capital depreciates 100 percent in every period: $\delta = 1$*.*

(iii) Output in the final-goods sector is Cobb Douglas in capital with multiplicative climate damages:

$$Y_t = (1 - D(S_t)) K_t^{\alpha} A_t(E_t),$$

7. The Supplementary Appendix explains why we prefer our assumption about the stage game. It also explains why the different assumptions alter the effect on the carbon tax of introducing a second policy instrument. We also illustrate the magnitude of the difference, showing its quantitative importance.

8. Traeger (2016) uses weaker assumptions that preserve the log-linear structure, and he also incorporates temperature inertia and different types of uncertainty.

where $A_t(E_t)$ is a time-dependent function that incorporates changes in technology and labour supply; the damage function is exponential:

$$1-D(S_t)=\exp\left(-\gamma(S_t-\bar{S})\right).$$

(iv) The climate vector, **S**, *affects output only via atmospheric carbon, denoted S, and S is linear in prior emissions:*

$$S_t-\bar{S}=\sum_{j=0}^{t+H}(1-d_j)E_{t-j},$$

where $\bar{S}$ is the preindustrial stock of atmospheric carbon, H is the number of periods between industrialization and time zero, $1-d_j$ is the portion of emissions remaining in the atmosphere after j periods, and $0\leq d_j\leq 1$.

3.2. *Equilibrium without commitment*

We first solve the model for the case where planners choose policy for a single period only. Proposition 1 shows that the limit equilibrium is unique, and it provides formulae for the equilibrium savings rule and the equilibrium carbon tax. The proofs are in Appendix A. We use the following notation:

$$\rho\equiv\sum_{t=1}^{\infty}\lambda_t \text{ and } \lambda_{l,m}\equiv\frac{\lambda_m}{\lambda_l} \text{ for } m\geq l.$$

We assume ρ is finite. $\lambda_{l,m}$ is the amount by which an agent in t would be willing to reduce utility in $t+l$ to obtain a one-unit increase in utility in $t+m$.

Proposition 1 *In the log-linear model (Definition 3) without commitment ($j=1$) the limit equilibrium is unique. The equilibrium carbon tax and equilibrium savings are constant fractions of income. The equilibrium competitive savings rule for an individual in t with capital k is*

$$k'=\frac{\rho}{1+\rho}r_t k,$$

where r_t is the rental rate (Definition 2); the aggregate (competitive) savings rule is

$$K'=sY_t, \text{ where } s\equiv\frac{\alpha\rho}{\rho+1}; \tag{3.9}$$

and carbon taxes are

$$\tau_t=\left[\sum_{k=0}^{\infty}\lambda_k(1-d_k)\gamma\cdot\Gamma_k\right]Y_t, \tag{3.10}$$

where

$$\Gamma_k\equiv\frac{\sum_{m=0}^{\infty}\alpha^m\lambda_{k,k+m}}{\sum_{n=0}^{\infty}\alpha^n\lambda_n}. \tag{3.11}$$

Corollary 1 Intra-temporal decoupling*: The equilibrium aggregate savings rate does not depend on emissions levels; equilibrium emission levels do not depend on the aggregate savings rate.*

Corollary 2 Inter-temporal decoupling: *The equilibrium emissions tax policy (the tax expressed as a fraction of income) in any period is independent of climate policy in all other periods.*

Corollary 3 Welfare ranking*: Equilibrium welfare is higher when planners choose an emission tax only, compared to the two-instrument case where they also control investment.*

The equilibrium carbon tax is a constant fraction of output, where the constant depends on the path of time preference rates, the damage elasticity parameter, γ, the carbon cycle parameters, and the Cobb Douglas coefficient on capital in final-goods production. Notably, the tax does not depend on beliefs about future technology or emissions. The same features hold for the optimal tax in Golosov et al. (2014) with the exception that the formula there does not depend on the Cobb Douglas coefficient on capital. With constant time preference, $\Gamma_k = 1$, and the expression in (3.10) reduces to the tax formula in Golosov et al. (2014). When time preference rates decline, $\Gamma_k > 1$. In this case, the equilibrium tax is greater than the tax one would obtain if the constant-discounting time weights in the Golosov et al. (2014) formula were simply replaced with non-constant-discounting time weights.

The functional forms in Definition 3 produce intertemporal decoupling between current and future emissions decisions (Corollary 2). This corollary—and our assumption that the planner takes the current savings rule (not the current level of savings) as given—lead to the decoupling of policies within a period (Corollary 1). This corollary underlies the welfare ranking (Corollary 3).

The Supplementary Appendix contains a heuristic derivation of the equilibrium carbon tax. The derivation provides intuition for the relationship between the optimal taxes under constant and non-constant time preference. In both cases, the carbon tax equals the social cost of carbon, defined as the stream of future damages (measured in units of the consumption good) due to an extra unit of emissions, weighted by the appropriate marginal rates of substitution. With constant discounting, the relevant marginal rate of substitution is the familiar ratio of marginal utilities of consumption in two periods. With NCTP, in contrast, the relevant marginal rate of substitution involves ratios of two shadow values of income. The shadow value of income is the derivative with respect to income of the following value function:

$$V_t(Y_t) \equiv \sum_{v=0}^{\infty} \lambda_v \ln C_{t+v},$$

where the relationship between Y_t and future consumption (C_{t+v}) can be pinned down using the formulae in Proposition 1. The two definitions of marginal rates of substitution are equivalent under constant discounting but differ under NCTP because their future savings rules are not optimal from the perspective of the current planner.

3.3. *Equilibrium with commitment*

Next, we consider the problem when the planner in $t=0$ chooses policy for j periods. The equilibrium is equivalent to one in which an appropriately defined sequence of planners choose policy without commitment, so we can solve for the commitment equilibrium by applying Proposition 1.

For each period $k<j$, when policy would be chosen by the initial planner with commitment, let policy instead be chosen by a "pseudo-planner" without commitment, endowed with the sequence of time preference rates that the initial planner would want the pseudo-planner to use. The pseudo-planner at k employs time weights $\{\lambda_{k,k+m}\}_{m=0}^{\infty}$. By Corollary 1, changing from the setting without commitment to the setting with commitment—thus, changing the path

of equilibrium emissions—does not affect the equilibrium aggregate savings rate. The pseudo-planner's optimized continuation welfare cannot be higher than the corresponding continuation value if policy were chosen by the initial generation because that would contradict optimality of the initial generation's policy choice with commitment. Similarly, the continuation value when policy is chosen by the initial generation cannot be higher, because that would contradict optimality of the pseudo-planner's problem. Thus, because optimal policy is unique, the policy rules must be the same.

Given this equivalence, we can obtain the equilibrium with commitment by merely substituting the discount factors $\{\lambda_{k,k+m}\}_{m=0}^{\infty}$ into the definitions in (3.11) for the first j periods. For simplicity, we present the case where planners beyond the initial commitment interval choose policy for one period only, though due to Corollary 2 the assumption does not affect the optimal policy within the commitment window.

Proposition 2 *Suppose the planner at period $t=0$ has a commitment device for $j>1$ periods, and thus can choose policy for $v=0,\dots,j-1$. In equilibrium, savings in every period are determined by the aggregate savings rule in* (3.9)*; carbon taxes in periods $v>j-1$ are given by* (3.10) *(with $v=t$). In addition, within the commitment interval ($v=0,\dots,j-1$), equilibrium carbon taxes are given by*

$$\tau_v = \left[\sum_{k=0}^{\infty} \lambda_{v+k}(1-d_k)\gamma \cdot \tilde{\Gamma}_k^v\right] Y_v, \tag{3.12}$$

where

$$\tilde{\Gamma}_k^v = \frac{\sum_{m=0}^{\infty} \alpha^m \lambda_{v+k,v+k+m}}{\sum_{n=0}^{\infty} \alpha^n \lambda_{v+n}}.$$

The two forms of decoupling in the log-linear model simplify the equilibrium construction. Commitment has no effect on the savings rule, and it alters the tax rules only during the commitment interval. These changes in the tax rules alter emissions during these periods, so they change the climate state and the capital stock inherited at period $t+j$. Subsequent tax levels, but neither the tax rules nor the emissions levels, change after $t+j$.

Formula (3.12) implies that the tax during the first period of a commitment interval does not depend on the length of that interval. Let t be the time index at the beginning of a commitment interval of length $j \geq 1$. Denote the equilibrium tax in t given commitment period j as $\tau_t^{(j)}$.

Corollary 4 *The equilibrium tax in the first period of the commitment interval (time t) is independent of j: $\tau_t^{(1)} = \tau_t^{(j)}$ for $j \geq 1$.*

The result follows because $\tilde{\Gamma}_k^0 = \Gamma_k$. The corollary shows that time consistency is not an issue for near-term policy in the log-linear setting. Time inconsistency arises with NCTP when the agent incorrectly assumes a commitment device (often referred to as the "naive" scenario). Without commitment, the agents' decisions in the future differ from the choices earlier generations would like them to make. This difference typically changes the marginal payoffs of initial-period actions—but not in the log-linear model. This independence arises because equilibrium flow payoffs are linear in prior emissions. Consequently, commitment-induced changes in future emissions do not change the optimal emission decision at the beginning of the commitment interval. Comparison of (3.10) and (3.12) shows that subsequent tax rules during the commitment window can differ substantially.

Corollary 4 reflects a more general feature of the log-linear model: today's policy decisions are independent of future generations' time preference paths (Iverson, 2012). This feature has a

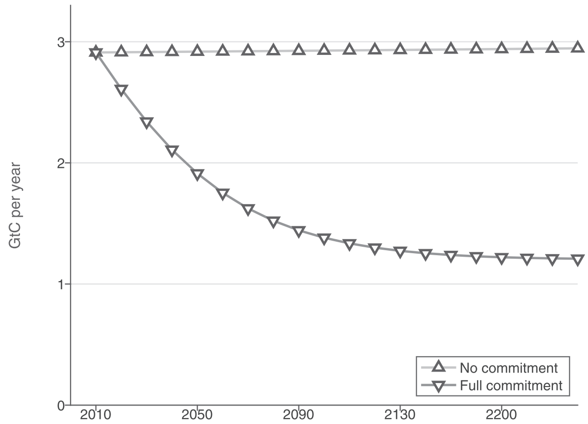

FIGURE 1

Equilibrium emissions with and without full commitment ($j=\infty$). Log-linear model with TFP growth using baseline calibration from Section 5.1.

useful implication. Because agents know less about future generations' preferences than about their own, they might entertain a wide range of beliefs about the former (Beltratti et al., 1998). In most settings, but not in the log-linear model, the initial carbon price is sensitive to these beliefs.

3.4. *The value of commitment*

This section uses Propositions 1 and 2 to study the value of commitment for the log-linear case. The quantitative results use the calibration in Section 5.1, with a 10-year time step.

To motivate the study of commitment value, Figure 1 compares the path of emissions in the equilibrium without commitment (upward-pointed triangles) with the corresponding path when the initial generation chooses climate policy in all future periods (downward-pointed triangles). The comparison shows that, absent commitment, future generations do considerably less to combat climate change than the initial generation would like.[9] The difference in emission levels suggests that the ability to commit to future policy could have substantial value for the initial generation.

We define the value of a j-period commitment technology as the increment in first-period consumption needed in the scenario without commitment ($j=1$) to raise welfare to the level with the commitment technology. To calculate this equivalent variation, we divide the difference in welfare, with and without commitment, by the initial period marginal utility of consumption in the no-commitment scenario. The permanent commitment device illustrated in Figure 1 is worth about 16 trillion USD (2.3% of the initial decade's Gross World Product). Permanent commitment would be valuable, but it is almost certainly infeasible. History is replete with examples of the difficulty of maintaining earlier commitments, even over short periods of time. The increasing difference between future generations' equilibrium emissions levels and the levels the current generation would like them to choose (Figure 1) compounds this general difficulty: more distant generations have greater incentive to break the commitment device.

9. Absent commitment, emissions are constant in the log-linear model because we assume that the TFP paths in different energy sectors are constant. If TFP grew more rapidly in the clean than in the dirty sector, emissions in the equilibrium without commitment would decline. Section 5.1, where we calibrate the model, describes these sectors.

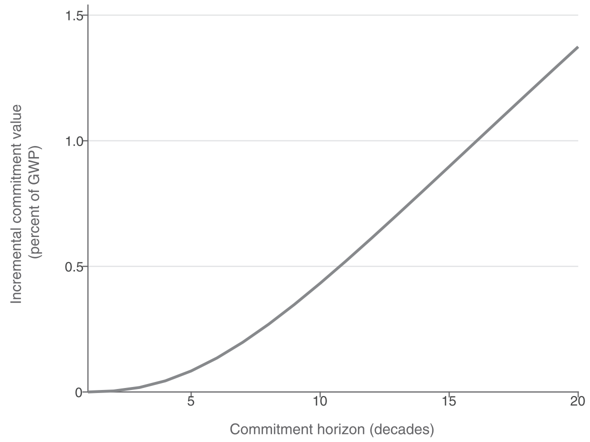

FIGURE 2
Commitment value as a function of commitment horizon. Log-linear model. Baseline calibration from Section 5.1, including base case discounting function.

Due to the difficulty of making a long-lasting commitment, we study the value of partial commitment. Figure 2 plots commitment value as a function of the commitment horizon. The slope of the curve is initially very small, and it is convex for the first 120 years. The curve shows that a small amount of commitment provides only a small welfare increase. Even a non-negligible five-decade commitment window is worth only about 0.1% of the first decade's Gross World Product, a small fraction of the value of permanent commitment. Given the difficulty of achieving long-term commitment and the small value of limited commitment, we conclude that efforts to create a commitment technology for climate policy are unwarranted.

While the commitment value curve is initially convex, it eventually becomes concave and converges to a maximum. Two offsetting forces explain this shape. We refer to them as the *decision-conflict effect* and the *present-value effect*. The decision-conflict effect stems from the difference between the discount rate that the agent in t uses to evaluate a utility transfer N periods ahead (in $t+N$) and the discount rate that the agent in period $t+N$ uses to evaluate the same transfer. When the discount rate falls over time, this difference grows in N. As a result, the degree of decision conflict between agents also grows in N. Because an added decade of commitment resolves a greater conflict, the value of commitment tends to be convex in N, not merely monotonic. The present-value effect works in the opposite direction. It arises because more serious conflicts occur further in the future and are therefore discounted more heavily. Discounting reduces the present value benefit of the resolution of a more distant conflict, and it tends to make the value-of-commitment curve concave.

Proposition 3 formalizes this intuition, showing that the value-of-commitment graph is convex in the commitment period if and only if the decision-conflict effect exceeds the present-value effect. The proposition uses the function V_{t+N}^{N}, defined as the current value (from the perspective of the agent at t) of the stream of utility from $t+N$ onward; the superscript N denotes the number of periods of commitment. The increase in this continuation value due to an additional commitment period is

$$\Delta V_{t+N}^{N} \equiv V_{t+N}^{N+1} - V_{t+N}^{N}.$$

Proposition 3 *For the log-linear model, a necessary and sufficient condition for the commitment value to be convex in the commitment horizon is that the rate of change of ΔV_{t+N}^{N} exceeds the* *discount rate that the planner at t applies between periods $N-1$ and N:*

$$\frac{\Delta V^{N+1}_{t+N+1}-\Delta V^{N}_{t+N}}{\Delta V^{N}_{t+N}} \geq \frac{\lambda_{N-1}-\lambda_N}{\lambda_N}. \tag{3.13}$$

The left side of (3.13) measures the current value of the decision-conflict effect, and the right side, equal to the period-N discount rate, measures the present-value effect.[10]

3.5. *The value of climate policy*

Even though increased intergenerational commitment produces only modest benefits, the welfare gain from climate policy (without commitment) is large.

To explain the quantitative results in this section, we introduce the discount function:[11]

$$\lambda_t = (1-w)e^{-r_{\text{short}}t} + we^{-r_{\text{long}}t}, \tag{3.14}$$

with $r_{\text{long}} \leq r_{\text{short}}$. Section 5.1 discusses the calibration. The instantaneous discount rate at $t=0$ is $wr_{\text{long}}+(1-w)r_{\text{short}}$ and the asymptotic rate as $t\to\infty$ is r_{long}. We hold r_{long} fixed and vary w, making compensating changes in r_{short} to keep the $t=0$ rate constant. A larger value of w thus lowers the discount rate at every time except $t=0$ and $t=\infty$.

Figure 3 plots the value of climate policy as a function of the commitment period, for three values of w.[12] Lower discount rates (higher w) increase the value of climate policy for any level of commitment. Policy without commitment is worth 8.5% of the first decade of Gross World Product when $w=0.13$, 15.3% when $w=0.23$, and 21.8% when $w=0.33$.[13] A larger value of w also increases the absolute value of commitment, but by a much smaller amount. Therefore, the ratio of the value of additional commitment relative to the value of policy with no commitment is small and falls with w.[14]

For our baseline calibration, climate policy without commitment reduces first-period annual emissions from the zero-tax business as usual level of 8.7 GtC/year to 2.9 GtC/year. Permanent

10. Under quasi-hyperbolic (β-δ) discounting, the right side of (3.13) equals the constant $\frac{1-\delta}{\delta}>0$. With a constant discount rate after the first period, the decision-conflict does not change over time, so the left side of (3.13) tends to be close to zero. (It need not be exactly zero because changes in technology, capital, and the climate cause the marginal utility of income to change over time.) In this case, inequality (3.13) is not satisfied: the value of commitment is concave, so a small amount of commitment provides a first-order welfare gain. The Supplementary Appendix shows that, for our calibration, the value of commitment is globally concave under β-δ discounting. This difference provides another example of a situation where models with ever-decreasing discount rates and models with quasi-hyperbolic discounting might have different policy implications. See footnote 1.

11. The simplicity and flexibility of the sum of exponentials provides the most direct justification for this functional form. In addition, Ekeland and Lazrak (2010) show that an overlapping generations model with altruistic agents and constant mortality rate produces a discount function of this form. The Supplementary Appendix discusses the modification needed to include privately owned capital to their overlapping generations model, and also explains why Section 5.1 uses a different basis for calibration. The sum of exponentials is also often used to aggregate welfare of agents with constant but different pure rates of time preference. However, that interpretation leads to a different aggregate savings rate than the one studied in this article, and we do not pursue it.

12. The value of climate policy without commitment equals the difference in welfare without commitment ($j=1$) and under zero-tax business as usual, divided by the marginal utility of consumption in the first decade under $j=1$.

13. The value of climate policy here is high relative to Nordhaus (2013) because the two models use different elasticities of intertemporal substitution and different discounting assumptions. Results in Section 5.3 show that a more conventional elasticity of substitution (0.6 instead of 1 as in the current section) brings our estimated value of climate policy much closer to previous estimates.

14. For example, if commitment increases from one decade to one century ($j=10$ versus $j=1$), this ratio is 0.04 when $w=0.13$, 0.03 when $w=0.23$, and 0.02 when $w=0.33$.

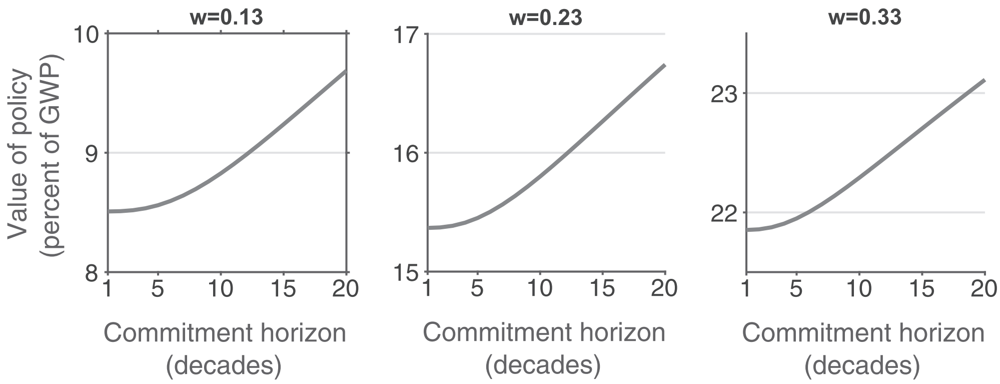

FIGURE 3

Policy value as a function of commitment horizon. Log-linear model. The three calibrations of the discount function are presented in Section 5.1. The base case has $w=0.23$.

commitment does not change current emissions, but reduces emissions in the distant future to slightly above 1 GtC/year. Thus, climate policy leads to a large reduction (relative to elasticity of intertemporal substitution) in cumulative emissions over the next several centuries, while permanent commitment leads to a relatively modest further reduction. This difference in emissions explains the large welfare difference noted above. Meanwhile, due to the convex shape of the commitment value curve, an arguably more realistic commitment period of (say) 50 years produces a negligible increase in welfare. It follows that achieving international cooperation needed to implement global climate policy is much more important than establishing a mechanism to commit policy over time.

3.6. *Limitations of the log-linear model*

The decoupling properties of the log-linear model lead to its tractability, resulting in the analytic formula for the optimal carbon tax. This tractability makes it easy to perform comparative static experiments and to examine the welfare effects of different assumptions about commitment. However, the model's strengths are closely tied to its limitations.

The chief reason to use an integrated assessment model, where savings are endogenous, is that savings and the climate are both sufficiently important macro forces that one is likely to affect the other. In decoupling the savings and climate decisions, the log-linear model breaks this connection, thus shutting down potentially important general equilibrium interactions.

In addition, the log-linear model decouples current and future climate policy. With NCTP, equilibrium policy is the outcome of a game. If current planners are unable to directly choose future policies, they might try to manipulate future planners by altering state variables such as capital or the climate stock. The log-linear model shuts down these strategic interactions, which in a more general setting might affect the value of commitment and other quantitative results.

Finally, the restricted functional form eliminates two other forces known to be quantitatively important to climate policy. First, with logarithmic utility, the income and substitution effects exactly cancel, making the carbon tax and the value of commitment invariant to exogenous growth. It is well known that, with a lower elasticity of intertemporal substitution, exogenous growth reduces the optimal carbon tax. We show below that, with growth, the lower elasticity also reduces the value of commitment. Second, the exponential damage function in Definition 3.ii, together with other parts of the definition, cause the social cost of carbon to be independent of the carbon stock. This feature means that we cannot represent the situation where damages are strongly convex, as with tipping points.

To move beyond the log-linear model, we need to be able to solve a more general model numerically. We turn to this challenge next.

4. NUMERICAL METHODS

A growing literature examines climate and macroeconomic linkages in general settings with constant discounting (Nordhaus, 2013; Lemoine and Traeger, 2014; Kelly and Tan, 2015; Hiraguchi, 2014; Lontzek et al., 2015). However, the general (non-log-linear) equilibrium growth problem with quasi-hyperbolic discounting, even for a single state variable, is unsolved. Our model accommodates arbitrary non-constant discounting, non-stationarity, three aggregate state variables and, most importantly, a coupled equilibrium problem that jointly determines competitive savings and climate policy. To solve this problem, we overcome two distinct challenges.

First, because agents in different periods disagree about the value of subsequent utility streams, each agent aggregates those streams using a different continuation value function. Standard dynamic programming methods cannot be used here because those methods involve a single continuation value function for each time period. To proceed, Proposition 4 modifies the standard finite horizon dynamic programming algorithm to track separately the distinct continuation value functions of agents in different periods.

Second, polynomial approximation methods applied to non-stationary dynamic programs suffer from instability problems even in simple settings. This instability occurs because repeated iteration of the dynamic programming equation can amplify small wobbles in the approximated value function, causing it to lose shape—even causing the approximation of a convex function to cease to be convex (Cai and Judd, 2014).

In our case, the instability is particularly severe because calculating the competitive equilibrium savings rule requires that we solve a fixed-point problem at each time step.[15] The vector of state variables for the representative agent includes her individual capital, k, and aggregate capital, K. In equilibrium, $k=K$ at every time step. However, we have to allow the agent to deviate by saving a non-equilibrium amount, causing $k \neq K$ in the next period. Therefore, we need to obtain the continuation values for all feasible (not merely for all equilibrium) values of (k,K).[16] Proposition 5 provides a simple solution by producing a semi-analytic numerical procedure that exploits the structure of the model when utility is isoelastic. The procedure uses an insight from Krusell et al. (2002), who note that the value function for agents with isoelastic utility and a linear constraint can be expressed in closed form in individual (but not aggregate) capital. Krusell et al. (2002) use this fact to calculate the steady state for a class of models. We adapt it to approximate the dynamic equilibrium for a broader class. We verify the accuracy of the approach by comparing the numerical and the analytic solutions for the log-linear case.

15. When solving the two-instrument problem, the planner effectively chooses investment in each period, so one can obtain the investment policy without solving a fixed-point problem. This simplifies the numerical challenge. Nevertheless, to *implement* the planner's choice of savings using an investment tax, we would still need the agent's savings rule, which would require solving the equilibrium problem considered here.

16. Linear splines or shape-preserving spline approximation methods (which currently have been developed only for problems with a single state variable) help with stability in some cases (Cai and Judd, 2014). These methods do not help in our setting, where the difficulty arises because of the need to solve an equilibrium problem and also equate individual and aggregate savings. Furthermore, a shape preserving procedure would not be useful because we do not know the "correct" shape, *e.g.*, whether the continuation functions are concave.

4.1. *The generalized dynamic programming equation*

Proposition 4 presents the modified dynamic programming algorithm that we use to compute the Markov equilibrium. We present the algorithm with only enough generality for our model. Iverson (2012) provides a more general version, requiring additional notation. The proposition considers the savings decision, taking climate policy as given. To solve for the full Markov equilibrium of the model, the two decisions must be considered simultaneously.

The model has $T+1$ periods. Working backwards, the iteration index i identifies the number of periods to go. Fixing the initial period at $t=0$, the iteration index also accounts for calendar time, thus incorporating non-stationarity. The agent applies the discount factor $\lambda_j \equiv \Pi_{i=0}^{j}\beta_i$—where β_j is the single period discount factor and $\beta_0=1$—to utility flows j periods in the future. In iteration i, the algorithm computes the continuation value associated with the stream of payoffs from $T-i$ through T. The challenge that arises with non-constant discounting stems from the need to account for the distinct time perspective of agents in different periods. For example, the agent in period $T-i-1$ evaluates the stream of utility from $T-i$ onwards differently than does an agent in $T-i-2$. This difference arises because the two agents apply different incremental discount rates to subsequent payoffs.

To account for these differences, we construct a series of "auxiliary value functions" to separately track the appropriate continuation value for each prior generation. For $j=i+1, i+2...T-1$, $W_{T-j}^{(i)}(k,K,\mathbf{S})$ gives the value that the agent at $T-j$ attaches to the stream of payoffs from $T-i$ through T. We need to construct as many auxiliary value functions as there are periods over which the rate of time preference declines. With quasi-hyperbolic discounting, we need one auxiliary value function, in addition to the usual value function that gives equilibrium welfare as a function of the current state variable (Harris and Laibson, 2001). Because computing the auxiliary value functions requires no additional optimization, the computational intensity is only modestly greater than for a comparable-sized dynamic program.

Proposition 4 *At iteration i the agent solves*

$$k'^{*} \equiv \operatorname{argmax}_s \left(U\left(R_i k + w_i + (1-\delta)k - k'\right) + \beta_1 W_{T-i}^{(i-1)}\left(k', K', \mathbf{S}'\right)\right). \tag{4.15}$$

Using k'^{*} to denote the equilibrium savings rule (as distinct from the level of savings), the updating equation for the auxiliary value function is

$$W_{T-j}^{(i)}(k,K,\mathbf{S}) = U\left(R_i k + w_i + (1-\delta)k - k'^{*}\right) + \beta_{j-i+1} W_{T-j}^{(i-1)}\left(k'^{*}, K', \mathbf{S}'\right) \tag{4.16}$$

for $j=i+1, i+2...T-1$ with boundary condition

$$W_{T-j}^{(-1)}(k,K,\mathbf{S}) \equiv 0, \textit{ for } j=0,1,...T-1. \tag{4.17}$$

For $i=0$ we impose the constraint $k' \geq 0$.

To apply the algorithm, we first solve the static problem in the last period. We then iterate backwards, using the appropriate auxiliary value function when evaluating the "next-period outcome" from the perspective of each earlier generation.[17] Iverson (2012) used this algorithm to

17. The discount rate corresponding to (3.14) becomes constant only asymptotically, so we would require as many continuation functions as there are periods. In implementing the algorithm, we approximate this discount function by making the discount rate constant after 30 periods (three centuries).

compute Markov equilibrium for the two-instrument problem in which the planner also controls investment. To compute the Markov equilibrium climate policy for the single-instrument problem, we use the next proposition to compute the equilibrium savings rule.

4.2. *Semi-analytic numerical procedure*

Proposition 5 uses structure from the household savings problem to derive semi-analytic expressions for the auxiliary value functions and the saving rules used in Proposition 4. After presenting the proposition, we turn to the accompanying climate policy problem, and then use plots to illustrate the procedure's stability.

Denote $R(K,E,\mathbf{S},i)$ and $w(K,E,\mathbf{S},i)$ as the equilibrium rental rate and transitory income in period i. These functions depend on the aggregate state variables and current emissions, and are solutions to the static equilibrium conditions in Definition 2.i. The representative agent, with isoelastic utility and capital k, takes the rental rate and transitory income as given and has the linear constraint $k'=R(K,E,\mathbf{S},i)k+w(K,E,\mathbf{S},i)-c+(1-\delta)k$. The auxiliary value functions are power functions (thus analytic) in own-capital, k, and the savings rules are linear in own-capital. The coefficients of the value functions and savings rules depend on aggregate state variables and current emissions, but not on own-capital. We obtain recursive formulae for those coefficient functions.

Proposition 5 *Taking as given the decision rules for climate policy, $E_i \equiv E_i(K,\mathbf{S})$, the equilibrium savings rule for the agent is linear in own-capital, $k'=s_i(k+\xi_i)+(1-\delta)k$, and the auxiliary value function has the form:*

$$W^{(i)}_{T-j}(k,K,E,\mathbf{S})=b_{i;j}(K,E,\mathbf{S})\,\frac{(k+f_i(K,E,\mathbf{S}))^{1-\eta}}{1-\eta}, \tag{4.18}$$

for $i=0,1,2..T$; for each i, the index j runs over $i+1...T-1$. At $i=0$, the agent saves nothing (so net savings equal $-(1-\delta)k$):

$$s_0=-(1-\delta) \;\; \textit{and} \;\; \xi_0=0. \tag{4.19}$$

For $i\geq 1$ the coefficients of the savings rule are

$$s_i=\frac{R_i-(1-\delta)\left(\left(\beta_1 b_{i-1;i}\right)^{-\frac{1}{\eta}}\right)}{1+\left(\beta_1 b_{i-1;i}\right)^{-\frac{1}{\eta}}} \tag{4.20}$$

$$\xi_i=\frac{w_i-\left(\beta_1 b_{i-1;i}\right)^{-\frac{1}{\eta}}f_{i-1}}{R_i-(1-\delta)\left(\left(\beta_1 b_{i-1;i}\right)^{-\frac{1}{\eta}}\right)}. \tag{4.21}$$

Appendix A.4 provides the recursions that determine the functions $b_{i;j}$ and f_i. The formulas differ in the log case, which we present in Appendix A.5.

The proposition is useful for three reasons. First, it means that we need to approximate functions that depend only on the aggregate states, not on individual capital, thus reducing the number of state variables by one. Second, it implies that we need only to solve a fixed-point problem from $\Re^2 \rightarrow \Re^2$ at each stage, instead of a more complicated mapping in the space

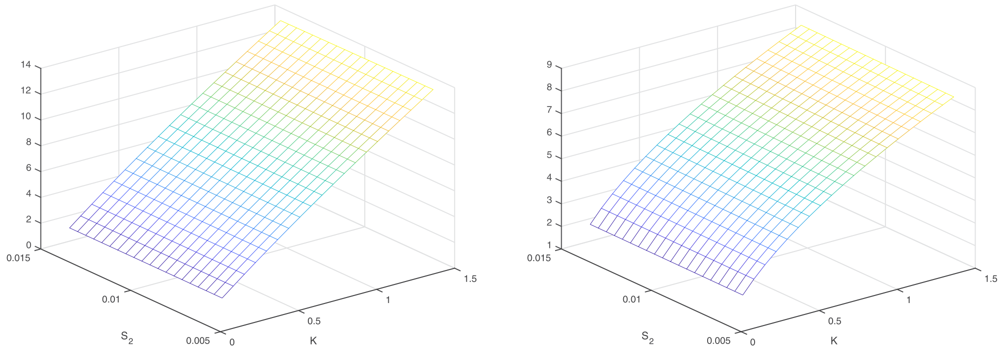

FIGURE 4

The left panel shows $f_i(K,\mathbf{S})$, while the right panel shows $b_i(K,\mathbf{S})$. Values are shown for the initial period when the baseline model is solved with a time horizon of 2,000 years.

of functions. Finally, we need only approximate the coefficient functions, which in practice are much flatter than the auxiliary value functions that they feed into (Figures 4 through 5). The flatter objects can be well-approximated with low-dimensional polynomials, increasing the stability of the procedure.

To demonstrate the stability of the approach, Figures 4 and 5 show the relationship between the coefficient functions, b_i and f_i, and the auxiliary value function, $W^{(i)}_{T-j}$. The plots use the baseline calibration and iterate for 200 time steps, corresponding to the 2,000 year time horizon in the application. Figure 4 plots $b_i(K,S)$ and $f_i(K,S)$. The state variable S corresponds to the transient stock of atmospheric carbon; the plots hold emissions and the carbon stock in the permanent reservoir fixed. Section 5.1 describes the two carbon stocks. Figure 5 plots the corresponding auxiliary value function. The plots show that the numerical procedure is highly stable, owing in part to the low degree of curvature in the b_i and f_i functions, along with the fact that the auxiliary value functions (thus, equilibrium welfare) inherit concavity directly from the analytic portion of the semi-analytic expression in equation (4.18).

Next, we turn to the determination of climate policy within each stage. The planner takes the aggregate states, $(K,\mathbf{S})$, all savings rules, and future climate policy rules as given. The equilibrium current savings rule, $s_i(K,E,\mathbf{S})[K+\xi_i(K,E,\mathbf{S})]$, depends on current emissions, via its effect on output. The planner chooses current emissions, $E=E_i(K,\mathbf{S})$, to maximize welfare for the representative agent, so the value functions and savings rules are evaluated where $k=K$. The planner's problem is

$$\max_E \frac{\left(R_i(K,\mathbf{S},E)K+w_i(K,\mathbf{S},E)-K'\right)^{1-\eta}}{1-\eta}+\beta_1 b_{i-1;i}\left(K',\mathbf{S}'\right)\frac{\left(K'+f_{i-1}(K',\mathbf{S}')\right)^{1-\eta}}{1-\eta}$$

$$\text{subject to} \qquad .$$

$$K'=s_i(K,E,\mathbf{S})(K+\xi_i(K,E,\mathbf{S}))+(1-\delta)K,$$
$$\mathbf{S}'=H(\mathbf{S},E).$$

Current emissions affect the next period states: $\mathbf{S}'$ through the impact on the stock accumulation equation—here denoted $H(\cdot)$; and K' through the impact of emissions on current savings.

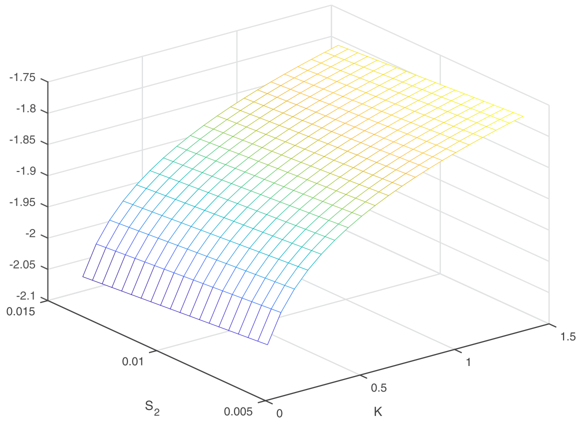

FIGURE 5
Value function W corresponding to the f_i and $b_{i;j}$ functions in Figure 4.

5. RESULTS FOR THE GENERAL MODEL

We use the methods developed above to relax the most restrictive assumptions: log-linear damages, log utility, and 100% depreciation of physical capital.[18] We first present the model and discuss the calibration. We then discuss strategic incentives in the general setting, where the decoupling properties of Section 3 do not hold. Next, we turn to the value of commitment. A final section collects sensitivity results.

5.1. *Assumptions and calibration*

We consider both the original Golosov et al. (2014) damage function—$D(\cdot)$ in Definition 3.iii—and an alternative that transforms $D(\cdot)$ using the convex function

$$\mathcal{F}(D(S)) \equiv \theta_1 D^{\theta_2} + \theta_3. \tag{5.22}$$

The transformation accommodates criticisms that mainstream models understate likely damages at high temperatures (Weitzman, 2012). For example, with the damage function in Golosov et al. (2014), a 6°C higher temperature lowers output by only 4%. To calibrate $(\theta_1, \theta_2, \theta_3)$ in (5.22), we assume damages are three times higher than this (12%) when temperature rises 6°C.[19] In addition, we equate the level and slope of $\mathcal{F}(D(S))$ with the level and slope of $D(S)$ at the current carbon stock. Figure 6 shows the calibrated damage functions. We refer to the Golosov et al. (2014) damage function as "linear" because damages are approximately linear in the stock of carbon.[20] Section 5.2 shows that the choice between linear and convex damages has an important effect on the nature of strategic interactions among generations.

18. Barrage (2013) uses numerical methods to evaluate robustness of the log-linear model to alternative functional form assumptions under constant time preference.

19. Setting climate sensitivity (the change in steady state temperature due to a doubling of atmospheric carbon) to 3°C, the increase in temperature is $T = 3.0 \ln(S/\bar{S})/\ln 2$. A temperature increase of 6°C corresponds to an increase in $S - \bar{S}$ of about 1,700 GtC.

20. This linearity is consistent with a damage function that is convex in temperature because temperature is logarithmic in carbon; the composition of the convex and concave functions is approximately linear.

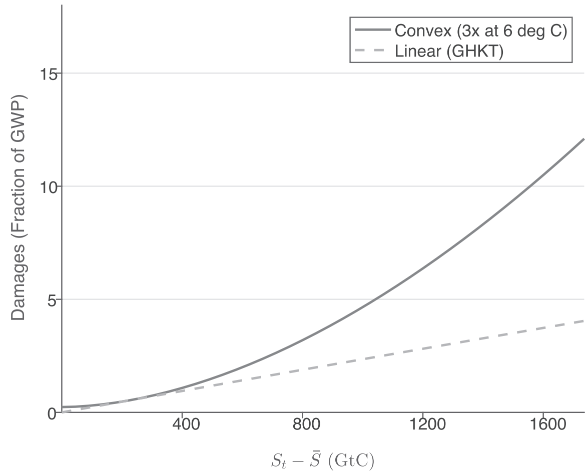

FIGURE 6

Linear climate damages follow calibration in Golosov et al. (2014); convex damages apply the calibrated convex transformation.

We replace the assumption of log utility with the isoelastic utility function in (2.1) using $\eta \neq 1$, and we allow for incomplete depreciation of physical capital: $0 < \delta \leq 1$. Our baseline uses $\eta = 1.7$ for an elasticity of intertemporal substitution of $1/\eta = 0.59$. The general model reduces to the log-linear case when $(\theta_1, \theta_2, \theta_3, \delta, \eta) = (1, 1, 0, 1, 1)$. We use this special case to validate the numerical code, verifying that solutions computed through the numerical and analytic approaches coincide almost exactly.

The model abstracts from resource scarcity, so it is sufficient to include a single composite fossil energy sector ("coal"; $i = 1$) and a single composite clean energy sector ("wind"; $i = 2$). By choice of units, emissions equal fossil fuel production: $E \equiv E_1$. Following Golosov et al. (2014), output in each energy sector i is linear in labour,

$$E_{i,t} = A_{i,t} N_{i,t},$$

and the constant elasticity of substitution energy composite equals

$$\tilde{E}_t = (\kappa_1 E_{1,t}^{\rho} + \kappa_2 E_{2,t}^{\rho})^{1/\rho}.$$

Output is Cobb Douglas in capital, labour, and energy,

$$Y_t = [1 - \digamma(D(S))] A_{0,t} K_t^{\alpha} N_{0,t}^{1-\alpha-\nu} \tilde{E}_t^{\nu}.$$

Labour is mobile across sectors with $N_{0,t} + N_{1,t} + N_{2,t} = 1$.

To calibrate the energy-sector model, we maintain the same calibration assumptions stated in pages 68–69 of Golosov *et al.* (2013). The constant elasticity of substitution energy composite function in our model differs from theirs because we assume two energy sectors (coal and renewables), while they assume three (oil, coal, and renewables). Nevertheless, by maintaining the same assumptions on the relative price of coal and wind, and the same assumptions on current energy demand, we can back out the values for κ_1 and κ_2 that are applicable in our setting.

The carbon cycle model sets current atmospheric carbon equal to the sum of carbon in permanent ($S_{1,t}$) and transient ($S_{2,t}$) reservoirs:

$$S_t = S_{1,t} + S_{2,t}.$$

TABLE 2
Baseline calibration summary

ϕ	ϕ_L	ϕ_0	α	ν	κ_1	κ_2	ρ	δ	K_0
0.0228	0.2	0.393	0.3	0.04	0.2	0.8	−0.058	0.65	164,030
$S_0(S_{1,0})$	$A_{0,0}$	$g_{A,0}$	δ_A	$A_{1,0}$	$A_{2,0}$	θ_1	θ_2	θ_3	η
802(684)	17887	0.1617	0.0616	7,693	1,311	40.6	1.82	0.00236	1.7

The fraction ϕ_L of global carbon emissions enters the permanent reservoir

$$S_{1,t} = S_{1,t-1} + \phi_L E_t.$$

The fraction ϕ_0 of remaining emissions enter the transient reservoir and decay at the rate ϕ,

$$S_{2,t} - \bar{S} = \phi(S_{2,t-1} - \bar{S}) + \phi_0(1-\phi_L)E_t.$$

Following Nordhaus (2013), we assume that TFP grows at a declining rate:

$$A_{0,t} = A_{0,t-1}(1+g_{A,t}), \quad g_{A,t} = \frac{g_{A,t-1}}{1+\delta_A},$$

where δ_A and $g_{A,t}$ are decadal rates. To simplify the interpretation of results, we hold constant the technology terms for energy production, $A_{i,t}$, $i=1,2$. We assume that capital depreciates at 10% per year, implying a decadal rate of $\delta = 0.65$. Table 2 collects baseline parameter values.

To calibrate the discount function, (3.14), we choose r_{long} using the idea that people are less able to distinguish between adjacent generations the farther away those generations are from the current time (Layton and Brown, 2000; Rubinstein, 2003). People might distinguish between the generation in 25 years and today, without distinguishing between the generations in 200 years and in 225 years. Under the assumption that people lose the capacity to distinguish between generations in the very long run, we set r_{long} equal to the pure rate of time preference in Stern (2007), who argues that there is no ethical basis for distinguishing among different generations apart from extinction risk. He calibrates the pure rate of time preference to reflect an extinction risk of $1/1,000$ per period. Using this assumption, but only in the long run, we set $r_{\text{long}} = 0.001$.

To select baseline values of r_{short} and w in (3.14), we require (as in Nordhaus 2013) that the average real return on capital over the first 50 years equals 5.5% per year when $\eta = 1.7$. Our baseline uses $r_{\text{short}} = 0.03$ and $w = 0.23$. As discussed in Section 3.4, we generate alternative paths of time preference rates by changing w and making offsetting changes in r_{short} to maintain a constant instantaneous rate at $t = 0$. Figure 7 shows these three trajectories.

The planning horizon equals 2,000 years, and we follow Barrage (2013) in assuming that coal becomes fully decarbonized after 300 years. Because carbon persists in the atmosphere, emissions during the first 300 years affect welfare over the full planning horizon.

5.2. *Strategic interactions*

Corollary 2 showed that there are no intertemporal strategic interactions in the log-linear setting. Here, we study strategic interactions in the general setting, providing intuition for the main results in Sections 5.3 and 5.4.

Figure 8 illustrates the magnitude of strategic interactions in the model with convex damages. We do not show results for the linear-damage case because the corresponding impulse response bars are approximately two orders of magnitude smaller and thus not visible on the graph if we

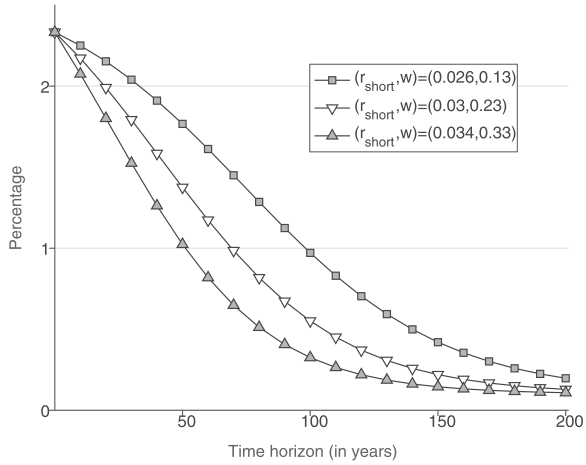

FIGURE 7

Trajectories for alternative calibrations of the discount function. r_{short} is the short-run discount rate and w is the weight on the long-run rate.

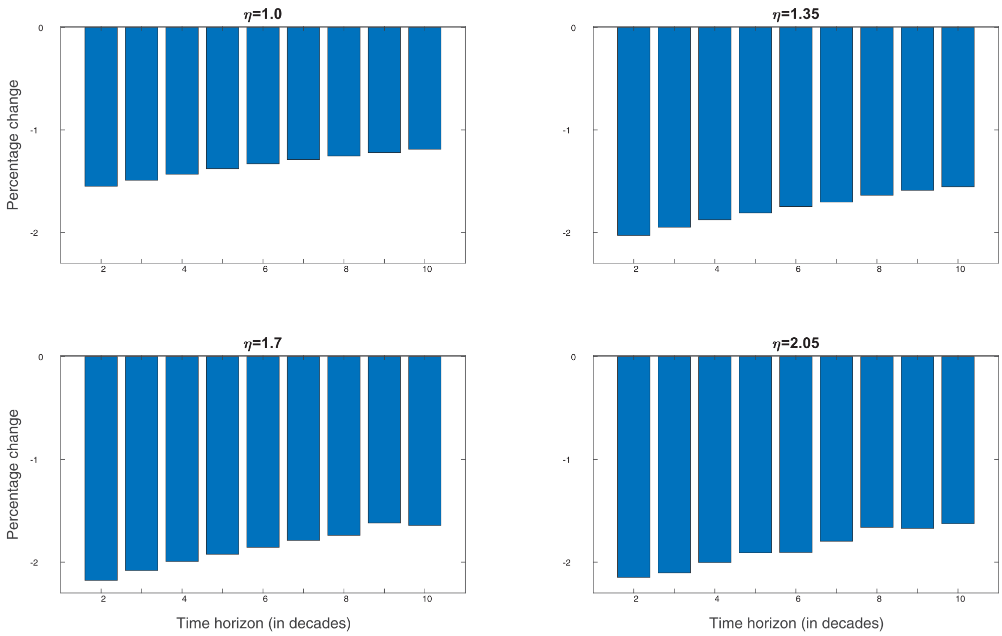

FIGURE 8

Impulse responses under convex damages following an exogenous doubling of first-period emissions. The impulse responses for the linear case are orders of magnitude smaller and therefore not shown.

keep the y-axis scale constant. The impulse response functions in Figure 8 show the equilibrium response of future generations to an exogenous doubling of current emissions. These are computed using the equilibrium policy rules for different values of η, while fixing the discounting parameters, depreciation rate and TFP growth at the baseline assumptions.

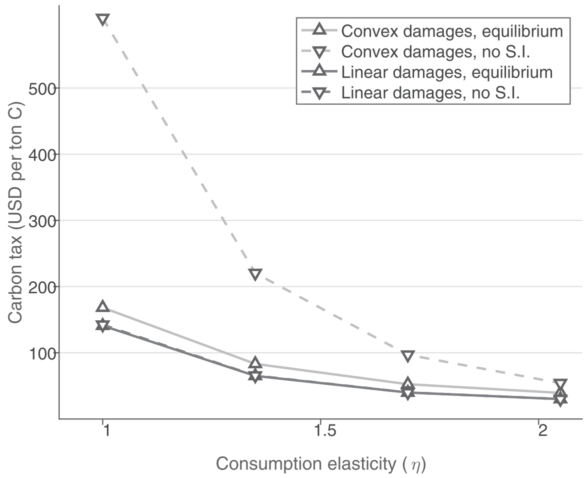

FIGURE 9

Solid and dashed black lines compare the equilibrium and restricted taxes when damages are linear. Solid and dashed grey lines compare the equilibrium and restricted taxes when damages are convex.

The results are starkly different depending on whether damages are linear or convex. With linear damages, strategic interactions remain quantitatively unimportant, even as we relax the assumptions of log utility and 100% depreciation. With convex damages, however, strategic interactions are quantitatively significant for all values of η. When $\eta = 1.7$, doubling first-period emissions reduces emissions in each of the next nine periods by roughly 2%. With convex damages, emissions are dynamic strategic substitutes: higher current emissions cause future generations to face higher marginal damages, inducing them to emit less. This strategic effect causes the Markov equilibrium carbon tax to be smaller than it otherwise would be.

To quantify the impact of strategic interactions on the optimal carbon tax, we compare the first-period optimal (Markov equilibrium) tax with the tax in a restricted problem that mechanically shuts down the strategic response of future emission decisions to changes in earlier period emission levels. This experiment enables us to compare policy with and without strategic incentives, holding fixed the convexity of damages.

Conditional on the period 1 state, $(K_1, \mathbf{S}_1)$, the restricted problem uses the equilibrium saving rate and labour allocations and the equations of motion for capital and the carbon stocks to construct the consumption trajectory, $\{\hat{C}_t\}_{t=1}^{T}$, and the restricted continuation value, $W^R(K_1, \mathbf{S}_1) = \sum_{t=1}^{T} \lambda_t u(\hat{C}_t)$. The initial decision in the restricted problem solves

$$\max_{N_0, N_1, N_2} u(C_1) + \beta_1 W^R(K_1, \mathbf{S}_1)$$

subject to the equations of motion for capital and the carbon stocks.

Figure 9 presents the results. The two black curves, corresponding to linear damages with and without strategic effects, are indistinguishable on the graph. In contrast, the grey curves, corresponding to convex damages with and without strategic effects, are far apart. This comparison shows that intergenerational strategic incentives do not meaningfully impact the carbon tax with linear damages, but they do when damages are convex. For our baseline calibration ($\eta = 1.7$) with convex damages, strategic incentives reduce the optimal tax by 45%; the reduction is much larger with log utility. Absent strategic incentives, moving from linear to convex damages leads to a large increase in the tax. But the strategic incentives that operate in the Markov equilibrium

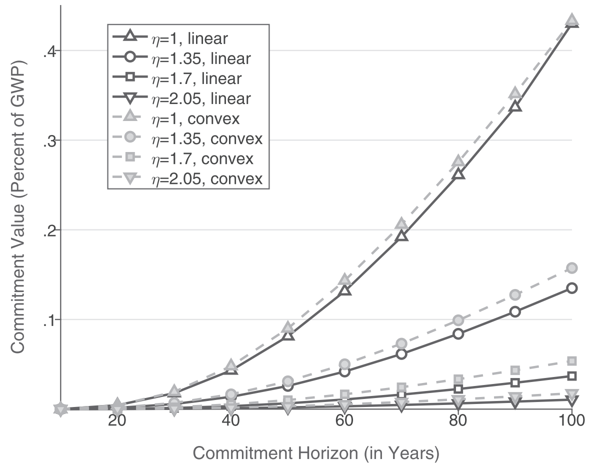

FIGURE 10

Commitment value curve for different η. Model assumes TFP growth, $\delta=0.65$, near-term discount rate of 3% and Stern weight of 0.23. Black lines show case with linear damages; grey lines show case with convex damages.

substantially offset the direct effect of higher damages that arise with convex damages. As a result, moving from linear to convex damages has only a modest effect on the Markov equilibrium tax.

5.3. *The value of commitment*

Next, we extend the analysis of commitment value to the general setting, confirming the robustness of the findings in Section 3.4. The value of commitment remains small for short commitment horizons and initially grows at an increasing rate. The intuition for the convex relationship remains approximately correct, and the incremental value of commitment is small compared to the value of climate policy without commitment.

Figure 10 generalizes Figure 2, showing the value of commitment as a function of the commitment horizon for the linear and the convex damage functions and for different values of the elasticity of intertemporal substitution. Here, we hold the discount function, depreciation rate, and TFP growth fixed at their baseline values. Solid lines indicate linear damages; dashed lines indicate convex damages. Marker shapes show different values of the elasticity of intertemporal substitution. The varying assumptions do not change the convex shape of the commitment value curve within the first century, although they can have a big effect on its magnitude.

The elasticity of intertemporal substitution is by far the most important factor influencing the value of commitment. A higher elasticity of intertemporal substitution (lower η) increases both the value of commitment and the equilibrium carbon tax (next section). The intuition for both results is essentially the same, and in the case of the tax is familiar from models with constant time preference. With TFP growth (leading to consumption growth), a higher EIS lowers the consumption discount rate and thus increases the present value of future benefits obtained from current climate policy. These higher future benefits increase the equilibrium tax and also increase the welfare gain from being able to commit to future taxes.

In comparison, the quantitative effect of moving from linear damages to convex damages is small. The strategic interactions discussed in Section 5.2 help to explain this relationship. Strategic interactions are important only when damages are convex, where emissions in different periods are dynamic strategic substitutes. Successors' equilibrium responses undercut the effectiveness

TABLE 3

The value of one century of commitment ($j=10$), as a percent of first decade of Gross World Product (first row) and the value of climate policy without commitment (second row); the third row compares these two values. Values computed for different damage scenarios and different discounting weights with $\eta=1.7$.

	Linear damages			Convex damages		
	$w=0.13$	$w=0.23$	$w=0.33$	$w=0.13$	$w=0.23$	$w=0.33$
Commitment value ($j=10$)	0.022%	0.037%	0.047%	0.034%	0.054%	0.063%
Value of climate policy	0.80%	1.4%	2.0%	2.7%	4.6%	6.5%
Relative value of commitment	2.8%	2.6%	2.3%	1.3%	1.2%	0.9%

of earlier policies. Because a commitment device shuts down strategic interactions within the commitment horizon, we might expect the presence of those interactions to increase the value of commitment. However, the strategic substitutability increases the surge of emissions in the post-commitment period.[21] The net effect of damage convexity on the value of commitment is small and the direction of the effect is ambiguous (Figure 10). As a result, the intuition for the convex shape of the commitment value curve derived in Section 3.4, which hinged on the absence of strategic interactions in the log-linear setting, remains approximately true in the general setting.

Rows 1 and 2 of Table 3 show, respectively, the incremental value of a 100-year commitment device ($j=10$) and the value of climate policy ($j=1$) (relative to elasticity of intertemporal substitution), expressed as a percent of Gross World Product during the first decade. Row 3 shows the value of commitment as a percent of the value of policy without commitment. We compute these values for both the linear- and convex-damage cases and for three values of the discounting weight, w. In all cases, the value of commitment is less than 3% of the value of the policy. These results reinforce the conclusion that search for a commitment technology to overcome decision conflict among generations with NCTP is not a first-order policy concern. Overcoming the impediments to global coordination of current climate policy, in contrast, is extremely important.

5.4. *Equilibrium carbon taxes*

Here we show how η (the inverse elasticity of intertemporal substitution), the path of time preference, and the damage function affect the equilibrium carbon tax.

Figure 11 shows the effect on the equilibrium tax of varying η. To illustrate the role of income and substitution effects, we plot the first-period tax as a function of η with TFP growth (upward-pointed triangles) and without TFP growth (downward-pointed triangles). Without TFP growth, the economy converges to a steady state where future generations are only modestly richer than the present, due to capital accumulation. In contrast, technical change makes future generations much richer.

Increased consumption in the future has no effect on the equilibrium tax when $\eta=1$ because income and substitution effects cancel. When $\eta>1$, the income effect dominates the substitution effect; here, the richer their successors, current agents are less willing to sacrifice to avoid climate-related losses in the future. In this case, TFP growth lowers the optimal carbon tax. When $\eta<1$, the substitution effect dominates the income effect, and TFP growth increases the optimal tax.[22]

21. The Supplementary Appendix includes a simulation to demonstrate the post-commitment surge in emissions after a commitment horizon of 5 and 10 years, respectively.

22. Other simulations (not included) show that relaxing the assumption of 100% depreciation has negligible effect when $\eta=1$, and a slightly larger but still small impact when $\eta>1$. The Supplementary Appendix contains additional

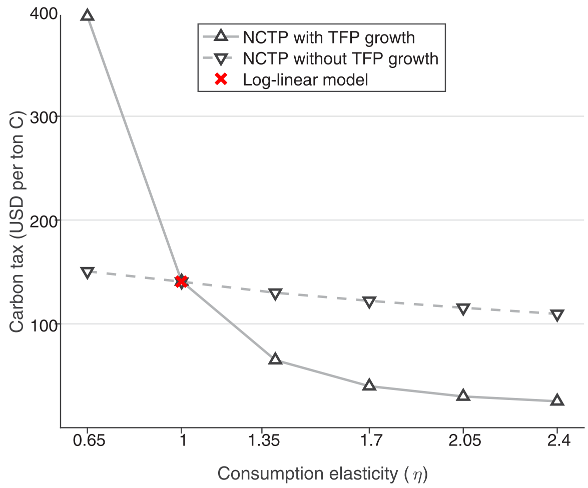

FIGURE 11

Optimal first-period carbon tax for log-linear case and alternatives. Both lines have $\delta = 0.65$, linear damages, and the base case path of time preference. The difference between downward- and upward-pointed triangles shows the effect of TFP growth.

TABLE 4

Optimal first-period carbon tax for different discounting assumptions and different damage functions. Assumes $\eta = 1.7$, $\delta = 0.65$, and TFP growth. w is the weight on the long-run discount factor.

	$w=0$	$w=0.13$	$w=0.23$	$w=0.33$
Linear damages	29	35	40	45
Convex damages	37	45	53	59

Table 4 shows the first-period Markov equilibrium tax for different values of the discounting parameter, w, and both damage functions. For each value of w, the other two discount function parameters are calibrated using the procedure described in Section 5.1. The other parameters are set at their baseline values, including $\eta = 1.7$ and $\delta = 0.65$.

Across all four values of w, the Markov equilibrium tax with convex damages is approximately 30% higher than the tax with linear damages.[23] The rough constancy of the proportional increase in the tax when moving from linear to convex damages reflects the net effect of two offsetting forces as w increases. First, higher w increases the strength of strategic interactions over time. Because emissions in different periods are dynamic strategic substitutes with convex damages, this effect causes the optimal tax to decrease. Second, higher w increases the weight on future damages, which causes the optimal tax to increase by more under convex damages.

simulation results, showing equilibrium trajectories for the savings rate, the carbon tax, and the stocks of capital and atmospheric carbon, under different assumptions about damages and the elasticity of intertemporal substitution.

23. Dietz and Stern (2015) show that moving from linear to convex damages has a much larger effect on the first period policy when climate damages reduce the growth rate, not merely the level of output.

6. CONCLUSION

We solve the first integrated assessment model with general functional forms and non-constant time preference. The setup requires that we accommodate differences in continuation values across many generations when solving for the Markov perfect equilibrium among generations and also that we solve the savings decision in a competitive equilibrium where the planner is unable to control investment (*e.g.* by means of an investment tax). These features of the model create technical difficulties that we overcome using Propositions 4 and 5.

With NCTP, the current generation has an incentive to devise commitment mechanisms to constrain their successors. Although a commitment device lasting hundreds of years would be highly valuable, one lasting only a few decades is worth little. In contrast, climate policy without commitment remains highly valuable, pointing to the importance of developing institutions that promote international cooperation.

In assuming the existence of a global planner, we ignore the possibility that intertemporal commitment might affect the equilibrium level of international cooperation. Battaglini and Harstad (2016) study a dynamic participation game in which intertemporal commitment makes it possible to solve a hold-up problem, thereby leading to a large and effective climate agreement. In their setting, intertemporal commitment promotes international cooperation. Karp and Sakamoto (2018) consider a dynamic participation game where beliefs, instead of the ability to commit, produce a broad and effective agreement. Both models assume that countries have constant discount rates. Given these differing conclusions under constant discounting, the effect of intertemporal commitment on equilibrium international cooperation is likely also model-dependent under NCTP.

In addition to solving a general Integrated Assessment Model, the article provides an analytic solution to a log-linear version. The primary appeal of the log-linear model is tractability, but it is important to ask whether the policy implications that result from it are robust. We find, consistent with results under constant discounting, that the magnitude of the optimal carbon tax is robust to changes in the depreciation rate, but not to changes in the elasticity of intertemporal substitution or the damage function. A decoupling property of the log-linear model, which shuts down strategic interactions among generations, approximately holds for other values of the EIS and of the depreciation rate, provided damages are roughly linear in the stock of carbon. But if damages are strongly convex, policies are dynamic strategic substitutes, creating strategic incentives that reduce the equilibrium carbon tax.

Some extensions would be simple, and others complex. For example, including additional state variables to capture temperature delay would be costless in the log-linear setting. Using the new algorithms, a higher-dimensional state would create the same computational issues in our problem as in standard dynamic programs, and could be remedied using the same methods (*e.g.* sparse grids). Adding a finite resource stock would require the incorporation of an additional Hotelling constraint, causing the equilibrium tax (even for the log-linear model) to no longer be independent of state variables. The resulting Hotelling problem with NCTP is more complex, and has not been solved. Incorporating uncertainty, *e.g.,* about future technology, would require fundamental changes. That extension would make it possible to quantify the relative importance of two distinct reasons for a declining consumption discount rate: hyperbolic discounting and changes to the term structure of interest rates due to uncertainty.

A. APPENDIX

The appendices provide proofs of the propositions stated in the text. We also extend Proposition 5 to include the case with logarithmic preferences. We omit the formal proof of Proposition 2 because the text above that proposition sketches the argument.

A.1. *Proof of Proposition 1 and Corollaries 1–3*

We solve for the unique Markov equilibrium that arises in the finite horizon version of the problem, then take the limit as the time horizon goes to infinity. Lemma 1 presents the equilibrium savings rule in the non-stationary, finite-horizon setting. Hiraguchi (2014) studies the analogous equilibrium problem in a stationary, infinite-horizon (one-state) economy with general NCTP. To link the savings equilibrium with the climate policy equilibrium in our setting and to solve for the limit equilibrium, we need to study the finite horizon equilibrium problem explicitly. For this reason, the proof we present differs significantly from that in Hiraguchi (2014). Nevertheless, because the result coincides with Hiraguchi (2014) in the infinite horizon limit, we relegate our proof of the lemma to the Supplementary Appendix.

We employ the following notation. The horizon at time 0 is H, so the remaining horizon at time t is $T=H-t$.

Lemma 1 *In the log-linear model (Definition 3), suppose that climate policy in each period is independent of the inherited states. Then the unique equilibrium savings rule is*

$$K'=G(K,T)Y_t=\frac{\alpha\rho(T)}{1+\rho(T)}Y_t=\frac{\rho(T)}{1+\rho(T)}r_tK \tag{A.1}$$

with the definition

$$\rho(T)\equiv\sum_{\tau=1}^{T-1}\lambda_\tau. \tag{A.2}$$

Building on (A.1), define

$$s_t\equiv\frac{\alpha\rho(H-t)}{1+\rho(H-t)}. \tag{A.3}$$

Proof. (Proposition 1)

The climate planner equilibrium is constructed using an inductive proof. In period t, the inductive hypothesis states that for all subsequent periods $\tau>t$ optimal emissions, E_τ, are independent of the inherited state variables, K_τ and $S_{\tau-1}$. In addition, from Lemma 1 the sequence of savings rules $\{s_{t+j}\}_{j=0}^{T-t}$ are stock invariant so the climate planner in t takes them as given also.

The hypothesis is easy to verify for the last period. Suppose that in a given period t, it holds for all subsequent periods. Then the planner in t anticipates capital will accumulate according to

$$K_{\tau+1}=s^\tau K_\tau^\alpha A_\tau(E_\tau)\exp(-\gamma S_\tau),\ \tau=t,\ldots,T-1. \tag{A.4}$$

Taking logs gives a first-order linear difference equation in the log of capital. Iterating this equation, the log of capital in $\tau>t+1$ can be written

$$\ln(K_\tau)=\alpha^{\tau-(t+1)}\ln(K_{t+1})+\sum_{j=0}^{\tau-t-2}\alpha^{\tau-t-2-j}[\ln(s_{t+1+j})+\ln(A_{t+1+j}(E_{t+1+j}))-\gamma S_{t+1+j}]. \tag{A.5}$$

Ignoring variables that are exogenous to the decision-maker in t, this can be written

$$\ln(K_\tau)=\alpha^{\tau-(t+1)}\ln(K_{t+1})-\sum_{j=0}^{\tau-t-2}\alpha^{\tau-t-2-j}\gamma(1-d_{1+j})E_t+\text{“terms”}. \tag{A.6}$$

Even though the decision-maker in t does not control the savings rate in t, they still influence K_{t+1} via their influence on Y_t. Using (A.6), the flow payoff in τ can be written

$$\begin{aligned}\ln(C_\tau)&=\ln[(1-s_\tau)K_\tau^\alpha A_\tau(E_\tau)\exp(-\gamma S_\tau)]\\&=\alpha\ln(K_\tau)-\gamma S_\tau+\text{“terms”}\\&=\alpha[\alpha^{\tau-(t+1)}\ln(K_{t+1})-\sum_{j=0}^{\tau-t-2}\alpha^{\tau-t-2-j}\gamma(1-d_{1+j})E_t+\text{“terms”}]-\gamma S_\tau+\text{“terms”}.\end{aligned}$$

Combining gives

$$\ln(C_\tau)=\alpha^{\tau-t}\ln(K_{t+1})-\sum_{j=0}^{\tau-t-1}\alpha^{\tau-t-1-j}\gamma(1-d_{1+j})E_t+\text{“terms”}. \tag{A.7}$$

Using this and substituting for Y_t in the planner's objective function, the planner's problem in t can be written

$$\max_{E_t} \ln\left((1-s_t)K_t^\alpha A_t(E_t)\exp\left(-\gamma\sum_{j=0}^{t+T}(1-d_j)E_{t-j}\right)\right)+$$
$$\sum_{\tau=t+1}^{T}\lambda_{\tau-t}\left[\text{“terms”}+\alpha^{\tau-t}\ln(s_tK_t^\alpha A_t(E_t)\exp\left(-\gamma\sum_{j=0}^{t+T}(1-d_j)E_{t-j}\right))-\sum_{j=0}^{\tau-t-1}\alpha^{\tau-t-1-j}\gamma(1-d_{1+j})E_t\right].$$

Letting F_t denote the final-good production function in t, the first-order condition with respect to E_t becomes

$$\frac{\partial F_t/\partial E_t}{F_t}\left(1+\sum_{\tau=t+1}^{T}\lambda_{\tau-t}\alpha^{\tau-t}\right)=\sum_{\tau=t+1}^{T}\lambda_{\tau-t}\sum_{j=0}^{\tau-t-1}\alpha^{\tau-t-1-j}\gamma(1-d_{1+j})$$
$$+\left(1+\sum_{\tau=t+1}^{T}\lambda_{\tau-t}\alpha^{\tau-t}\right)\gamma(1-d_0).$$

Combining the RHS terms gives

$$\sum_{\tau=t}^{T}\lambda_{\tau-t}\sum_{j=0}^{\tau-t}\alpha^{\tau-t-j}\gamma(1-d_j).$$

Taking the limit as $T\to\infty$, the first-order condition can be written

$$\frac{\partial F_t/\partial E_t}{F_t}(1+\theta)=\sum_{\tau=t}^{T}\lambda_{\tau-t}\sum_{j=0}^{\tau-t}\alpha^{\tau-t-j}\gamma(1-d_j). \tag{A.8}$$

Thus,

$$\partial F_t/\partial E_t=\frac{\sum_{\tau=t}^{T}\lambda_{\tau-t}\sum_{j=0}^{\tau-t}\alpha^{\tau-t-j}\gamma(1-d_j)}{1+\theta}Y_t\equiv\Lambda_t^s. \tag{A.9}$$

Standard arguments show that competitive firms will respond to a tax on emissions by setting $\partial F_t/\partial E_t$ equal to the tax so the equilibrium tax in t is Λ_t^s. Uniqueness of the equilibrium tax follows by construction.

To verify the inductive hypothesis, rewrite (A.9) as

$$A_t'(E_t)/A_t(E_t)=\frac{\sum_{\tau=t}^{T}\lambda_{\tau-t}\sum_{j=0}^{\tau-t}\alpha^{\tau-t-j}\gamma(1-d_j)}{1+\theta}. \tag{A.10}$$

This gives a deterministic function that defines E_t as a stock-invariant quantity that depends only on model parameters.

Next, we rewrite the expression for Λ_t^s to get the formula in the proposition. Specifically, for the infinite horizon case , we employ the following double sum identity:

$$\sum_{p=0}^{\infty}\sum_{q=0}^{p}a_{q,p-q}=\sum_{m=0}^{\infty}\sum_{n=0}^{\infty}a_{n,m}.$$

The identify can be proved by listing the terms in a two by two grid, where the index of the first sum comprise the rows and the index of the second sum comprise the columns. The left-hand side is obtained by summing the rows of the grid, while the right-hand side is obtained by summing the same set of terms diagonally.

When $t=0$ and $T=\infty$,

$$\sum_{\tau=t+1}^{T}\lambda_{\tau-1}\sum_{j=1}^{\tau-t}\alpha^{\tau-t-j}\gamma(1-d_j)=\sum_{p=0}^{\infty}\lambda_p\sum_{j=1}^{p+1}\alpha^{p+1-j}\gamma(1-d_j)\qquad=\sum_{p=0}^{\infty}\sum_{q=0}^{p}\lambda_{p-q+q}\alpha^{p-q}\gamma(1-d_{1+q}).$$

Letting $n=q$ and $m=p-q$,

$$
\begin{aligned}
&= \sum_{m=0}^{\infty}\sum_{n=0}^{\infty} \lambda_{m+n}\alpha^m \gamma(1-d_{1+n}) \\
&= \sum_{m=0}^{\infty}\sum_{n=0}^{\infty} \lambda_n \lambda_{n+1,n+m}\alpha^m \gamma(1-d_{1+n}) \\
&= \sum_{n=0}^{\infty}\sum_{m=0}^{\infty} \lambda_n \lambda_{n+1,n+m}\alpha^m \gamma(1-d_{1+n}) \\
&= \sum_{n=0}^{\infty} \lambda_n \gamma(1-d_{1+n}) \sum_{m=0}^{\infty} \alpha^m \lambda_{n+1,n+m} \\
&= \sum_{n=0}^{\infty} \lambda_n \gamma(1-d_{1+n}) \Gamma(n),
\end{aligned}
$$

where

$$\Gamma(n) = \sum_{m=0}^{\infty} \alpha^m \lambda_{n+1,n+m}.$$

Combining this with (A.9) gives the expression for the equilibrium tax stated in the expression. ■

Proof. (Corollary 1) This Corollary follows from Lemma 1, together with the finding that the path of stock-invariant aggregate savings rates drops out from the determination of the equilibrium tax. This step uses our assumption that the climate planner takes the equilibrium savings rule, not the level of either consumption or savings as given. ■

Proof. (Corollary 2) This Corollary follows from the inductive hypothesis, which was verified above. ■

Proof. (Corollary 3) This Corollary follows from the fact (established in the proof of the inductive hypothesis), that the continuation value consists of additively separable functions involving the capital stock, the climate variables, future climate policies, and future savings policies. The form of these functions are invariant to the savings rules. For example, conditional on the levels of capital and the climate stock, the shadow values of capital and of the climate stock are the same regardless of whether the planner uses only an emissions tax or both an emissions and an investment tax. By Corollaries 1 and 2, the climate polices are invariant to the savings rule, so the components of the continuation value involving the climate policies are the same in the two settings. Therefore, the welfare comparison depends only on a comparison of the components of the continuation value involving the savings rule. Hiraguchi's Proposition 3 shows that welfare is higher in the competitive setting, apart from set of measure 0, where the two equilibria result in the same level of welfare. ■

A.2. *Proof of Proposition 3*

Proof. (Proposition 3) Let $\{c_{t+s}^N\}_{s\geq 0}$ denote the sequence of equilibrium consumption when the commitment period is N. Generation t's present value of the continuation payoff beginning in $t+N$ is

$$W_{t+N}^N \equiv \sum_{s=0}^{\infty} \lambda_{N+s} \ln(c_{t+N+s}^N),$$

and the current value, viewed from the perspective of the generation in t, is

$$V_{t+N}^N \equiv \left(W_{t+N}^N\right)/\lambda_N. \tag{A.11}$$

The statement that the commitment value is convex in the commitment horizon means that the increment in commitment value when moving from N-period commitment to $N+1$-period commitment increases in N. Suppose commitment is imposed in period 0. With N-period commitment, the taxes in $t=0$ through $N-1$ are τ_t^N, where the "N" superscript indicates that the tax is chosen by the initial generation with N-period commitment. After period $N-1$, taxes are $\tau_t^{1,N}$, where the "1,N" superscript indicates that taxes are chosen by the contemporaneous generation without commitment in the equilibrium in which the initial generation commits policy N periods.

With $N+1$-period commitment, the taxes are $\left(\{\tau_t^{N+1}\}_{t=0}^{N-1}, \tau_N^{N+1}, \{\tau_t^{1,N+1}\}_{t=N+1}^{\infty}\right)$. It follows from Corollary 1 that $\tau_t^N = \tau_t^{N+1}$ for $t=0,\dots,N-1$ and $\tau_t^{1,N} = \tau_t^{1,N+1}$ for $t>N$. In addition, due to Corollary 2, savings rates are the same in all periods regardless of commitment horizon. It follows that consumption levels are the same for the first N periods with N and $N+1$ period commitment. Letting c_t^N denote equilibrium consumption in t with N-period commitment, we have

$$c_t^{N+1} = c_t^N \text{ for } t = 0, \ldots, N-1. \tag{A.12}$$

Define welfare for the generation in period 0 with j period commitment as $\tilde{W}_0^j$. Then the value of commitment is[24]

$$\tilde{W}_0^N - \tilde{W}_0^1.$$

Then the increment in commitment value when going from N-period commitment to $N+1$-period commitment is

$$\text{INC}_{N+1} \equiv (\tilde{W}_0^{N+1} - \tilde{W}_0^1) - (\tilde{W}_0^N - \tilde{W}_0^1) = \tilde{W}_0^{N+1} - \tilde{W}_0^N.$$

Using (A.12),

$$\begin{aligned}
\text{INC}_{N+1} &= \tilde{W}_0^{N+1} - \tilde{W}_0^N \\
&= \sum_{t=0}^{\infty} \lambda_t \ln(c_t^{N+1}) - \sum_{t=0}^{\infty} \lambda_t \ln(c_t^N) \\
&= \left(\sum_{t=0}^{N-1} \lambda_t \ln(c_t^{N+1}) + \sum_{t=N}^{\infty} \lambda_t \ln(c_t^{N+1}) \right) - \left(\sum_{t=0}^{N-1} \lambda_t \ln(c_t^N) + \sum_{t=N}^{\infty} \lambda_t \ln(c_t^N) \right) \\
&= \left(\sum_{t=0}^{N-1} \lambda_t \ln(c_t^N) + \sum_{t=N}^{\infty} \lambda_t \ln(c_t^{N+1}) \right) - \left(\sum_{t=0}^{N-1} \lambda_t \ln(c_t^N) + \sum_{t=N}^{\infty} \lambda_t \ln(c_t^N) \right) \\
&= \sum_{t=N}^{\infty} \lambda_t \ln(c_t^{N+1}) - \sum_{t=N}^{\infty} \lambda_t \ln(c_t^N) \\
&= \lambda_N \left[\sum_{s=0}^{\infty} \frac{\lambda_{N+s}}{\lambda_N} \ln(c_{N+s}^{N+1}) \right] - \lambda_N \left[\sum_{s=0}^{\infty} \frac{\lambda_{N+s}}{\lambda_N} \ln(c_{N+s}^N) \right] \\
&= \lambda_N \left[V_N^{N+1} - V_N^N \right],
\end{aligned}$$

where V_t^N is defined in (A.11).

It follows that $\text{INC}_{N+1} \geq \text{INC}_N$ if and only if

$$\frac{\Delta V_{N+1}^{N+1}}{\Delta V_N^N} \geq \frac{\lambda_{N-1}}{\lambda_N} \equiv 1 + r_N, \tag{A.13}$$

where r_N is the discount rate applied by the initial generation between period $N-1$ and period N. This condition is equivalent to (3.13). ■

A.3. *Proof of Proposition 4*

Proposition 4 generalizes the standard dynamic programming equation. Harris and Laibson (2001) use this approach for quasi-hyperbolic discounting. Fujii and Karp (2007) use a variation as a basis for numerical work for a one-state variable stationary problem with general hyperbolic discounting.

Proof. (Proposition 4) We use a proof by induction. At $i=0$, (4.17) and the fact that there are no subsequent payoffs imply that the agent solves the optimization problem in the first line of (4.15) (under the constraint $k' \geq 0$). Therefore, the decision rule is given by the second line of (4.15) for $i=0$.

The inductive hypothesis is that the first line of (4.15) holds at iterations $i-1$. We need to establish that this hypothesis and the updating equation 4.16 imply that the first line of (4.15) gives the correct optimization problem at iteration i. To establish this claim, it is necessary and sufficient to confirm that $\beta_1 W_{T-i}^{(i-1)}$ is the correct continuation value at iteration i. The rest of the proof establishes this claim.

24. For simplicity, we denominate commitment value in utility units here, though the argument is unchanged when welfare changes are denominated in consumption units.

The agent's evaluation, at iteration i, of an arbitrary sequence of flow payoffs $\left\{U^{(m)}\right\}_{m=0}^{i}$ is

$$U^{(i)}+\beta_1 U^{(i-1)}+...+(\beta_1\beta_2...\beta_i)\,U^{(0)}=\sum_{t=0}^{i}\left(\Pi_{n=0}^{t}\beta_n\right)U^{(i-t)}, \tag{A.14}$$

where the equality uses $\beta_0=1$. Denote the *equilibrium* flow payoff at iteration i as $\tilde{U}^{(i)}$. The equilibrium payoff at iteration i depends on the agent's level of capital, and current and future variables that the agent takes as exogenous (captured by the iteration index). Thus, the equilibrium values $\left\{\tilde{U}^{(m)}\right\}_{m=0}^{i}$ depend on k_i and on variables that the agent takes as exogenous.

Replacing the arbitrary functions $U^{(i)}$ with equilibrium levels of utility, $\tilde{U}^{(i)}$ and using (4.16) and (4.17) and repeated substitution, we write

$$\begin{aligned}
W_{T-j}^{(i)}&=\tilde{U}^{(i)}+\beta_{j-i+1}W_{T-j}^{(i-1)}\\
&=\tilde{U}^{(i)}+\beta_{j-i+1}\left[\tilde{U}^{(i-1)}+\beta_{j-i+2}W_{T-j}^{(i-2)}\right]\\
&=\tilde{U}^{(i)}+\beta_{j-i+1}\left[\tilde{U}^{(i-1)}+\beta_{j-i+2}\left(\tilde{U}^{(i-2)}+\beta_{j-i+3}W_{T-j}^{(i-3)}\right)\right]\\
&\;\;\vdots\\
&\;\;\vdots\\
&=\tilde{U}^{(i)}+\textstyle\sum_{t=1}^{i}\left(\Pi_{n=0}^{t-1}\beta_{j-i+1+n}\right)\tilde{U}^{(i-t)}.
\end{aligned}$$

Thus,

$$W_{T-j}^{(i-1)}=\tilde{U}^{(i-1)}+\textstyle\sum_{t=1}^{i-1}\left(\Pi_{n=0}^{t-1}\beta_{j-i+2+n}\right)\tilde{U}^{(i-1-t).}$$

Setting $j=i$ (the largest value of j when the superscript is $i-1$) gives

$$\begin{aligned}
W_{T-i}^{(i-1)}(k_{i-1})&=\tilde{U}^{(i-1)}+\textstyle\sum_{t=1}^{i-1}\left(\Pi_{n=0}^{t-1}\beta_{2+n}\right)\tilde{U}^{(i-1-t)}\Rightarrow\\
\beta_1 W_{T-i}^{(i-1)}(k_{i-1})&=\beta_1\left[\tilde{U}^{(i-1)}+\textstyle\sum_{t=1}^{i-1}\left(\Pi_{n=0}^{t-1}\beta_{2+n}\right)\tilde{U}^{(i-1-t)}\right]\\
&=\textstyle\sum_{t=1}^{i-1}\left(\Pi_{n=1}^{t}\beta_n\right)\tilde{U}^{(i-t)}.
\end{aligned} \tag{A.15}$$

Here, we include the argument k_{i-1}, the agent's stock of capital at iteration $i-1$, in the function $W_{T-i}^{(i-1)}$ in order to emphasize that the equilibrium sequence of current and future flow payoffs depend on this stock.

At iteration i, using (A.14), the agent chooses k' to solve

$$\begin{aligned}
&\max_{k'}\left(U^{(i)}\left(k,k'\right)+\beta_1\textstyle\sum_{t=1}^{i}\left(\Pi_{n=0}^{t}\beta_n\right)\tilde{U}^{(i-t)}\right)=\\
&\max_{k'}\left(U^{(i)}\left(k,k'\right)+\beta_1 W_{T-i}^{(i-1)}(k_{i-1})\right),
\end{aligned}$$

where the equality uses (A.15). This equation establishes that the first line of (4.15) gives the correct maximand at iteration i. The agent in the current period can deviate from equilibrium, so the current flow payoff is $U^{(i)}\left(k,k'\right)$. However, the sequence of future flow payoffs, $\tilde{U}^{(i-t)}$, are evaluated in equilibrium; these payoffs are functions of k'.) ■

A.4. *Proof of Proposition 5*

Before sketching the proof, we provide the recursions that determine the endogenous functions $b_{i;j}(K,E,\mathbf{S})$ (for $j=i+1,...T-1$) and $f_i(K,E,\mathbf{S})$.

$$\begin{aligned}
b_{i;j}&=\left(\frac{R_i+(1-\delta)}{1+(\beta_1 b_{i-1;i})^{-\frac{1}{\eta}}}\right)^{1-\eta}\left[\left(\beta_1 b_{i-1;i}\right)^{-\frac{1-\eta}{\eta}}+\beta_{j-i+1}b_{i-1;j}\right]\\
f_i&=\frac{w_i+f_{i-1}}{R_i}
\end{aligned} \tag{A.16}$$

with boundary conditions

$$f_0=\frac{w_0}{R_0+1-\delta}\text{ and }b_{0;j}=(R_0+1-\delta)^{1-\eta}. \tag{A.17}$$

We use Proposition 4 to develop formulae for the endogenous savings rule, a linear savings rule whose coefficients depend on variables that the agent takes as exogenous. Those variable might be either exogenous or endogenous to the model (*e.g.* technological change versus climate change). Single-period utility is

$$U(c) = \frac{c^{1-\eta}}{1-\eta} \text{ for } \eta \neq 1; \; U(c) = \ln(c) \text{ o/w} \Rightarrow U'(c) = c^{-\eta}.$$

A linear net savings rule, $s_i(k+\xi_i)$ implies next period capital for the agent is

$$k' = s_i(k+\xi_i) + (1-\delta)k. \tag{A.18}$$

Consumption is

$$c_i = (R_i - s_i)k + w_i - s_i\xi_i. \tag{A.19}$$

Utility under this linear savings rule equals

$$\frac{((R_i-s_i)k+w_i-s_i\xi_i)^{1-\eta}}{1-\eta} = \frac{(R_i-s_i)^{1-\eta}\left(k+\frac{w_i-s_i\xi_i}{R_i-s_i}\right)^{1-\eta}}{1-\eta} \tag{A.20}$$

The subscript i recognizes that the net savings rate changes with the iteration number. That change potentially arises for several reasons: changes in variables that are exogenous to the model (*e.g.* technology), changes in state variables that are exogenous to the agent but endogenous to the model (*e.g.* climate and climate policy) and also because (for finite terminal time, T) the distance from the current to the final period changes with the iteration index.

For the climate problem, it is important to include the climate state and climate policies as arguments of the endogenous functions, but in the interest of clarity we begin with a simpler problem in which the only exogenous states (from the agent's perspective) are aggregate capital, K, and the index i. In this case, $R_i = R(K,i)$ and $w_i = w(K,i)$ are known functions, depending on aggregate capital stock and i (to capture exogenous-to-the-model changes).

Proof. (Sketch of Proposition 5) The proof is algebra-intensive, so we relegate the details to the Supplementary Appendix. Here, we describe the straightforward logic. We use an inductive proof. In the last period ($i=0$) savings equal zero. With this fact, it is easy to confirm (A.17) and, for $i=0$, (4.18). We then use the inductive hypothesis (equation 4.18 holds for $i-1$) to write the agent's saving problem at iteration i. The first order condition to this problem implies the linear savings rule, with s_i given by (4.20) and ξ_i given by (4.21). The agent's problem is concave iff $b_{i-1;i} \geq 0$; we confirm this inequality numerically. We substitute this savings rule into the agent's dynamic programming equation to establish the recursions in the two lines of (A.16). ■

A.5. *Logarithmic utility*

Using the linear savings and consumption rules, equations A.18 and A.19, utility under logarithmic preferences equals

$$\begin{aligned} \ln((R_i-s_i)k+w_i-s_i\xi_i) &= \ln\left((R_i-s_i)\left[k+\tfrac{w_i-s_i\xi_i}{(R_i-s_i)}\right]\right) \\ &= \ln(R_i-s_i) + \ln\left(k+\tfrac{w_i-s_i\xi_i}{(R_i-s_i)}\right). \end{aligned} \tag{A.21}$$

Proposition 6 *The auxiliary value functions have the form*

$$W^{(i)}_{T-j}(k,K) = a_{i;j} + b_{i;j}\ln(k+f_i). \tag{A.22}$$

with coefficients

$$\begin{aligned} a_{i;j} &= \ln(R_i - s_i) + \beta_{j-i+1}\left[a_{i-1;j} + b_{i-1;j}\ln(s_i+1-\delta)\right] \\ b_{i;j} &= \left(1+\beta_{j-i+1}b_{i-1;j}\right) \text{ and } f_i = \tfrac{w_i+f_{i-1}}{R_i+(1-\delta)} \end{aligned} \tag{A.23}$$

and boundary conditions

$$\begin{aligned} a_{0;j} &= \ln(R_0+1-\delta) \\ b_{0;j} &= 1 \text{ and } f_0 = \tfrac{w_0}{R_0+1-\delta}. \end{aligned} \tag{A.24}$$

At $i=0$ gross savings equal zero; for $i \geq 1$ the equilibrium savings rule is

$$k' = s_i(k+\xi_i) + (1-\delta)k \tag{A.25}$$

with (for $i \geq 1$)

$$s_i = \frac{(\beta_1 b_{i-1;i}R_i) - (1-\delta)}{(1+\beta_1 b_{i-1;i})} \tag{A.26}$$

$$\xi_i = \frac{w_i\beta_1 b_{i-1;i} - f_{i-1}}{(\beta_1 b_{i-1;i}R_i) - (1-\delta)} \tag{A.27}$$

The proof parallels the case with $\eta \neq 1$, and is available upon request.

Acknowledgments. We thank David Antoff, William Brock, Reyer Gerlagh, Christian Gollier, John Hassler, Jonathan Hughes, David Kelly, Derek Lemoine, Matti Liski, Antony Millner, Salvador Ortigueira, Armon Rezai, Hiroaki Sakamoto, Christian Traeger, Sammy Zahran, and four referees for discussions and comments. An early version of the paper was circulated under the title "Optimal Carbon Taxes with Non-Constant Time Preference."

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

REFERENCES

AUGENBLICK, N., NIEDERLE, M. and SPRENGER, C. (2015), "Working Over Time: Dynamic Inconsistency in Real Effort Tasks", *Quarterly Journal of Economics*, **56**, 1067–1115.

BARRAGE, L. (2013), "Sensitivity Analysis for Golosov, Hassler, Krusell, and Tsyvinski (2013): Optimal Taxes on Fossil Fuel in General Equilibrium" (Technical Notes).

BARRO, R. (1999), "Ramsey Meets Laibson in the Neoclassical Growth Model", *The Quarterly Journal of Economics*, **114**, 1125–1152.

BATTAGLINI, M. and HARSTAD, B. (2016), "Participation and Duration of Environmental Agreements", *Journal of Political Economy*, **124**, 160–204.

BELTRATTI, A., CHICHILNISKY, G. and Heal, G. (1998), "Preservation, Uncertain Future Preferences and Irreversibility", in Chichilnisky, G. and Heal, G., and Vercelli, A. (eds) *Sustainability: Dynamics and Uncertainty* (Kluwer Academic Publishers).

BROCAS, I., Carrillo, J. and Dewatripont, M. (2004), "Commitment Devices under Self-control Problems: An Overview", *The Psychology of Economic Decisions*, **2**, 49–67.

BROCK, W. and Mirman, L. (1972), "Optimal Economic Growth and Uncertainty: The Discounted Case", *Journal of Economic Theory*, **4**, 479–513.

BRYAN, G., KARLAN, D. and NELSON, S. (2010), "Commitment Devices", *Annual Review of Economics*, 2, 671–698.

CAI, Y. and JUDD, K. (2014), "Advances in Numerical Dynamic Programming and New Applications", *Handbook of Computational Economics*, **3**, 479–516.

CROPPER, N., AYDEDE, S. and PORTNEY (1994), "Preferences for Life Saving Programs: How the Public Discounts Time and Age", *Journal of Risk and Uncertainty*, **8**, 243–265.

DIETZ, S., GOLLIER, C. and KESSLERI, L. (2018), "The Climate Beta", *Journal of Environmental Economics and Management*, **87**, 258–274.

DIETZ, S. and STERN, N. (2015), "Endogenous Growth, Convexity of Damage and Climate Risk: How Nordhaus' Framework Supports Deep Cuts in Carbon Emissions", *The Economic Journal*, **125**, 574–620.

DRUPP, M., FREEMAND, M., GROOM, B. and NESJE, F. (2014), "Discounting Disentangled: An Expert Survey on the Determinants of the Long Term Social Discount Rate" (London School of Economics Working Paper).

EKELAND, I. and LAZRAK, A. (2010), "The Golden Rule When Preferences Are Time Inconsistent", *Mathematical and Financial Economics*, **4**(1), 29–55.

FANG, H. and SILVERMAN, D. (2009), "Time-Inconsistency and Welfare Program Participation: Evidence from the NLSY", *International Economic Review*, **50**, 1043–1077.

FANG, H. and WANG, Y. (2015), "Estimating Dynamic Discrete Choice Models with Hyperbolic Discounting, with an Application to Mammography Decisions", *International Economic Review*, **56**, 565–596.

FREDERICK, S., LOEWENSTEIN, G. and O'DONOGHUE, T. (2002), "Time Discounting and Time Preference: A Critical Review", *Journal of Economic Literature*, **40**, 351–401.

FUJII, T. and Karp, L. (2008), "Numerical Analysis of Non-constant Pure Rate of Time Preference: A Model of Climate Policy", *Journal of Environmental Economics and Management*, **56**, 83–101.

GERLAGH, R. and LISKI, M. (2017), "Consistent Climate Policies", *Journal of the European Economic Association*, **16**, 1–44.

GIGLIO, S., MAGGIORI, M., STROEBEL, J. et al. (2015), "Climate Change and Long Run Discount Rates: Evidence from Real Estate" (NBER Working Paper 21767).

GOLOSOV, M., HASSLER, J., KRUSELL, P. et al. (2014), "Optimal Taxes on Fossil Fuel in General Equilibrium", *Econometrica*, **82**, 41–88.

HARRIS, C. and LAIBSON, D. (2001), "Dynamic Choices of Hyperbolic Consumers", *Econometrica*, **69**, 935–957.

HARSTAD, B. (2020), "Technology and Time Consistency", *Journal of Political Economy*, **128**, 2653–2689.

HIRAGUCHI, R. (2014), "On the Neoclassical Growth Model with Non-constant Discounting", *Economics Letters*, **125**, 175–178.

IVERSON, T. (2012), "Optimal Carbon Taxes with Non-constant Time Preference" (MPRA Working Paper No. 43264).

IVERSON, T., DENNING, S. and ZAHRAN, S. (2015), "When the Long Run Matters", *Climatic Change*, **129**, 57–72.

KARP, L. (2005), "Global Warming and Hyperbolic Discounting", *Journal of Public Economics*, **89**, 261–282.

KARP, L. (2017), "Provision of a Public Good with Multiple Dynasties", *The Economic Journal*, **127**, 2641–2664.

KARP, L. and SAKAMOTO, H. (2018), "Sober Optimism and the Formation of International Environmental Agreements" (Working Paper).

KARP, L. and TSUR, Y. (2011), "Time Perspective and Climate Change Policy", *Journal of Environmental Economics and Management*, **62**, 1–14.

KELLY, D. and TAN, Z. (2015), "Learning and Climate Feedbacks: Optimal Climate Insurance and Fat Tails", *Journal of Environmental Economics and Management*, **72**, 98–122.

KOOPMANS, T. (1960), "Stationary Ordinal Utility and Impatience", *Econometrica*, 287–309.

KRUSELL, P., KURUŞÇU, B. and SMITH, A. (2002), "Equilibrium Welfare and Government Policy with Quasi-geometric Discounting", *Journal of Economic Theory*, **105**, 42–72.

KRUSELL, P. and SMITH, A. (2003), "Consumption-savings Decisions with Quasi-geometric Discounting", *Econometrica*, **71**, 365–375.

LAIBSON, D. (1997), "Golden Eggs and Hyperbolic Discounting", *The Quarterly Journal of Economics*, **112**, 443–478.

LAIBSON, D., MAXTED, P., REPETTO, A. and TOBACMAN, J. (2017), "Estimating Discount Functions with Consumption Choices over the Lifecycle" (Mimeo, Harvard University).

LAYTON, D. and LEVINE, R. (2003), "How Much Does the Far Future Matter? A Hierarchical Bayesian Analysis of the Public's Willingness to Mitigate Ecological Impacts of Climate Change", *Journal of the American Statistical Association*, **98**, 533–544.

LAYTON, D. F. and Brown, G. (2000), "Heterogeneous Preferences Regarding Global Climate Change", *Review of Economics and Statistics*, **82**, 616–624.

LEMOINE, D. and TRAEGER, C. (2014), "Watch Your Step: Optimal Policy in a Tipping Climate", *American Economic Journal: Economic Policy*, **6**, 1–31.

LONTZEK, T., CAI, Y., JUDD, K. et al. (2015), "Stochastic Integrated Assessment of Climate Tipping Points Indicates the Need for Strict Climate Policy", *Nature Climate Change*, **5**, 441–444.

MAHAJAN, A. and TAROZZI, A. (2012), "Time Inconsistency, Expectations and Technology Adoption: The case of Insecticide Treated Nets" (Duke University Working Paper).

MONTIEL OLEA, J. and STRZALECKI, T. (2014), "Axiomatization and Measurement of Quasi-Hyperbolic Discounting", *The Quarterly Journal of Economics*, **129**, 1449–1499.

NORDHAUS, W. (2013), *The Climate Casino: Risk, Uncertainty, and Economics for a Warming World* (Yale University Press).

PASERMAN, M. (2008), "Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation", *Economic Journal*, **118**, 1418–52.

PHELPS, E. and POLLAK, R. (1968), "On Second-best National Saving and Game-Equilibrium Growth", *The Review of Economic Studies*, **35**, 185–199.

REZAI, A. and VAN DER PLOEG, F. (2017), "Climate Policy with Declining Discount Rates in a Multi-region World: Back of the Envelope Calculations" (Working Paper).

RUBINSTEIN, A. (2003), "Economics and Psychology? The Case of Hyperbolic Discounting", *International Economic Review*, **44**, 1207–1216.

SAMUELSON, P. (1937), "A Note on Measurement of Utility", *The Review of Economic Studies*, **4**, 155–161.

SHUI, H. and AUSEBEL, L. (2005), "Time Inconsistency in the Credit Card Market" (Mimeo, University of Maryland).

STERN, N. (2007), *The Economics of Climate Change: The Stern Review* (Cambridge, UK: Cambridge University Press).

STROTZ, R. (1955), "Myopia and Inconsistency in Dynamic Utility Maximization", *The Review of Economic Studies*, **23**, 165–180.

TRAEGER, C. (2014), "Why Uncertainty Matters - Discounting under Intertemporal Risk Aversion and Ambiguity", *Economic Theory*, **56**, 627–664.

———. (2016): "ACE—An Analytic Climate Economy (with Temperature and Uncertainty)" (Working Paper).

WEITZMAN, M. (2012), "GHG Targets as Insurance Against Catastrophic Climate Damages", *Journal of Public Economic Theory*, **14**, 221–244.