Lund University
Department of Economics

Modeling Time-inconsistent Climate Policy

Master Essay

May 24, 2010

Author:
Claes Ek

Supervisor:
Tommy Andersson

Abstract

This paper examines the degree to which dynamic inconsistency in climate policy presents a serious challenge in the sense that exclusively pursuing near-term cost-efficiency implies strong disincentives for policy makers to attain long-term abatement goals. We model technological change and cost-efficient emissions cuts of roughly 80% in the EU-27 energy sector during the forty-year period 2010-2050. Our model features two broad cases, as we contrast 'standard' long-term policy optimization with more short-term, sequential optimization (i.e. across subsections of the 2010-2050 period). We find that changing the policy time frame in this manner has the unfortunate effect of working as a strong disincentive for both long-term and short-term targets, perhaps even making them politically—though not theoretically—impossible to attain under most scenarios. One conclusion is that although policy makers may now commit to ambitious emissions reductions deemed realistic based on a long-term analysis, unless targets and policies are sufficiently 'locked in' the credibility of those commitments may be doubted. Another is that the idea that climate policy should be cost-efficient in the near term is inherently contradictory, because a decision maker who truly cares only about the near term has little incentive to undertake climate policy to begin with.

Keywords: climate change, dynamic consistency, cost-efficiency, energy, long-term targets

Contents

List of Figures

List of Tables

1 Introduction

This paper deals with the problem of finding an optimal policy response to climate change; specifically, we examine by means of a simple model of the EU-27 energy sector the implications of *time-inconsistent* climate policies.

What is meant in this case by time, or dynamic consistency? At its core, it is the requirement that decision makers should never deviate from previously made plans which were deemed optimal (even considering the future) at the time. On the level of the individual, microeconomists have long noted that dynamic inconsistency may arise due to 'myopia' (Strotz, 1956) such as in the case of addiction or the procrastination of unpleasant tasks (Asheim, 1997). There is also experimental evidence of dynamic inconsistency in the context of risk and repeated gambles, where participants may reverse their plans mid-gamble (such as in Johnson and Busemeyer, 2001).

Within macroeconomics, the best-known example of time inconsistency is optimal monetary policy, following the seminal papers by Kydland and Prescott (1977) and Barro and Gordon (1983). These argued that monetary policy may be understood as a repeated game between policy makers and agents in the wider economy. In short, because monetary policy is only effective as a means of increasing output if its use is not expected by economic agents, policy makers may be tempted to unexpectedly deviate from previously announced monetary policy rules so as to generate surprise inflation and thereby increase output; and such a strategy may be suboptimal if economic agents have forward-looking expectations and thus 'expect the unexpected'.

The modeling exercise of this paper, however, includes no strategic interaction between policy makers and economic agents, but rather treats the EU-27 energy-economy as essentially a deterministic system whose state is contingent only upon policies made. Thus, we examine a kind of time consistency which is somewhat different from the specific macroeconomic application just mentioned, yet nonetheless holds true to the principle of the concept.

In our case, the basic idea is as follows. Suppose a social planner (most likely, government) commits to ambitious climate policy such as reducing domestic greenhouse gas (GHG) emissions by 80% until 2050, but then opts to implement this policy in a stepwise manner: rather than sticking to what is optimal (i.e. least-cost) policy for the entire period up until 2050, policy is sequentially optimized over, say, one political term or commitment period of a Kyoto Protocol-like treaty at a time, each time with a different portion of the ultimate 80% objective in mind. Although it is the intention of policy makers to meet abatement targets, there is no policy 'lock-in'; thus, we assume that governments may procrastinate, deviating from policy commitments, if meeting them is not sufficiently politically attractive in the near term. Qualitatively speaking, it may be realistic to assume that this is how climate policy, and politics in general, works in practice: for instance, if the outcome of the next democratic election is uncertain, an incumbent administration may not find it rational to implement costly long-term policies producing benefits only during subsequent terms. This is the specific flavor of dynamic (in)consistency under scrutiny in our model.

The question is what effect such an approach might have on the feasibility of actually meeting long-term targets. In a 2005 paper, Sandén and Azar make the case that long-term emission abatement targets require that focus be put not only on using carbon pricing policies for stimulating the adoption of existing low-carbon technologies in the near term, but also on more long-term policies for technological development. Under a short-term policy time frame, excessive emphasis might initially be placed on the former; specifically, they warn that

> without these new technologies, stricter emission reduction targets may be considered impossible to meet by the government, industry and the general public, and therefore not adopted... There is a risk that the society in its quest for cost-efficiency in meeting near-term emissions targets, becomes blindfolded when it comes to the more difficult, but equally important issue of bringing more advanced technologies to the shelf (Sandén and Azar, 2005, pg. 1).

Thus, although policy makers may now commit to ambitious emissions reductions deemed realistic based on a long-term analysis, if climate policy is time-inconsistent the credibility of those commitments may be doubted.

In the spirit of Sandén and Azar (2005), our model separates 'green' technologies into low-carbon technologies offering modest short-run emissions reductions at low cost, and no-carbon technologies, which come at a significant initial cost but whose ultimate deployment are the only means of attaining long-term abatement targets. However, due to the structure of our model, our results generally do not capture the mechanism mentioned by Sandén and Azar (2005). Rather, they suggest that shortening the time frame within which climate policy is considered has the unfortunate effect of working as an overall disincentive for dealing with climate change, adversely affecting not only long-term policies but those of the short term as well.

Our model employs a cost-efficiency framework insofar as it treats a certain targeted terminal state (in our case, 80% reductions) as given, and then proceeds to find the cheapest path towards meeting it. It may be argued that this amounts to skipping the vital first step of deciding which degree of abatement, as well as the timing of it, would actually be appropriate. Doing so, one might argue,

> requires the comparison of the benefits and the costs of moving the global economy to some agreed-upon preferred state, by comparing the costs of the current state or trajectory (or business-as-usual) against the benefits and costs of specific actions designed to reach the preferred state (Dore, 2009, pg. 2).

If extended to an entire portfolio of abatement targets or 'preferred states', such an approach might provide a more exhaustive and seemingly precise account of the economics of climate change. The class of so-called Integrated Assessment Models (IAM) such as the Dynamic Integrated Climate-Economy (DICE) model of Nordhaus (first presented in Nordhaus, 1992) rely on this cost-benefit approach to climate policy; based on the trade-off between abatement costs and a damage function describing the impacts of unchecked climate change, they tend to indicate as optimal various scenarios featuring a 'policy ramp' in which abatement starts off slow and gradually accelerates (for an overview and critique of the integrated assessment literature, see Ackerman et al., 2009).

Why then have we decided in effect to disregard such broad cost-benefit analysis? After all, deciding *a priori* that climate change is a problem which demands decisive (or otherwise) action may appear as a foregone conclusion. We believe that the recent argument made by Weitzman (2009), however, suggests that our second-best approach may in practice be the most feasible one.

In his 'Dismal Theorem', Weitzman makes the case that the 'fat tail' of negative climate change impacts tends to dominate any cost-benefit analysis. That is, if there is a worst-case outcome implying near-zero (fatal) consumption levels worldwide, and we have reason to believe that the probability of that outcome is nonzero (though it may be extremely small), then avoiding it becomes an overriding priority regardless of the characteristics of more likely impacts and/or outcomes. Thus, the Dismal Theorem is a kind of formalized version of the precautionary principle.

Weitzman then argues that the very reason why integrated assessment models are able to draw exact conclusions regarding the optimal speed and scale of abatement is precisely because they tend to exclude the full implications of catastrophic low-probability, high-impact events from their damage functions. But because IAM conclusions are based on a false premise (extreme climate change is highly unlikely and can thus be disregarded, rather than placed at the center of analysis), we may have little choice but to adopt abatement targets based on some decision criterion other than (a narrow) application of cost-benefit analysis. This is why for simplicity we treat our admittedly arbitrary abatement targets as given.

1.1 Aim, methodology, and structure

The aim of this paper is to examine the degree to which time-inconsistency in climate policy presents a serious challenge in the sense that exclusively pursuing near-term cost-efficiency implies strong disincentives for policy makers to attain long-term abatement goals. This is done by means of a simple energy model: we model technological change and cost-efficient emissions cuts of roughly 80% in the EU-27 energy sector during the forty-year period 2010-2050. Our model features two broad cases, as we contrast 'standard' long-term policy optimization with more short-term, sequential optimization (i.e. across subsections of the 2010-2050 period). We explore several such sub-period durations; thus, our approach allows us to isolate various 'degrees' of time inconsistency in climate policy.

The remainder of this paper is structured in the following way. Section 2 outlines the energy-sector model used and describes its long-term and short-term formulations corresponding to different policy time frames. Section 3 presents the results of a baseline scenario, and shows how short-term cost-minimization produces disincentives for meeting our assumed abatement target. Section 4 provides sensitivity analysis for our dynamic consistency results by repeating the analysis of Section 3 under a number of alternative scenario formulations. Section 5 discusses model mechanics and assumptions, presents some policy recommendations, and concludes the paper.

2 The model

As stated, our model describes technological change and cost-efficient emissions cuts in the EU-27 energy sector. The term 'energy sector' should be given quite a broad interpretation, as it includes energy use in all sectors, such as industry, households, and transport. In 2007, this accounted for roughly 80% of the total EU-27 GHG emissions of just over 5000 million tons carbon dioxide equivalents, or CO_2e (European Environment Agency, 2009).

In our long-term case, we take the entire forty-year period under consideration, finding the trajectory in $t \in [0, 40]$ along which emissions are reduced by 80% at the lowest possible cost. In the short-term case, that more traditional approach is contrasted with one where the 2010-2050 period is split up into shorter, sequential sub-periods. For instance, if the duration T of each sub-period is five years, then the short-term approach implies eight consecutive and separate cost-minimization problems, with the first covering only the (five) years 2010-2014.[1]

[1] The starting point where $t = 0$ is at the very beginning of 2010.

We explore several such sub-period durations.

In this paper, we will only consider a continuous-time version of the model. However, its actual solution was produced using a computer spreadsheet based on discrete time, with intervals of 1/5 year; that is, the spreadsheet version covered $5 \cdot 40 = 200$ discrete periods, divided into a variable number of sub-periods, as just described.

2.1 Modeling the energy sector

EU-27 total primary energy supply (TPES) E is exogenously assumed to change by some constant annual proportion ρ throughout the period 2010-2050. This proportion captures, among other things, the progress being made in improving energy efficiency. Energy consumption is divided over three main technology types, all of which are measured in EJ (10^{18} J). These are: fossil fuels (F), low-carbon technologies (x_1), and no-carbon technologies (x_2). F is determined endogenously as the difference between TPES and that which is provided by x_1 and x_2. That is,

$$F = E - x_1 - x_2 = E_0 \cdot (1 + \rho)^t - x_1 - x_2 \tag{1}$$

where E_0 is initial EU-27 TPES (in the year 2010).

In the model, it is possible to increase x_1 and x_2 through investment. If $\rho > 0$ and no investments are made, then instead F will increase along with E. All per-energy investment costs are measured in €million/EJ and are considered to be additional to the cost of fossil energy; a cost of zero would then mean simply that there is no cost differential compared to fossil energy. Thus, investment in F is treated as if free.

We assume that low-carbon technologies are 'mature', so their cost per energy c_1 is constant throughout 2010-2050; because of this constant cost, we have chosen to include among other things current nuclear and hydroelectric power in this technology group despite the low emissions profiles of such technologies. For x_2, there is a relatively high initial investment cost of $c_2 > c_1$ which over time falls towards the cost of fossil energy by a constant annual proportion μ. For instance, the learning curve of photovoltaic solar technologies is such that the price per unit of energy falls by roughly 20% every time cumulative production doubles (Bhandari and Stadler, 2009, pg. 3). In the model however, cost reductions through research, learning and economies of scale are driven not by cumulative production but by time itself. Although convenient, this does have the drawback of not allowing for direct policy influence on the cost dynamics of no-carbon technologies through, for example, a 'big push' investment strategy specifically designed to lower costs.

Technologies are also characterized by a carbon intensity, measured in million tons CO_2e/EJ. The carbon intensity of fossil fuels is given by a parameter θ. As might be expected, low-carbon technologies are associated with a carbon intensity σ, with $0 < \sigma < \theta$; and the carbon intensity τ of no-carbon technologies is zero.

Further descriptions of the three technology types are given in Table 1.

2.2 Modeling the cost of emissions

Using (1), we find that total energy-sector emissions P are given by

$$P = \theta F + \sigma x_1 + \tau x_2 = \theta(E_0 \cdot (1 + \rho)^t - x_1 - x_2) + \sigma x_1 \tag{2}$$

In addition, it is assumed there exists some exogenously given emissions pathway to be followed; indeed, our assumed policy objective is to follow it as cheaply as possible. This pathway

Table 1: Main features of technological categories.

Technology	Includes	Carbon intensity	Per-energy cost
Fossil energy, F	Coal, peat, crude oil, petroleum products as well as Canadian tar sands and other unconventional oil resources that may see future large-scale exploitation (Mohr and Evans, 2010).	High	Treated as if free
Low-carbon technologies, x_1	Natural gas, most current biofuel technologies, carbon capture and storage (CCS) as well as current nuclear and hydroelectric power.	Medium	c_1, a constant
No-carbon technologies, x_2	Wind, solar, geothermal, tidal and wave power, advanced biofuel technologies, and advanced nuclear technologies possibly including fusion power.	Zero	Initial cost of $c_2 > c_1$. Cost falls over time towards the cost of fossil energy.

is given by

$$P_{path} = P_0 \cdot 0.96^t \tag{3}$$

with P_0 Mton CO_2e the initial amount of annual emissions from EU-27 energy use. Over the course of forty years, this pathway translates into roughly 80% emissions reductions on 2010 levels.[2]

Of course, as the model covers only the EU-27 region, from this it is not possible to draw conclusions concerning global GHG concentrations. However, it may be noted that our emissions path is at least broadly consistent with low to medium stabilization levels. For instance, most studies find that in order to stabilize atmospheric concentrations at 450 ppm CO_2e, developed (Annex I) countries as a group will need to reduce economy-wide emissions by 80–95% on 1990 levels by 2050 (IPCC, 2007, chap. 13). Between 1990 and 2010, EU-27 energy emissions were reduced by approximately 6.5% (European Environment Agency, 2009); therefore our energy-sector 80% target corresponds to roughly 81% on 1990 levels.

A penalty is now imposed on exceeding ('overshooting') the emissions pathway. To ensure commensurability with investment costs (see the next subsection), we use a time-constant shadow price of emissions γ measured in €/ton CO_2e. The penalty is thus treated as a monetary cost, and reducing emissions faster than is necessary as given by the pathway yields a negative cost. It is important to note that the shadow price does *not* represent any attempt to quantify the actual negative effects of climate change, as in a 'damage function' of the

[2] $P_0 \cdot 0.96^{40} \approx 0.195 P_0$.

sort used in IAM models. However, the carbon price does drive emissions reductions in our model; hence, it is by changing γ that we are able to influence the modeled EU-27 terminal-point emissions. The implication is that we take a kind of second-best approach to building a framework for analyzing cost-efficiency; rather than explicitly locking in the abatement targets for 2050,[3] we will be adjusting the magnitude of the carbon price to ensure that they are met. The pros, cons, and implications of using this kind of framework are extensively discussed at the end of this paper.

The cost of deviating from the emissions pathway is given by

$$\gamma\left[\theta(E_0(1+\rho)^t - x_1 - x_2) + \sigma x_1 - P_0 \cdot 0.96^t\right] \tag{4}$$

which follows from combining (2) and (3).

2.3 The long-term cost-minimization problem

The model drives changes to the energy mix and emissions profile of the EU-27 by means of a cost-minimization problem which is most concisely described as a trade-off between on the one hand exceeding the emissions pathway (3), and on the other investing in alternative energy sources. Both of these alternatives imply significant costs. In the model, it is possible to reduce GHG emissions *only* through investment in alternative energy technologies; energy efficiency is not (endogenously) included, nor is REDD[4] and other measures related to land-use change, nor offsets, nor policies for consumer demand management such as individual carbon rationing. Consequently, the only relevant choice concerns when to invest, how much, and in which technology group.

Our solution to the cost-minimization problem uses the maximum principle of optimal control theory (Sydsæter et al., 2008). As such, we introduce control variables u_1 and u_2 governing how much is invested in x_1 and x_2, respectively, at any given point in time t. The problem is given by

$$\min \int_0^{40} \left(\gamma\left[\theta(E_0 \cdot (1+\rho)^t - x_1 - x_2) + \sigma x_1 - P_0 \cdot 0.96^t\right] + \alpha c_1 u_1^2 + \frac{\beta c_2 u_2^2}{(1+\mu)^t}\right) dt \tag{5}$$

$$\dot{x}_1 = \alpha u_1 - \delta, \qquad x_1(0) = x_1^0 > 0, \qquad x_1(40) \ is\ free \tag{6}$$

$$\dot{x}_2 = \beta u_2, \qquad x_2(0) = x_2^0 > 0, \qquad x_2(40) \ is\ free \tag{7}$$

$$u_1 \geq 0, \qquad u_2 \geq 0, \qquad u_1 + u_2 \leq 1 \tag{8}$$

Looking from left to right, we may note that the continuous-sum cost (5) to be minimized has three components. The first is the cost of emissions already identified in (4). The second and third terms are the costs associated with investment in x_1 and x_2, respectively. To see why, we first need to examine differential equations (6) and (7). These equations describe the processes of investment and depreciation; they are the means by which the control variables govern the two 'state' variables x_1 and x_2. Parameters α, β, and δ are all measured in EJ. α and β should be interpreted as maximum annual (gross) installed energy capacity. δ is

[3] Imposing 'pure state' terminal conditions such as an emissions target for 2050 was not possible during the time frame of writing this paper.

[4] Reducing emissions from deforestation and forest degradation.

an annual depreciation rate which has the additional interpretation of being a measure of institutional, political and technical ‘lock-in’ with regard to low-carbon technologies, making it more difficult to phase out such energy systems (Sandén and Azar, 2005); being composed of relatively new technologies, it is assumed x_2 does not depreciate. Now, note that if $u_1 = u_2 = 0$ (no investments are made), then $\dot{x}_1 = -\delta$ and $\dot{x}_2 = 0$. That is, over time x_1 will decrease linearly, x_2 will remain constant, and (as mentioned in Section 2.1) F and P will tend to grow if $\rho > 0$. Greater values of u_i must then imply more investment and faster growth in x_i, and also generally less emissions.

Now, because the second and third cost terms in (5) include u_1 and u_2, they are clearly linked to the volume of investment being made in x_1 and x_2 at any point in time. For instance, if $u_2 = 1$ for $t = 0$, then $\dot{x}_2 = \beta$, while the third term in (5) reduces to βc_2. Given that c_2 is a per-energy cost, βc_2 must be the total cost of the investment being made (yielding additional capacity β).

However, we may note that unlike in (6) and (7), u_1 and u_2 enter into the cost function as *squared* factors, reflecting the increasing marginal costs of investment for any given t; thus, the correspondence between any growth in ‘green’ energy production given by the (linear) differential equations and its (quadratic) associated cost in (5) is generally not perfect. In other words: because we defined c_1 and c_2 as constant per-energy cost parameters, we might expect costs to rise linearly with the scale of energy investment, with total investment costs simply being per-energy cost multiplied by energy quantity installed; yet the (quadratic) functional form of the second and third terms of (5) implies that the actual modeled relationship between scale and cost is nonlinear. In fact, our formulation implies that the correspondence of the example given above actually only applies when $u_i = 0$ or $u_i = 1$, although it is generally at least reasonably close for $u_i \in (0, 1)$.[5]

The three restrictions (8) on control vector $\mathbf{u}$ imply that there are limits to the speed of growth in $\mathbf{x}(t)$; intuitively, there are only so much investment funds available, making it impossible to phase in low-carbon and/or no-carbon technologies overnight. For instance, if $u_1 = 1$, then by necessity $u_2 = 0$, and vice versa.

Finally, if $\gamma = 0$ (no penalty is associated with GHG emissions), then intuitively as well as analytically, optimal cost-minimizing policy is to set $u_1 = u_2 = 0$ for all $t \in [0, 40]$. Climate-friendly energy technologies, then, are never deployed in the model because of any inherent profitability vis-à-vis fossil energy, but only as part of a conscious decision to lower emissions.

The full analytical solution to this problem is given in an appendix at the end of this paper. For any set of parameter values, the solution is given by an optimal and admissible pair $(\mathbf{u}^*(t), \mathbf{x}^*(t))$. That is, $\mathbf{u}^*(t)$ and $\mathbf{x}^*(t)$ must solve (5) and satisfy restrictions (6-8). The solution also involves finding ‘adjoint functions’ $\mathbf{p}(t)$ associated with $\mathbf{x}^*(t)$; $p_i(t)$ can be interpreted as a forward-looking shadow price (marginal cost) of $x_i^*(t)$.

[5]A short word on why the polynomial $f(u) = u^2$ was chosen, as opposed to some other functional form. In order to capture the increasing nature of marginal investment cost, we wished to find some simple increasing, convex, nonlinear function. In addition, for the reasons given above, we wished to remain reasonably close to the line $f(u) = u$, not least at the end points of the interval $[0, 1]$. Thus, for $u \in [0, 1]$, $f'(u) > 0$; $f''(u) > 0$; $f(0) \approx 0$; $f(1) \approx 1$; and more generally $f(u) \approx u$. $f(u) = u^2$ then seems a natural candidate.

2.4 The short-term cost-minimization problem

The only feature of the model that is changed for the short-term problem is the cost function to be minimized. In its short-term formulation, this can be generally written

$$\min \int_0^T \left(\gamma \left[\theta(E_0 \cdot (1+\rho)^{Tk+t} - x_1 - x_2) + \sigma x_1 - P_0 \cdot 0.96^{Tk+t} \right] + \alpha c_1 u_1^2 + \frac{\beta c_2 u_2^2}{(1+\mu)^{Tk+t}} \right) dt$$

with T the length in years of each sub-period, and k the number of subperiods preceding the one under consideration. For instance, for the third eight-year period $T = 8$, $k = 2$. The above expression ensures that time-dependent processes such as the emissions path and the learning curve of x_2 do not behave differently from the long-term case over the course of the forty-year period.

The short-term problem has been examined for the separate cases of $T = 4$, $T = 5$, $T = 8$, $T = 10$, and $T = 20$. Also, the long-term problem can be considered a special case of the short-term one, where $T = 40$ and $k = 0$. Note that in each case T is always constant throughout the 2010-2050 period; all individual sub-periods are of equal length.

2.5 Finding baseline parameter values

This section provides some reasoning for the specific parameter values chosen under the baseline scenario. These figures, which are given in Table 2, are used in the long-term case; the short-term problem is dealt with separately in Subsection 2.5.5.

2.5.1 Initial values and energy-sector growth

In setting starting values, data for the year 2007 has been consistently used as proxy for the year 2010. Values for E_0, x_1^0 and x_2^0 were taken from the International Energy Agency (IEA) online statistical database (IEA, 2010a) As indicated in Table 1, x_1^0 includes natural gas, nuclear and hydro as well as the category 'combustible renewables and waste'. x_2^0 includes only the category 'geothermal, solar, etc'. E_0 does not include the minor categories 'electricity' and 'heat'. The IEA also provided a source for the value of P_0 (IEA, 2009). Finally, ρ has been set in line with projections for 2030 by the European Commission (2008), and is also roughly consistent with post-1990 EU-27 TPES historical growth rates (IEA, 2010b).

2.5.2 The carbon intensity of technologies

Parameters θ and σ describing the carbon intensity (as measured in million tons CO_2e/EJ) of F and x_1 have been based on a simple calculation using IEA data (IEA, 2009, 2010a). The technological category F used in this model includes coal, peat, crude oil, and petroleum products. In 2007 total EU-27 emissions from these categories were 2893.9 million tons. The same year, a total of approximately 39.2 EJ of fossil energy (as just defined) was consumed. Dividing the one over the other, we find a fossil carbon intensity θ of 73.7 million tons/EJ.

Comparing emissions from peat, coal, and oil with P_0, we find that this leaves another 1032.5 million tons of energy emissions, mostly caused by combustion of natural gas. According to the IEA, the total 2007 energy output of the four technological categories included in x_1 was approximately 33.7 EJ. Thus, $\sigma = 30.7$ million tons/EJ.

The fact that θ and σ are constants implies that the carbon intensity of technologies is assumed not to change before 2050; F and x_1 will not see any aggregate carbon efficiency gains (or losses) being made despite CCS, tar sands and other new technologies or fuels possibly entering the market or becoming dominant in the future.

2.5.3 The cost of energy investment

Recognizing the difficulty in identifying a single parameter value for an entire range of technologies and countries, c_1 and c_2 are set somewhat arbitrarily. For example, the cost of solar PV technologies is generally much higher than that of on- and offshore wind power (REN21, 2008), yet both are included in the x_2 aggregate. The actual values chosen are inspired by the general levels of EU feed-in tariff price guarantees encouraging renewable energy deployment (Europe's Energy Portal, 2010). Of course, it should be noted that feed-in tariffs concern electricity generation rather than overall energy production. Nevertheless, it is assumed that the investment price of x_1 is €0.1/kWh, while the price of x_2 is €0.3/kWh. In addition, since c_1 and c_2 are defined as the price *differential* compared to conventional fossil energy, €0.05/kWh (representing the cost of conventional energy) has been subtracted from these figures (implying that $c_2 = 5c_1$). Finally, costs have been redefined as €million/EJ.

μ has been defined in such a way as to make it possible to specify the t at which

$$c_1 = \frac{c_2}{(1+\mu)^t} \tag{9}$$

meaning that the per-energy costs of x_1 and x_2 are equal. Solving for μ in (9), we find that

$$\mu(t_b) = \left(\frac{c_2}{c_1}\right)^{1/t_b} - 1 \tag{10}$$

where in this case t_b is the duration until the per-energy prices of x_1 and x_2 reach a break-even point. If for instance this happens in the year 2020, then $\mu = 5^{0.1} - 1 \approx 0.17$.

2.5.4 Growth, depreciation, and the shadow price of carbon

As suggested in Section 2.2, given that this paper looks at cost-effective abatement policies, the carbon price γ is not set independently of other parameters. The emissions trajectory is predetermined; the issue under scrutiny is rather what it takes to achieve it as cheaply as possible. Thus, γ is set only after all other parameters have been determined, and is adjusted iteratively to ensure that the modeled cost-minimization problem actually results in emissions reductions by 2050 that are in the neighborhood of the terminal point of the mandated emissions path (in other words, 'close enough' to 80% reductions). Unless stated otherwise, γ is always set to the integer value yielding the closest fit.

However, because the maximum linear growth and depreciation rates of x_1 and x_2 (α, β, δ) are difficult to objectively determine, they have in effect been set arbitrarily in such a way as to imply that the shadow price of carbon, once iteratively adjusted, takes some plausible value. Values of α and β too low even in theory to enable 80% emissions reductions by 2050, or producing carbon prices outside of the approximate interval €5-150/ton, have been excluded for all scenarios. For the baseline scenario, α, β, and γ have been chosen to yield a carbon price of €40/ton. Because of the guesswork involved, we do not claim that the baseline scenario represents a more probable future than any of the alternative scenarios examined in Section 4.

Although it is assumed that x_2 does not depreciate, it is consistently the case throughout this paper that $\beta = \alpha - \delta$; the maximum growth rate (when $u_i = 1$) is then identical for both technology categories, the difference being that the absolute amount of x_2 may never decrease.

E_0	x_1^0	x_2^0	P_0	ρ	θ	σ	c_1	c_2	μ	α	β	δ	γ
73.58	33.66	0.67	3926.4	0.005	73.73	30.68	69400	13900	~0.17	3.1	2.1	1	40

Table 2: Summary of parameter values used for the baseline scenario with $T = 40$. μ is exactly determined by (10); the break-even year is set to 2020, so $t_b = 10$.

2.5.5 The short-term case

The short-term case problem relies on two sets of parameters. In the first instance, the parameter values in Table 2 are transferred completely intact from the long-term model to the short-term one, allowing illustrative direct comparisons between the two regimes, not the least in terms of terminal-point emissions.

Next, starting from those parameter values, we compensate for any observed differences in terminal-point emissions by adjusting γ, so that emissions are reduced by 80% even in the short-term case. Note that the adjustment of γ nevertheless produces a single, 'blanket' carbon price throughout 2010-2050; that is, for any given T, γ is not readjusted within each individual subperiod. Then, by examining the resulting relative magnitude of carbon prices in the long-term and short-term cases, we may draw conclusions regarding the price incentives needed to achieve long-term abatement objectives under each regime, and thus whether short-term optimization is realistically compatible with meeting such targets. In summary, first γ is kept constant; then terminal-point emissions are kept constant (by variation of γ).

3 Results of the baseline scenario

Using the parameters given in Table 2, the long-term cost-minimization problem yields the emissions trajectory and investment costs shown in Figure 1. Clearly, terminal-point emissions are in very close alignment with those mandated by emissions path (3), although there is significant overshoot during preceding decades. This cumulative amount by which the emissions path is exceeded is roughly 17.4 Gigatons CO_2e.

Averaged over the entire forty-year period, investment in low-carbon and no-carbon technologies cost €53.95 billion every year in excess of equivalent fossil fuel investments. However, investment is markedly front-loaded, with much greater costs borne in the near term; annual (additional) investment costs peak immediately in 2010 at €171.75 billion, corresponding to roughly 1.5% of current (2009) EU-27 gross domestic product (Eurostat, 2010).

3.1 Constant carbon prices

Turning to the short-term problem, we find some starkly contrasting results. As previously indicated, we first transfer all parameters, including the carbon price γ, to the short-term framework. Starting with the case when $T = 4$, some strong qualitative results emerge. With the carbon price γ constant at €40/ton, as shown by Figure 2 the exclusive focus on short-term

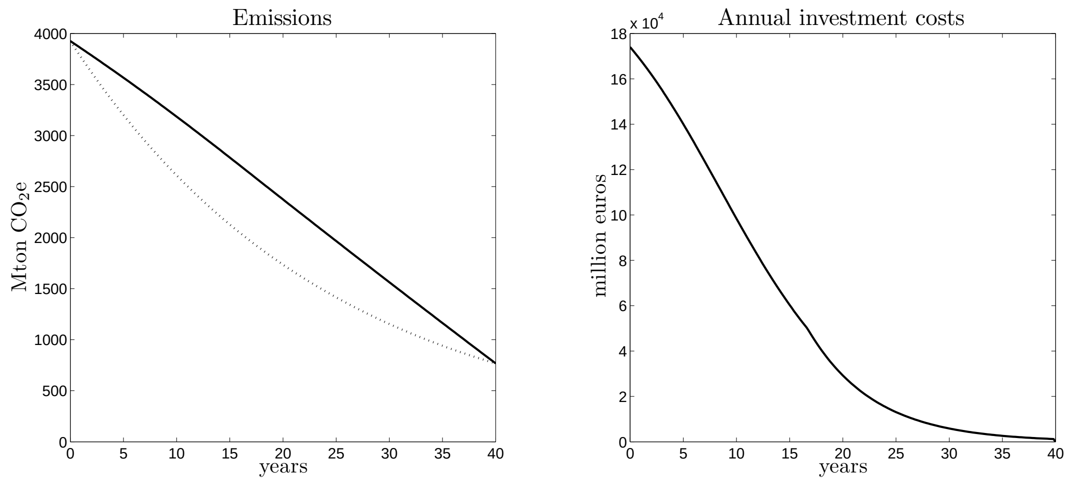

Figure 1: Baseline emissions and investment costs with $T = 40$. The left panel shows modeled emissions against the mandated emissions path (dotted line).

efficiency results in much lower investment rates and correspondingly less ambitious emissions reductions. Also, investment is no longer front-loaded, but highly irregular, peaking in the year 2022 at a much lower level than in the long-term case. This is an important result: for the baseline scenario at least, the effect of changing time frame without adjusting γ clearly is not negligible. In fact, it strongly dictates the terminal state.

To understand the underlying processes producing such radical differences, we may consider Figure 3. As might be expected, in the long-term case optimal controls $u_1^*(t), u_2^*(t)$ are continuous over time, resulting in smooth and rapid deployment of non-fossil energy sources (with some initial emphasis on x_1 while the cost of x_2 is still high). However, in the short-term case (here, with $T = 4$) the optimal controls are clearly discontinuous at the breaks between

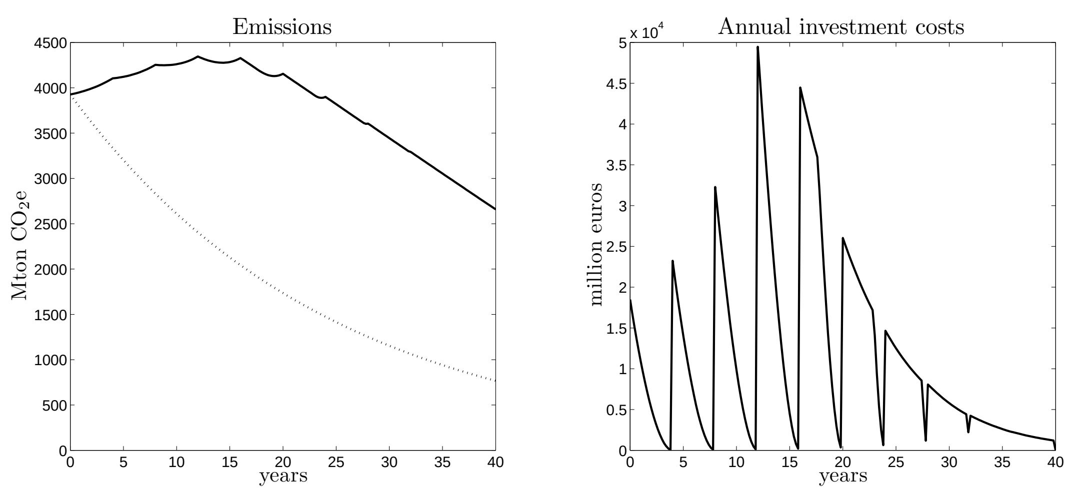

Figure 2: Baseline emissions and investment costs with $T = 4$. The left panel shows modeled emissions against the mandated emissions path (dotted line).

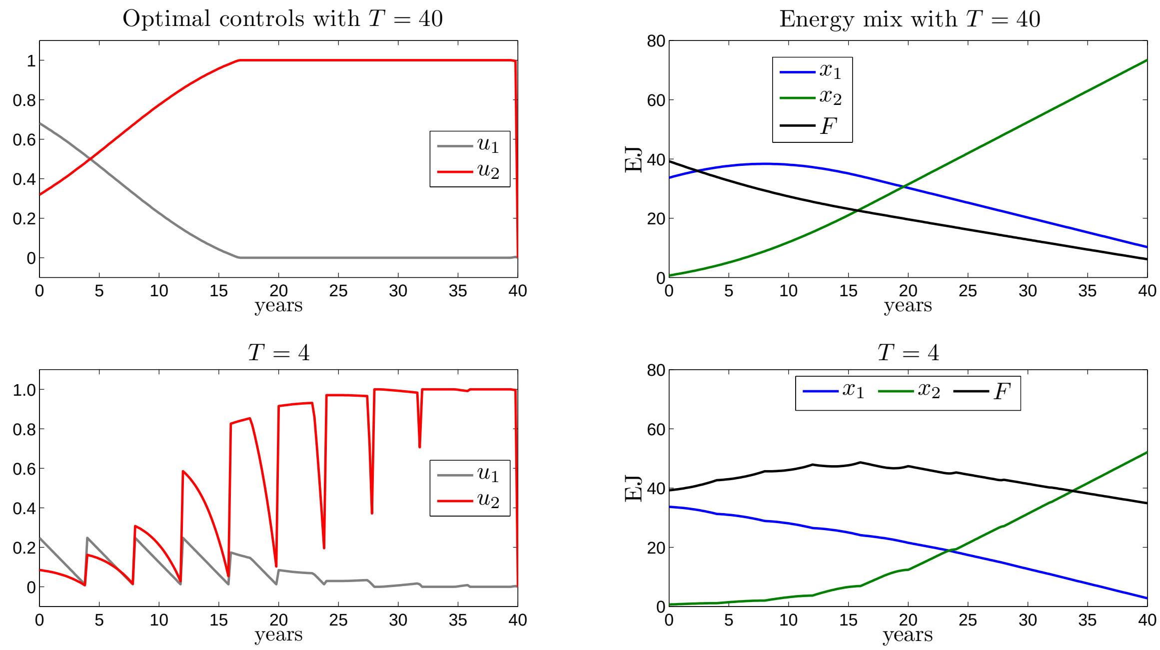

Figure 3: Baseline optimal control variables and resulting energy mix for $T = 40$ and $T = 4$.

periods, owing to the fact that $\mathbf{u} \to 0$ as $t \to T$.[6]

The intuition as to why this happens is that within each separate period, returns on a given investment—in terms of the avoided penalty for overshooting the mandated emissions path—are only counted up until the end of the period; what happens beyond T is not considered. The situation is somewhat similar to the standard intertemporal consumption problem where a rational consumer will consume all her wealth (rather than save any of it) in the final period (Varian, 1992). Analogously, in our model it pays considerably less well to invest as T grows nearer. Clearly, this is also the reason why overall incentives to invest are less in the short-term case (where $T = 4$, for instance) compared to the long-term one (where $T = 40$).

We may also note that short-term optimization affects not only long-term policies, but those of the short term as well; apparently, dynamic inconsistency creates general disincentives for *any* climate policy. The implication is that, somewhat paradoxically, short-term optimization discourages meeting short-term targets. The implications of this apparent contradiction are further discussed at the end of this paper.

Proceeding to examine other values of T, Table 3 indicates a monotonous relationship between the length of the sub-periods and the magnitude of investment as well as the amount of cumulative and final emissions. A graphic representation of the relationship between T and 2050 carbon emissions is shown in Figure 4; for clarity, the data points (in this and similar figures) have been connected by straight lines.

[6]The adjoint functions $\mathbf{p}(t)$ both include the factor $(T - t)$, which is carried over to formulae for $u_1^*(t)$ and $u_2^*(t)$. See the appendix at the end of this paper.

	$T = 40$	$T = 20$	$T = 10$	$T = 8$	$T = 5$	$T = 4$
2050 emissions	767.36	957.03	1679.03	1 902.91	2 424.60	2 658.26
Cumulative overshoot	17 380.43	21 682.23	45 633.02	52 639.00	68 529.36	75 043.83
Total investment cost	2 158.02	1 951.30	1 103.81	930.42	559.63	446.95
Average annual investment cost	53.95	48.78	27.60	23.26	13.99	11.17

Table 3: Model results from baseline scenario with carbon prices γ constant. All emissions in million tons CO_2e; all costs in €billion. Note that $P_{path}(40) = 767.09$.

3.2 Adjusted carbon prices

As described in Section 2.5.5, we now attempt to adjust (raise) γ to bring the modeled EU-27 economy back to the terminal point of the mandated emissions path (where $P = 767.09$), so that once again reducing emissions by 80% by 2050 minimizes overall costs. Depending on the degree to which the values of γ chosen can be considered realistic, we may draw conclusions regarding the severity of time-inconsistent climate policy under the baseline scenario. The question we are interested in is this: is the kind of disincentive just unearthed really a major problem? If it were the case that terminal-state emissions are sensitive to changes in γ, then perhaps no more than a minor adjustment to the carbon price might suffice to ensure that abatement targets are met. On the other hand, if very large and potentially unrealistic values of γ are required to compensate, then time-inconsistency has serious implications.

Unfortunately, it appears that the latter is true. As Table 4 and Figure 5 show, the

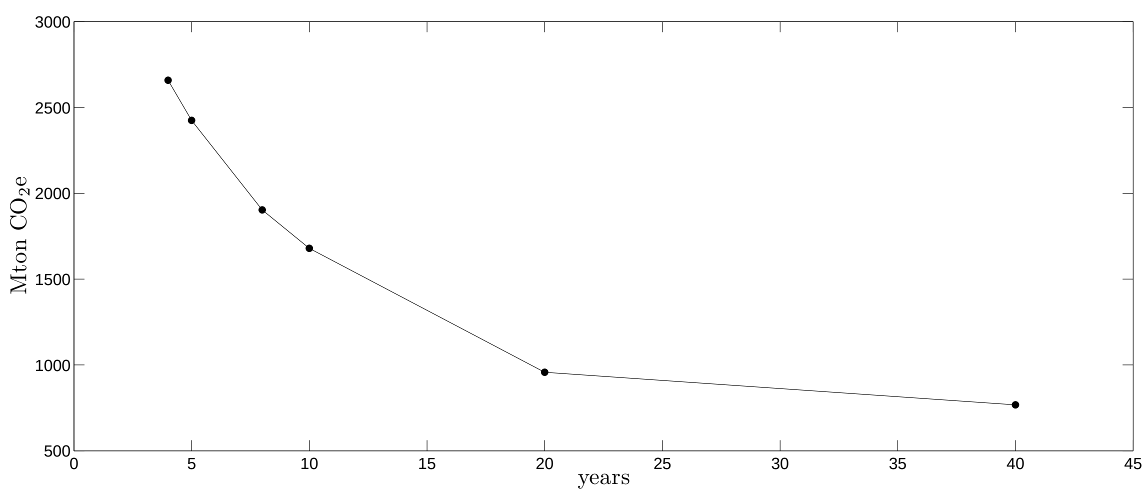

Figure 4: The relationship between sub-period duration T (x-axis) and terminal-point emissions (y-axis) under the baseline scenario.

	$T=40$	$T=20$	$T=10$	$T=8$	$T=5$	$T=4$
γ	40	170	458	548	873	1014
2050 emissions	767.36	767.05	767.15	767.16	767.04	767.09
Cumulative overshoot	17 380.43	17 024.33	17 270.69	17 316.14	17 370.30	17 401.20
Total investment cost	2 158.02	2 371.02	2 430.36	2 469.38	2 427.10	2 400.57
Average annual investment cost	53.95	59.28	60.76	61.73	60.68	60.01
Total cost	5 634.09	16 841.67	41 980.17	49 915.51	78 248.31	90 624.48

Table 4: Baseline model results with γ adjusted and terminal-point emissions constant. All emissions in million tons CO_2e; all costs in €billion. Note that $P_{path}(40) = 767.09$.

relationship between T and adjusted γ unsurprisingly is monotonous and has a similar shape to the one between terminal-point emissions and T in Figure 4. However, with T small, compensating for disincentives requires γ to be set at extremely high values, even in excess of €1000/ton, throughout the entire period 2010-2050. Compared with the EU-ETS carbon prices observed so far—of approximately €10-30/ton (European Climate Exchange, 2010)—these figures are so large as to be plausibly considered unrealistic. Our baseline scenario then suggests that time inconsistency has serious implications.

Figure 6 shows that there are remarkable similarities in terms of the modeled dynamic structure of the EU-27 energy sector through 2010-2050; with γ everywhere adjusted, we find

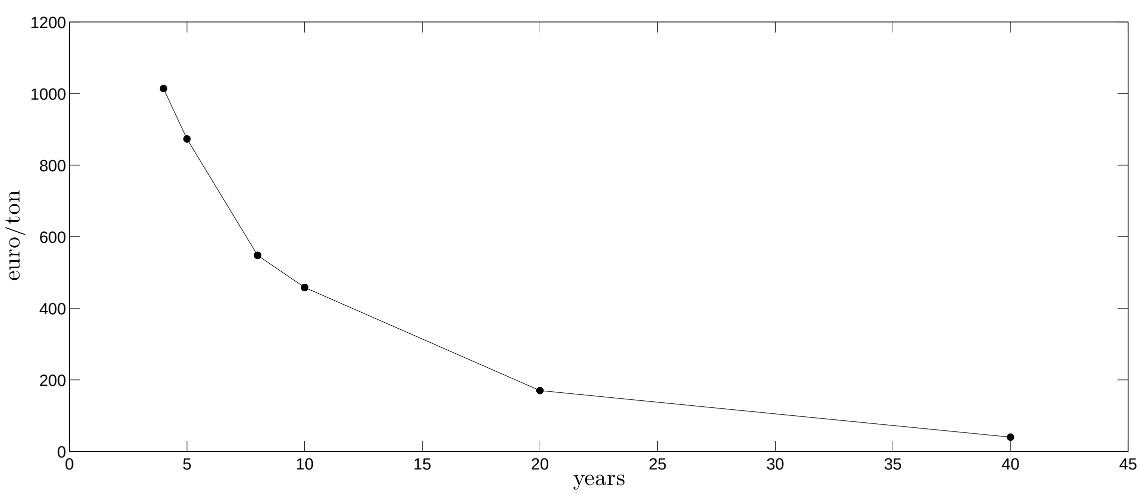

Figure 5: The relationship between sub-period duration T (x-axis) and adjusted carbon price γ (y-axis) under the baseline scenario.

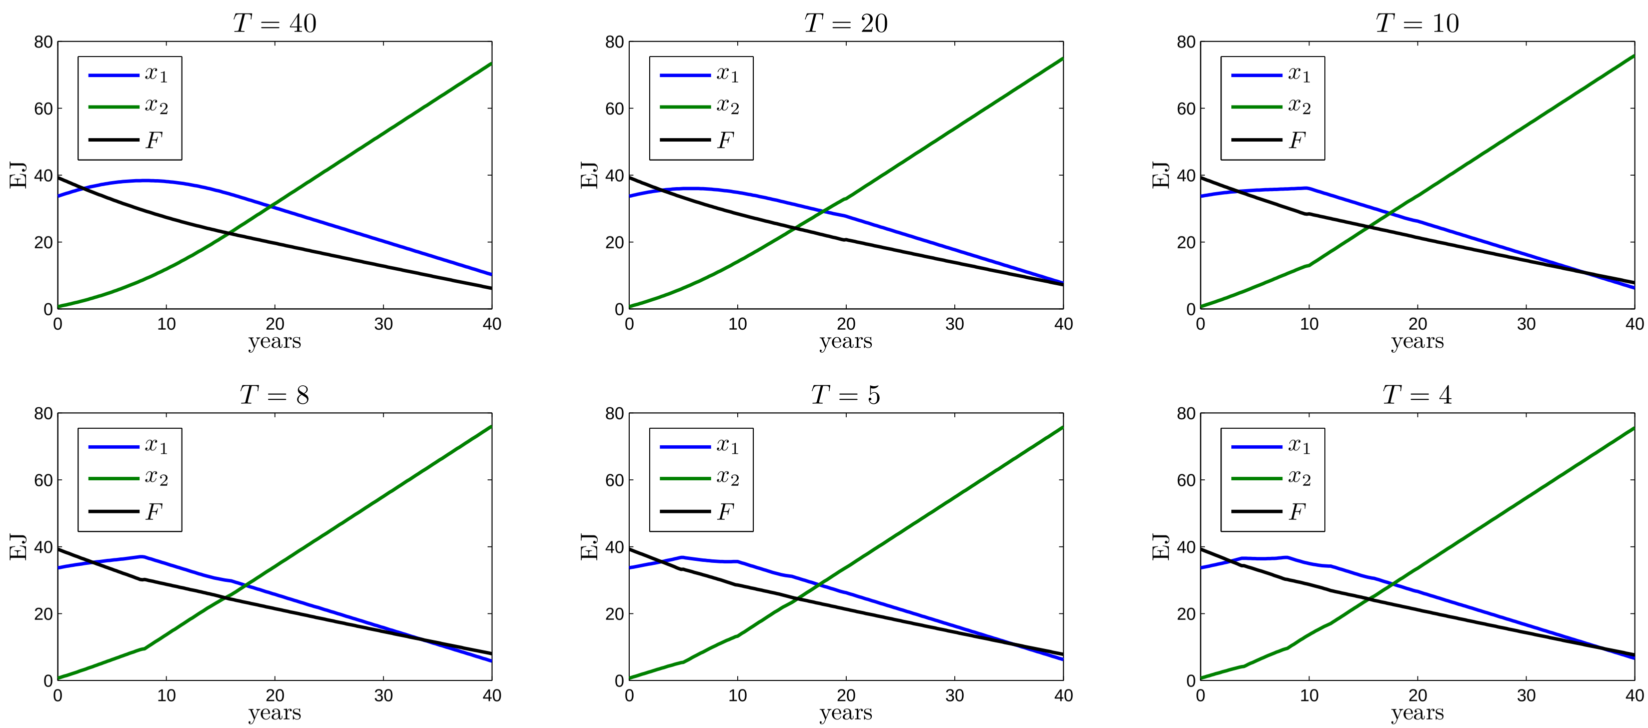

Figure 6: Modeled energy mix for various values of T under the baseline scenario, with γ adjusted to ensure that abatement targets are met.

near-identical results. F is always phased out starting immediately at $t = 0$; x_1 sees some initial growth yet starts to lose ground within 5–10 years; and x_2 grows rapidly throughout the entire forty-year period. There is then some indication that under the baseline there is in general only one way to cost-effectively attain long-term targets; and that, all else being equal, short-term optimization distracts from the task of following that one pathway.

Finally, because by realigning terminal-point emissions we make constant the effectiveness of policy, we may also examine the relative overall cost-efficiency of the long-term and short-term regimes simply by comparing their costs. As shown in Table 4, once γ is adjusted, total investment costs as well as cumulative overshoot are revealed as being broadly similar for most T, though investment costs are roughly 10% lower with $T = 40$. Thus, even if the costs of overshooting the mandated emissions path are not considered, setting $T = 40$ seems to yield somewhat more cost-efficient results over the entire forty-year period. Total costs, including the cost of overshooting, grow much higher with T small; however, this is overwhelmingly due to the extremely high carbon prices implied in such cases.

4 Sensitivity analysis

4.1 Long-term optimization

To examine whether the results of the previous section are generally true regardless of parameter values, we have explored several alternative narratives in addition to the baseline scenario (Table 5). For each of these scenarios some small subset of parameters is altered; for instance, the 'worst-' and 'best-case' scenarios entail changing only α, β, and the 'break-even point' t_b used to determine μ. The right column of Table 5 describes which parameters are changed in each case, and also the new value used.

Scenario name	Narrative	New parameter values
Best-case	Costs of no-carbon technologies decline rapidly, leading to a break-even point with low-carbon technologies in 2015; relatively high capacity to quickly implement abatement policies	$\alpha = 3.20$ $\beta = 2.20$ $\mu \approx 0.38$
Worst-case	Costs of no-carbon technologies decline slowly, leading to a break-even point with low-carbon technologies in 2025; relatively low capacity to quickly implement abatement policies	$\alpha = 3.08$ $\beta = 2.08$ $\mu \approx 0.11$
Lock-in	Due to institutional, political and technical path dependency, low-carbon technologies are phased out slowly in the absence of new investment	$\alpha = 2.85$ $\delta = 0.75$
High-efficiency	Much progress is made in improving energy efficiency, leading to negative growth in TPES	$\rho = -0.01$

Table 5: Scenario description. All parameters not included in the right column (except γ) are as given in Table 2.

As with the baseline scenario, we will use the long-term case where $T = 40$ as a starting point. Obviously, to ensure 80% emissions reductions, γ then needs to be readjusted under each new scenario to compensate for the other parameter values being altered. The new values of γ are in line with expectations: the best-case scenario, where lowering emissions is relatively cheap and easy, requires a comparably low shadow price of carbon to achieve abatement objectives. Conversely, the worst-case scenario implies quite a high carbon price. The lock-in scenario allows for a lower value of γ because a low depreciation rate δ implies that there is less need to invest in low-carbon technologies to achieve any given value of x_1.

However, slow depreciation of x_1 may imply other difficulties. Qualitatively speaking, for very low values of δ, additional emissions reductions towards the year 2050 become increasingly difficult because once F is mostly phased out, emissions associated with x_1 cannot quickly be reduced. In the end, the only way to achieve abatement objectives may then be to let $x_1 + x_2 > E$, which implies $F < 0$ and thus negative emissions from fossil fuels. Because there is no meaningful interpretation of $F < 0$, in effect this implies that the lock-in of low-carbon technologies would render long-term abatement objectives unattainable. However, in our lock-in scenario this effect is not severe, with F remaining positive throughout 2010-2050.[7]

	Baseline	Best-case	Worst-case	Lock-in	High-efficiency
γ	40	6	143	8	1.5
2050 emissions	767.36	735.99	767.29	781.37	773.88
Cumulative overshoot	17 380.43	21 631.65	16 560.59	23 572.53	34 876.70
Total investment cost	2 158.02	621.29	4 624.65	1 428.65	151.84
Average annual investment cost	53.95	15.53	115.62	35.72	3.8
Total cost	5 634.09	1 270.24	16 465.44	2 371.55	413.41

Table 6: Key scenario results for $T = 40$. All emissions in million tons CO_2e; all costs in €billion. Note that $P_{path}(40) = 767.09$.

We record the new values of γ for $T = 40$ under each scenario,[8] along with some other key statistics, in Table 6. The results of the baseline scenario are reproduced for quick reference.

As Figure 7 shows, there are marked differences between the scenarios in terms of both the cost and the optimal timing of energy investment. Most scenarios feature generally front-loaded investment, with the bulk of such costs borne in the first half of the forty-year period,

[7] In addition to the scenarios described in Table 5, several scenarios were examined and found not to yield meaningful results, one of which was a second, 'severe' lock-in scenario with $\delta = 0.35$ where the pattern described above did become an issue. We also examined a 'low-efficiency' scenario where $\rho = 0.02$; here, energy demand grows so quickly that even for extremely high carbon prices low-carbon and no-carbon technology investment rates are unable to keep up.

[8] Regarding the decimal point under the high-efficiency scenario: setting γ only to integer values yields quite inaccurate terminal-point emissions in some cases; thus, we have sometimes allowed for more finely-tuned carbon prices under this scenario.

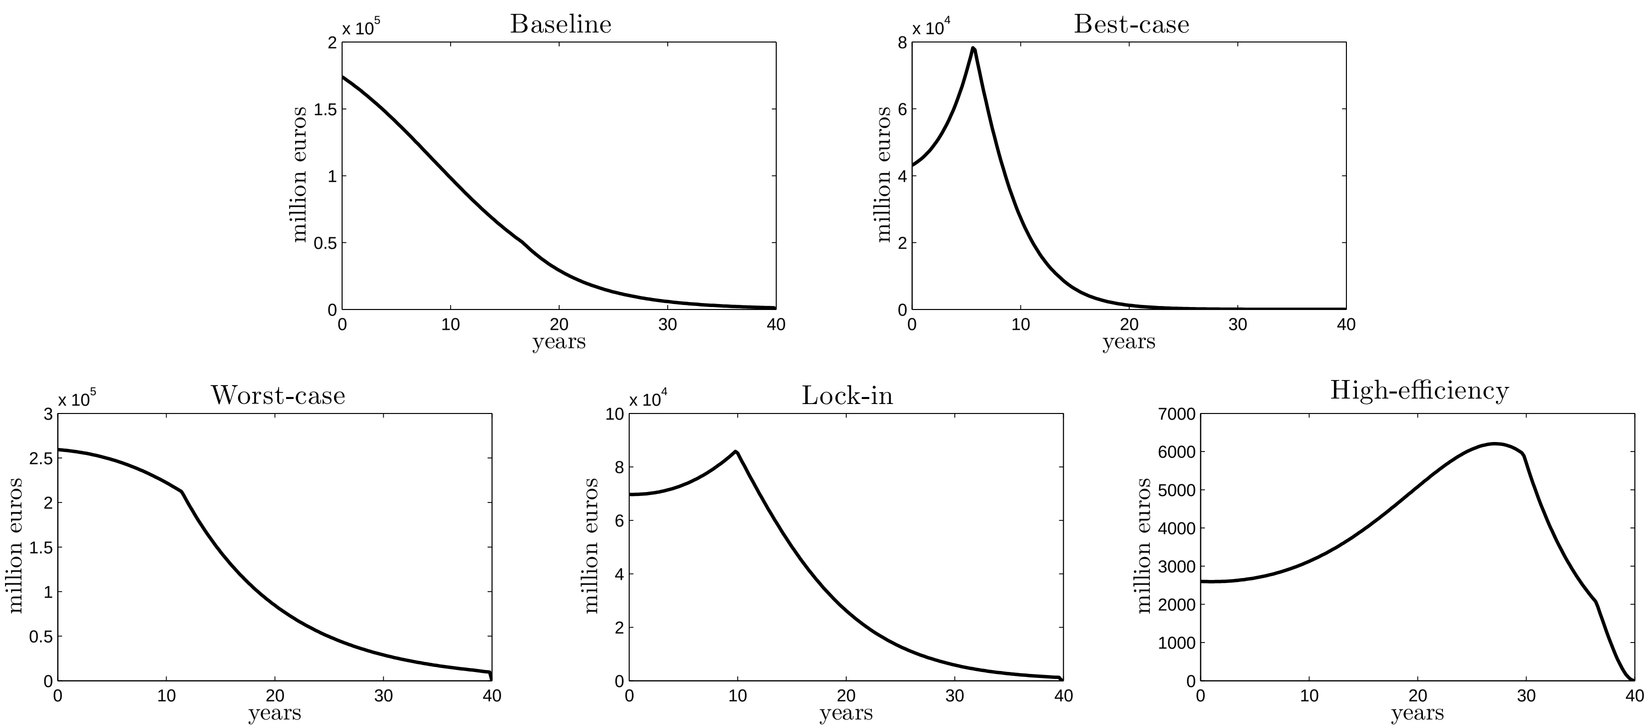

Figure 7: Annual investment for all scenarios. In all cases $T = 40$. Note that the cost scale differs across scenarios.

although investment peaking dates may differ. For instance, the worst-case scenario sees average annual costs in excess of €100 billion, with an initial cost of roughly €258.87 billion, or 2.2% of current EU-27 gross domestic product (Eurostat, 2010). In contrast, under the high-efficiency scenario where overall EU-27 TPES sees negative growth, it is optimal to delay much of the investment beyond 2030, and annual investment costs do not peak until the year 2037. Also, the total annual (additional) investment cost associated with abatement is much smaller, averaging at less than €4 billion and peaking at the miniscule amount of 6.2 billion. This seems to indicate that aggressive policies targeting energy efficiency (which of course in itself would imply significant costs that are not considered here) may be feasible in the near term as a substitute for rolling out low-carbon and no-carbon energy technologies. Still, allowing for delayed investment in low-carbon and no-carbon technologies does come at a price: the cumulative emissions overshoot is significantly greater than under any other scenario.

4.2 Short-term optimization

As before, we now transfer all parameter values from the long-term to the short-term framework and keep γ constant. This produces results that are qualitatively similar to those of the baseline scenario: there is a monotonous relationship between the length of each sub-period T and terminal-point emissions (Figure 8), though the exact shape of the relationship differs across scenarios. We may also conclude that for any given carbon price γ, shifting to smaller values of T once more results in discontinuities in $\mathbf{u}^*(t)$ between sub-periods, as well as overall disincentives both in the near term and long term.

Next, we again adjust γ to account for any differences in the terminal-point emissions reported in Figure 8. Figure 9 describes the resulting values of γ for all T and all scenarios. The left panel of Figure 9 shows these carbon prices as absolute values. For $T = 4$, values of γ span the interval between €2907/ton CO_2e for the worst-case scenario, and €15/ton for

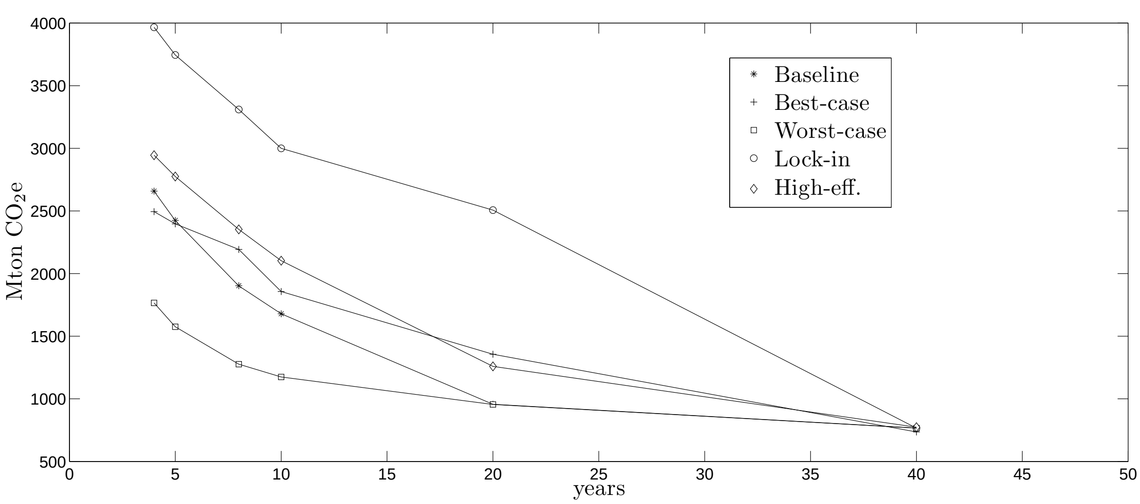

Figure 8: The relationship between sub-period duration T (x-axis) and terminal-point emissions (y-axis) under all scenarios, with γ constant.

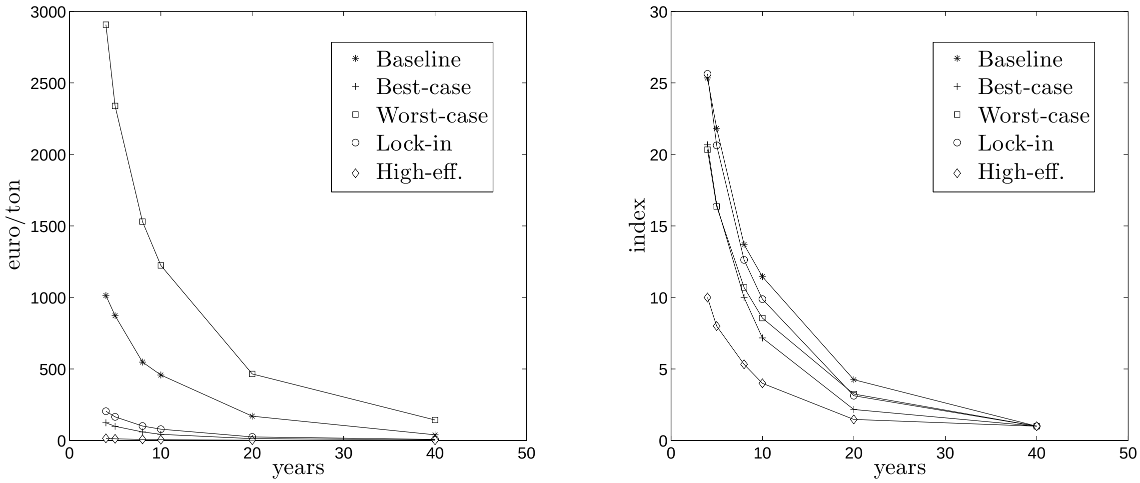

Figure 9: The relationship between sub-period duration T (x-axis) and adjusted carbon price γ (y-axis) under all scenarios; in absolute figures (left panel) and as an index, with the value for $T = 40$ set equal to 1 (right panel).

the high-efficiency scenario. When sub-periods are short ($T \leq 5$), all but the high-efficiency scenario report carbon prices around or in excess of €100/ton.

What is the general nature of the disincentives produced by short-term optimization? The right panel of Figure 9 shows adjusted carbon prices expressed as indexed values relative to the $T = 40$ case, which is normalized to 1. Expressed in this way, we see that the relationship between adjusted γ and T is remarkably consistent across scenarios, with only the high-efficiency scenario really standing out. There are clear indications towards the existence of some fairly uniform 'carbon price elasticity with respect to T'. For instance, disregarding the high-efficiency scenario[9] we find that regardless of the actual numerical value of γ used when $T = 40$, if we change the time frame to $T = 4$, compensating for disincentives requires increasing γ by a factor of roughly 20-25.

As was the case under the baseline scenario, once γ is adjusted the modeled energy mix under each scenario develops along similar lines regardless of T. However, as our discussion on low-carbon lock-in and efficiency policies suggested, there is some variation *between* scenarios. For instance, compared to the baseline, the worst-case, best-case, and especially the high-efficiency scenario sees less initial emphasis on low-carbon technologies,[10] yet the opposite is true for the lock-in scenario, just as one might expect. Most scenarios see fossil fuels phased out starting immediately in 2010, yet under the high-efficiency scenario the volume of F is relatively stable for approximately the first fifteen years. Thus, to some extent the different assumptions underlying our scenarios allow for diverging pathways towards meeting abatement targets. Nevertheless, by 2050 all scenarios feature no-carbon technologies as the

[9] As we have seen, this scenario features strong policies on energy efficiency that are not driven by the carbon price, but rather are given exogenously. Thus, there may be good reasons for ignoring it in this context.

[10] The high-efficiency scenario is in fact somewhat problematic due to x_1 consistently being negative for most of the final year (2049) at least; this is because low-carbon technologies are generally allowed to depreciate without being replaced under this scenario. Moreover, with γ unadjusted, the smaller the value of T, the more severe is this effect; for $T = 4$ and γ unadjusted, $x_1 < 0$ for $t \geq 34.2$.

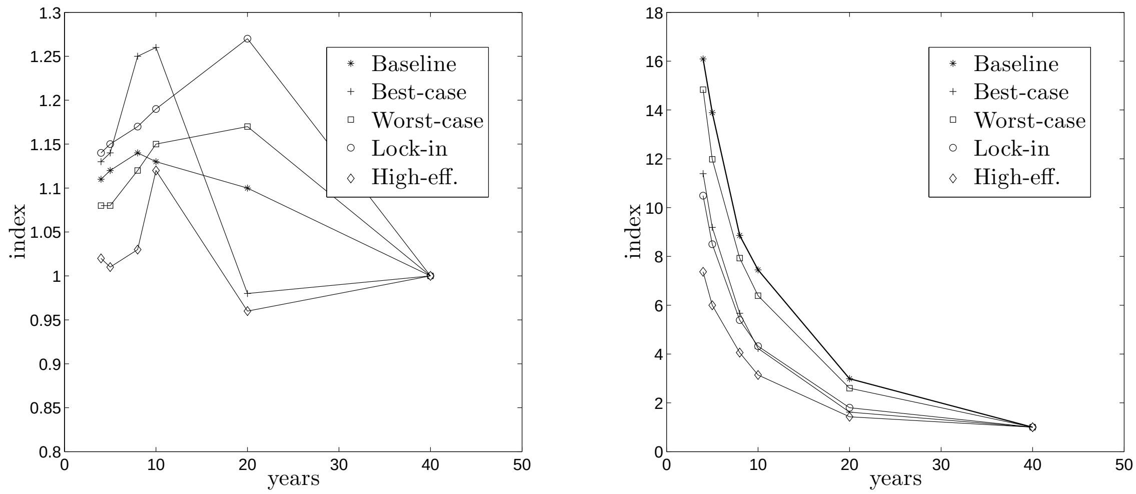

Figure 10: The relative cost-effectiveness of various values of T, under all scenarios. Left panel: sub-period length T (x-axis) and indexed investment costs (y-axis). Right panel: sub-period length T (x-axis) and indexed total costs (y-axis), under all scenarios.

dominant energy source, with both F and x_1 in steady decline.

Finally, we may examine the relative cost-effectiveness of various T within each scenario. Counting only direct investment costs and once again indexing against the case when $T = 40$, we find costs to be generally stable regardless of T, with $T = 40$ and $T = 20$ somewhat cheaper than shorter sub-period lengths (Figure 10, left panel). $T = 40$ thus seems consistently more cost-effective, by a factor of up to (roughly) 25%, than shorter sub-periods—with the possible exception of $T = 20$. This supports the general conclusions of the baseline scenario. As before, total costs—including the costs of overshooting the mandated emissions pathway (Figure 10, right panel)—are mainly the result of the adjusted values of γ and the cumulative overshoot.[11]

In summary, most conclusions from the baseline scenario have survived our sensitivity analysis intact: *ceteris paribus*, short-term optimization implies disincentives for sticking to both near-term and long-term abatement objectives, and correcting for those disincentives requires quite drastic readjustment of the carbon price driving investment. Also, once adjustments are made, short-term optimization, while still cost-effective on its own terms, is generally more expensive than long-term optimization when taken across the entire 2010-2050 period.

5 Concluding remarks

Procrastination in climate policy—i.e. continuing on the 'business-as-usual' emissions pathway even though some different and optimal abatement policy exists—may come at substantial cost, as shown in an IAM framework by Keller et al. (2007). Our modeling exercise provides one account of *why* such procrastination may occur, and the degree to which it may be likely unless issues of time inconsistency are addressed.

[11] The overshoot is consistently and significantly greater under the high-efficiency scenario; however, this fact is offset by the low carbon price used under this scenario, so there is little overall effect on calculated total costs.

The point of departure for this paper has been the assumption that long-term climate policy targets are time-inconsistent to some degree in the sense that policy makers will tend to seek cost-efficiency over policy periods that are shorter than the full time frame of abatement. Of course, doing so will be a major problem only if targets are found to be incompatible with sequential short-term optimization. Yet assuming that EU-27 carbon prices at or above €100/ton are unlikely to be politically realistic up until 2050, then all but one of our scenarios do indeed indicate that dynamic inconsistency is a serious issue, because attaining abatement objectives through strictly short-term ($T \leq 5$) cost-minimization requires carbon prices to be set above that threshold. In addition, the 'high-efficiency' scenario, where this was not the case, is somewhat problematic; and the relevance of its results may be compromised due to its assumption of exogenous, essentially unpaid-for aggressive policy on energy efficiency. Finally, even for longer sub-period durations, dynamic inconsistency was revealed as being problematic under some scenarios.

However, our model has also shown that unless under the short-term regime carbon prices are adjusted to implausible values, not only long-term but *near-term* policies are less ambitious compared to the $T = 40$ case. Clearly, in our model short-term optimization shifts the general focus, not from no-carbon technologies to relatively cheap low-carbon technologies—as one might imagine, and Sandén and Azar (2005) argued would be the case—but rather more generally from action to delay. Paradoxically, it seems the implication is that near-term optimization has the unfortunate and serious side effect of in fact making even near-term objectives *more* difficult—indeed all but impossible—to attain.

How may such an effect arise? To answer that question, we need to reexamine the process by which dynamic inconsistency generates disincentives for abatement. Crucially, although we have taken abatement targets 'as given' by adjusting the values of γ—in accordance with a cost-efficiency approach—our model really uses a cost-benefit framework in the sense that optimal policy follows from the trade-off between the benefits (i.e. avoided priced emissions) and (investment) costs of abatement. After all, a cost-efficiency framework would imply that a certain target must be met, if at all possible; the issue is then not *if*, but *how*. Yet assuming that policy targets will be met is of little use in explaining why, or to what degree, they might *not* be met due to dynamic inconsistency. Our model, then, is a kind of hybrid between the two approaches.

However, when abatement is framed as a cost-benefit issue, shifting to an exclusive focus on near-term costs to be minimized (as in our short-term cases) must imply placing correspondingly less weight on any and all long-term considerations, because all benefits and costs beyond the end of the sub-period are in effect discounted to zero. Yet climate change is of course at its core a long-term problem, the costs of which accumulate only slowly over the long run. Thus, our observed general disincentives can in fact be traced back to the cost-benefit framework used in the model.

It follows that implications for real-world policy optimization under a 'pure' cost-efficiency approach—where policy targets are truly taken for granted such that meeting them is a given precondition, rather than a matter of prioritization (i.e. adjusting the values of γ)—may be less clear-cut than this paper suggests; this is the essence of the paradox we observed. At the same time, we feel it is unlikely that policy makers operate within a framework where targets are set in stone; and it should be noted that even if the incentives for overall delay produced by our quasi-cost-benefit approach were eliminated, the incentives for focusing excessively on low-carbon technologies in the near term—to the detriment of long-term targets, as described by Sandén and Azar (2005)—might still remain. This paper has touched upon the issue of low-carbon 'lock-in' only briefly, yet qualitatively at least, what results there are suggest such concerns are legitimate.

In other words, there are two separate criticisms to be made towards the short-term framing of climate policy. The argument of Sandén and Azar (2005) is that although policies for averting dangerous climate change must certainly be cost-effective, pursuing cost-efficiency in achieving incremental near-term targets may lead to the premature lock-in of low-carbon technologies. And, as we have seen, there is the general cost-benefit point that addressing long-term issues, including climate change, will only be worthwhile if there is a sufficiently long-term perspective.

Our paper, then, has provided some illumination on the connections between these two arguments; and one conclusion must be that our model has revealed a fundamental contradiction inherent in the idea that near-term abatement should be cost-efficient, namely that any decision maker who truly cares only about the near term has little incentive to undertake climate policy to begin with. It makes little sense for a social planner who clearly considers the long term in deciding what policy objectives to pursue, to disregard the long term in choosing how to pursue them.

A few caveats may be made with regard to our results. Firstly, the exact degree to which climate policy is time-inconsistent—the degree to which policy makers actually practice strictly short term (as opposed to long or medium term) optimization—is uncertain. We have implicitly assumed throughout this paper that the strict short-term case where, say, $T \leq 5$, is the most relevant one. Yet in the real world, matters may be more complex.

Secondly, one might question whether our top-down 'social planner' approach to carbon pricing is in itself relevant and realistic, given that it abstracts from actual financial incentives on the ground for those private companies from which the bulk of climate-related investment will surely flow. That is, the carbon price used in our model has the interpretation of a 'weight' which the social planner places on reducing emissions, rather than as a policy instrument creating incentives for investment.

Finally, near-term options for abatement—notably energy efficiency—are either not included or treated as strictly exogenous, which is of course a dramatic oversimplification. Although simple, we nevertheless believe that our model does capture some important general aspects; specifically, no matter the number of low-carbon or no-carbon technologies to choose from, in order to attain ambitious abatement targets for 2050 most of the low-carbon and fossil technologies will have to be phased out, and the energy mix will need to go predominantly no-carbon. It seems clear that all else being equal, short-term optimization as defined in our model will detract from any such development.

If then we accept that dynamic inconsistency leads to disincentives for meeting abatement targets, what can be done? Within macroeconomics, the study of time-inconsistent optimal monetary policy has led to the conclusion that policy commitments need to be 'locked in'. In the words of Barro and Gordon (1983), "*the time-inconsistency of the optimal solution is either irrelevant—when commitments are feasible—or else this solution* [climate policy with long time frames, in our case] *does not solve the problem actually faced by the policymaker.*" There is a similar case to be made for locking in climate policy; and some progress has been made in this area, notably in the shape of the 2008 UK Climate Change Act setting a legally binding target for 2050 of reducing GHG emissions by 80% on 1990 levels (Crown, 2008). However, the Climate Change Act also divides the long-term target into sequential near-term ones; thus, although it solves our problem of overall disincentives by locking in the targets, it does not address the concerns of Sandén and Azar (2005) with regard to the pursuit of short-term cost-efficiency.

An alternative to legislation might perhaps be a kind of institutional regime analogous to that of central banks wielding independent control over monetary policy. Although it would certainly be desirable for such a climate policy body to be democratically accountable, some degree of separation of climate issues from the sphere of party politics might present promising avenues for locking in both long-term and short-term commitments while preempting the pursuit of exclusively short-term cost-efficiency in meeting them.

In conclusion, aforementioned conceptual and model-related issues notwithstanding, we would argue that because the exact difficulty and cost of ambitious climate policy is subject to uncertainty, and because there may be little margin for error in pursuing policies of climate stabilization, it is only prudent for climate change economists to explicitly consider issues of time consistency in climate policy. Our modeling exercise has provided some preliminary results underscoring this need.

6 Appendix. Solving the model for (u_1^*, u_2^*)

The problem is

$$\min \int_0^T \left(\gamma \left[\theta(E_0(1+\rho)^t - x_1 - x_2) + \sigma x_1 - P_0 \cdot 0.96^t \right] + \alpha c_1 u_1^2 + \frac{\beta c_2 u_2^2}{(1+\mu)^t} \right) dt$$

$$\begin{array}{lll} \dot{x}_1 = \alpha u_1 - \delta, & x_1(0) = x_1^0 > 0, & x_1(T) \ is \ free \\ \dot{x}_2 = \beta u_2, & x_2(0) = x_2^0 > 0, & x_2(T) \ is \ free \end{array}$$

$$u_1 \geq 0, \qquad u_2 \geq 0, \qquad u_1 + u_2 \leq 1$$

The Hamiltonian function associated with this problem is

$$H = -\gamma \left[\theta \left(E_0(1+\rho)^t - x_1 - x_2 \right) + \sigma x_1 - P_0 \cdot 0.96^t \right] - \alpha c_1 u_1^2 - \frac{\beta c_2 u_2^2}{(1+\mu)^t} + p_1(\alpha u_1 - \delta) + p_2 \beta u_2$$

Being a sum of concave functions, H is concave in $(\mathbf{x}, \mathbf{u})$ for all t in $[0, T]$. The control region U is a closed, convex set given by $\{(u_1, u_2) : 0 \leq u_1, \ 0 \leq u_2, \ u_1 + u_2 \leq 1\}$. Necessary and sufficient conditions for a solution follow.

(i) $(u_1, u_2) = (u_1^*, u_2^*)$ maximize

$$\alpha(u_1 p_1(t) - c_1 u_1^2) + \beta \left(u_2 p_2(t) - \frac{c_2 u_2^2}{(1+\mu)^t} \right)$$

for $(u_1, u_2) \in U$.

(ii) $p_1(t), p_2(t)$ satisfy

$$\begin{array}{ll} \dot{p}_1(t) = -(H'_{x_1})^* = -\gamma(\theta - \sigma) & p_1(T) = 0 \\ \dot{p}_2(t) = -(H'_{x_2})^* = -\gamma\theta & p_2(T) = 0 \end{array}$$

(iii) $x_1^*(t), x_2^*(t)$ satisfy

$$\begin{array}{ll} \dot{x}_1^* = \alpha u_1^*(t) - \delta & x_1^*(0) = x_1^0 \\ \dot{x}_2^* = \beta u_2^*(t) & x_2^*(0) = x_2^0 \end{array}$$

Note that both $x_1(T)$ and $x_2(T)$ are free; there are no restrictions on the terminal points. Consequently, as condition (ii) shows, $\mathbf{p}(T) = 0$ and thus $p_0 = 1$, so the Hamiltonian function above is correctly specified.

First, we find $p_1(t)$ and $p_2(t)$.

$$\begin{aligned} \dot{p}_1(t) &= -\gamma(\theta - \sigma) \\ p_1(t) &= -\gamma(\theta - \sigma)t + A \end{aligned}$$

Using condition (ii), we find that

$$p_1(T) = -\gamma(\theta - \sigma)T + A = 0$$

and thus

$$p_1(t) = \gamma(\theta - \sigma)(T - t) \tag{11}$$

Similarly, we find that

$$\begin{aligned}
\dot{p}_2(t) &= -\gamma\theta \\
p_2(t) &= -\gamma\theta t + B \\
p_2(T) &= -\gamma\theta T + B = 0
\end{aligned}$$

and thus

$$p_2(t) = \gamma\theta(T - t) \tag{12}$$

Inserting expressions (11) and (12) into condition (i) then gives the following maximization problem for u_1, u_2.

$$\max_{u_1,u_2} \phi = \max_{u_1,u_2} \left[\alpha(\gamma(\theta-\sigma)(T-t)u_1 - c_1u_1^2) + \beta\left(\gamma\theta(T-t)u_2 - \frac{c_2u_2^2}{(1+\mu)^t}\right)\right] \tag{13}$$

The control region U has the geometric interpretation of a triangle in $\mathbb{R}^2$ with corners in $(u_1, u_2) = (0, 0)$; $(u_1, u_2) = (1, 0)$; and $(u_1, u_2) = (0, 1)$. We wish to determine how $\mathbf{u}$ should be set in order to maximize the objective function ϕ given various parameter values and points in time t. To that end, we consider several cases.

Case I. Interior maximum
To find an interior solution, we check the first-order conditions of the partial derivatives of ϕ w.r.t. u_1 and u_2.

$$\begin{aligned}
\frac{\partial\phi}{\partial u_1} &= \alpha[\gamma(\theta-\sigma)(T-t) - 2c_1u_1] = 0 \\
\to u_1 &= \frac{\gamma(\theta-\sigma)}{2c_1}(T-t) \geq 0 \qquad (14) \\
\frac{\partial\phi}{\partial u_2} &= \beta\left(\gamma\theta(T-t) - \frac{2c_2u_2}{(1+\mu)^t}\right) = 0 \\
\to u_2 &= \frac{\gamma\theta(1+\mu)^t}{2c_2}(T-t) \geq 0 \qquad (15)
\end{aligned}$$

As ϕ is strictly concave in $\mathbf{u}$, the solution derived in this case is the global maximum of the objective function.

Note that (14) and (15) must be non-negative because all parameter values are positive (under all scenarios); $\theta > \sigma$; and it is always the case that $t \leq T$. Two of the three conditions on u_1 and u_2 used to delineate U are then fulfilled. Using the above expressions, the third condition ($u_1 + u_2 \leq 1$), requires that

$$(T-t)\left(\frac{\gamma(\theta-\sigma)}{2c_1} + \frac{\gamma\theta(1+\mu)^t}{2c_2}\right) \leq 1$$

Clearly, whether this is the case depends on what specific parameter values are chosen as well as on t. As a result, we will still need to examine the other cases. However, we will see below that in the absence of explicit parameter values these cases yield similarly inconclusive results.

Finally, we may note that the value of ϕ in the global maximum point is given by

$$\left(\frac{\alpha(\gamma(\theta-\sigma))^2}{4c_1}+(1+\mu)^t\frac{\beta(\gamma\theta)^2}{4c_2}\right)(T-t)^2 \tag{16}$$

Case II. Corners of U
First, we check $(u_1, u_2) = (1, 0)$. This produces

$$\phi(1,0)=\alpha\left[\gamma(\theta-\sigma)(T-t)-c_1\right] \tag{17}$$

Next, for $(u_1, u_2) = (0, 1)$,

$$\phi(0,1)=\beta\left(\gamma\theta(T-t)-\frac{c_2}{(1+\mu)^t}\right) \tag{18}$$

and finally, $(u_1, u_2) = (0, 0)$ naturally results in $\phi(0, 0) = 0$.

Case III. Edges of U
First, we set $u_2 = 0$. ϕ then reduces to

$$\phi(u_1,0)=\alpha[\gamma(\theta-\sigma)(T-t)u_1-c_1u_1^2]$$

We now need to find the point along the line $u_2 = 0$ where ϕ is maximized. If there is an interior maximum along this line, it will be found by imposing first-order conditions upon the derivative of ϕ w.r.t u_1. This implies that in this case u_1 will in fact be given by (14), and ϕ by

$$\frac{\alpha(\gamma(\theta-\sigma))^2}{4c_1}(T-t)^2 \tag{19}$$

The alternatives to an interior maximum are $u_1 = 1$ and $u_1 = 0$; these corner cases were previously examined.

Similarly, if $u_1 = 0$, then u_2 is given by (15), and ϕ by

$$\frac{\beta(1+\mu)^t(\gamma\theta)^2}{4c_2}(T-t)^2 \tag{20}$$

As before, the endpoints (corners) have already been examined.

Finally, we check the case when $u_2 = 1 - u_1$. The hypotenuse of the triangle-shaped control region U is a segment of this line. Along it, ϕ is given by

$$\phi=\alpha[\gamma(\theta-\sigma)(T-t)u_1-c_1u_1^2]+\beta\left(\gamma\theta(T-t)(1-u_1)-\frac{c_2(1-u_1)^2}{(1+\mu)^t}\right)$$

As before, this will either yield an interior maximum or a corner solution. Given that all corners have already been examined, we check first-order conditions.

$$\frac{\partial\phi}{\partial u_1} \quad = \quad (T-t)(\alpha\gamma(\theta-\sigma)-\beta\gamma\theta)-2\alpha c_1u_1+\frac{2\beta c_2(1-u_1)}{(1+\mu)^t}=0$$

Rearranging, we have

$$u_1 \quad = \quad \frac{\frac{(1+\mu)^t}{2}(T-t)(\alpha\gamma(\theta-\sigma)-\beta\gamma\theta)+\beta c_2}{\alpha c_1(1+\mu)^t+\beta c_2} \tag{21}$$

$$u_2 = 1-u_1 \quad = \quad \frac{(1+\mu)^t\left[\alpha c_1 - \frac{1}{2}(T-t)(\alpha\gamma(\theta-\sigma)-\beta\gamma\theta)\right]}{\alpha c_1(1+\mu)^t+\beta c_2} \tag{22}$$

Given how u_1 and u_2 are defined in this case, it is clear that the condition $u_1 + u_2 \leq 1$ must hold. However, without additional information on parameter values it is not apparent whether expressions (21) and (22) are nonnegative. That is, it is not clear whether the point in question actually lies on the hypotenuse of U, or somewhere else along the line $u_2 = 1 - u_1$.

The corresponding value of ϕ is quite lengthy and is not reproduced here. It is of course obtained by inserting (21) and (22) into ϕ.

In summary, we may conclude that there are a number of cases that may or may not apply for any given combination of parameter values and points in time t. Their values for ϕ were given by equations (16), (17), (18), (19), (20), the lengthy expression just mentioned but not specified; and finally, for $u_1 = u_2 = 0$, $\phi = 0$. The model may, and does, 'switch' between cases as t progresses. In this paper, maximizing ϕ and finding the right $\mathbf{u}^*(t)$ for every given point in time was (somewhat painstakingly) accomplished by plotting all the aforementioned expressions in an MS Excel spreadsheet and comparing numerical values across cases, while observing relevant restrictions on U. However, not all cases were relevant in practice; the model saw shifts only between the interior maximum of Case I, the corner $(u_1, u_2) = (0, 1)$ of Case II, and the hypotenuse maximum of Case III.

References

F. Ackerman, S.J. DeCanio, R.B. Howarth, and K. Sheeran. Limitations of integrated assessment models of climate change. *Climatic Change*, 95(3), 2009.

G.B. Asheim. Individual and Collective Time-Consistency. *The Review of Economic Studies*, 64(220), 1997.

R.J. Barro and D.B. Gordon. A Positive Theory of Monetary Policy in a Natural Rate Model. *Journal of Political Economy*, 91(4), 1983.

R. Bhandari and I. Stadler. Grid parity analysis of solar photovoltaic systems in Germany using experience curves. *Solar Energy*, 83(9), 2009.

Crown. Climate change act 2008, 2008. Available at `http://www.opsi.gov.uk/acts/acts2008/pdf/ukpga_20080027_en.pdf`. Last accessed 2010-05-12.

M.H.I. Dore. A Question of Fudge: Professor Nordhaus on climate change and mitigation policy. *World Economics*, 10(1), 2009.

European Climate Exchange. Prices, Volume & Open Interest: ECX EUA Futures Contract, 2010. Available at `http://www.ecx.eu/EUA-Futures`. Last accessed 2010-05-07.

European Commission Directorate-General for Energy and Transport. *European Energy and Transport. Trends to 2030 - update 2007*, 2008.

European Environment Agency. Annual European Community greenhouse gas inventory 1990–2007 and inventory report 2009. Submission to the UNFCCC Secretariat, 2009.

Europe's Energy Portal. Feed-in Tariffs, 2010. Available at http://www.energy.eu/. Last accessed 2010-04-19.

Eurostat. Gross domestic product at market prices, 2010. Available at http://epp.eurostat.ec.europa.eu/tgm/refreshTableAction.do?tab=table&plugin=1&pcode=tec00001&language=en. Last accessed 2010-04-19.

IEA. *CO_2 Emissions from Fuel Combustion. Highlights*, 2009.

IEA. 2007 Energy Balance for European Union - 27, 2010a. Available at http://www.iea.org/stats/balancetable.asp?COUNTRY_CODE=30. Last accessed 2010-04-19.

IEA. Total Primary Energy Supply. European Union - 27, 2010b. Available at http://www.iea.org/stats/pdf_graphs/30TPES.pdf. Last accessed 2010-04-19.

IPCC. *Contribution of Working Group III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, 2007.* Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 2007.

J.J. Johnson and J.R. Busemeyer. Multiple-stage decision making: the effect of planning horizon length on dynamic consistency. *Theory and Decision*, 51(2), 2001.

K. Keller, A. Robinson, D.F. Bradford, and M. Oppenheimer. The regrets of procrastination in climate policy. *Environmental Research Letters*, 2(2), 2007.

F.E. Kydland and E.C. Prescott. Rules Rather than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy*, 85(3), 1977.

S.H. Mohr and G.M. Evans. Long term prediction of unconventional oil production. *Energy Policy*, 38(1), 2010.

W.D. Nordhaus. An Optimal Transition Path for Controlling Greenhouse Gases. *Science*, 258(5086), 1992.

REN21. *Renewables 2007. Global Status Report*, 2008.

B.A. Sandén and C. Azar. Near-term technology policies for long-term climate targets—economy wide versus technology specific approaches. *Energy Policy*, 33(12), 2005.

R.H. Strotz. Myopia and Inconsistency in Dynamic Utility Maximization. *The Review of Economic Studies*, 23(3), 1956.

K. Sydsæter, P. Hammond, A. Seierstad, and A. Strøm. *Further Mathematics for Economic Analysis. Second Edition.* Pearson Education, Essex, United Kingdom, 2008.

H.R. Varian. *Microeconomic Analysis, Third Edition.* W.W. Norton & Company, New York, USA, 1992.

M.L. Weitzman. On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics*, 91(1), 2009.