Essays on optimal experimentation

Willems, T.

Publication date
2012

Chapter 2

Optimal Learning on Climate Change: Why Climate Skeptics Should Reduce Emissions

Uncertainty is worse than knowing the truth, no matter how bad.

The Magazine of Wall Street (1929)

2.1 Introduction

Although most climatologists agree that the upward trend in the average temperature of the Earth's atmosphere is driven by man-made emissions of greenhouse gases, policy makers, as well as many members of the public, are less convinced of the link with human activities: according to a 2007/8 survey conducted by Gallup Polls, 97 percent of all US adult citizens say to be aware of global warming, but only 49 percent of them believe that it is anthropogenic. Similarly, in 2006, the US president of that time (George W. Bush) expressed his concerns on global warming, but simultaneously stated that "there is a debate over whether it is man-made or naturally caused".[1] Corresponding views can be heard among Chinese policy makers: there, Xie Zhenhua (China's lead negotiator in the last three UN Climate Change Conferences) has said that China is keeping an "open attitude" on

[1] See http://georgewbush-whitehouse.archives.gov/news/releases/2006/06/20060626-2.html.

the causes of global warming.[2] Other examples of countries that have had openly climate skeptic leaders in the recent past include Australia (in the form of their former (1996-2007) prime minister John Howard), the Czech Republic (their current president Václav Klaus) and Russia (Vladimir Putin). More generally, virtually all countries have their climate skeptical political parties, Members of Parliament, et cetera.

Since the climate skeptic position is so widely represented in reality, this paper investigates what the optimal policy for these skeptics actually looks like.[3] In practice, skeptical policy makers typically argue that the possibility that global warming is caused by exogenous factors (such as increases in solar activity) implies that we should not take additional action towards reducing greenhouse gas emissions: Australia for example has a political party named "No Carbon Tax Climate Skeptics", while Mitt Romney (currently running to become the Republican candidate for the 2012 US elections) has made similar arguments. In October 2011 he stated that "we do not know what is causing climate change on this planet" and that "the idea of spending trillions and trillions of dollars to try to reduce CO2 emissions is not the right course for us".[4] Since this position opposes that of "climate believers" (who are convinced of the anthropogenic nature of climate change and therefore argue in favor of emission reductions), these contrasting views on the causes of global warming have led to a fierce policy debate.

Although the argument that uncertainty on the causes of global warming weakens the case for emission reductions seems to make intuitive sense at first sight, this paper shows that it is fallacious. More specifically, it is shown that the policy implications of the different positions in the climate debate actually coincide. The reason is that the typical argument put forward by skeptics neglects the production of information and the accompanying learning process on how our climate functions. Once this learning process is taken into account, it is shown that even skeptical policy makers have an incentive to reduce greenhouse gas emissions rel-

[2] See "Skeptics turn up the heat" in the *China Daily* of February 2, 2010. The article also mentions a Chinese study that debates the link between global warming and human activity.

[3] I define a climate skeptic as someone who is uncertain on the causes of global warming. Note that this is different from someone who is 100 percent certain that global warming is exogenous. This extremer position is relatively rare among policy makers (*cf.* the aforementioned quotes, where they all talk about "a debate" or "an open attitude").

[4] See "Mitt Romney Embraces Climate Denial: 'We Don't Know What's Causing Climate Change'" in *The Huffington Post* of October 28, 2011.

ative to current levels - the reason being that current emission levels are in their eyes apparently not informative enough on the nature of climate change.

In such a situation, reducing greenhouse gas emissions facilitates the learning process on whether global warming is anthropogenic.[5] Since the optimal policy differs depending on whether this is the case or not, the causes of global warming are valuable to know. Although an equally-sized *increase* in emissions could be just as informative, that strategy suffers from the fact that emitting these gases is irreversible. Consequently, there is the risk that decision makers cannot adjust their policy to the additional knowledge that they have acquired over time - thereby rendering this information useless. After all, if we (after consciously emitting more greenhouse gases) learn that global warming indeed is caused by this channel, we cannot undo the previous policy by removing the gases (of which we have by then found out that they are harmful) from the atmosphere.

A cautious policy on the other hand leaves all options open and is therefore more robust to misspecification: if this policy teaches us that there indeed is a link between the atmospheric stock of greenhouse gases and global temperature, the prudence was justified. On the other hand, we can always increase future consumption of greenhouse gases if the cautious policy tells us that there is no such link. So under the cautious strategy, the information that is produced over time is actually useful in a sense that decision makers can improve their future actions by exploiting this information.

Summarizing, this paper thus shows that uncertainty on the impact of a control variable gives agents an incentive to experiment by taking somewhat more extreme actions. Simultaneously, the irreversibility induces them to experiment "in the safe direction". In this sense, the present paper also develops a theory of optimal experimentation under irreversibilities (an issue that has not been studied before to the best of my knowledge), which may be of independent interest.

At this stage, I wish to emphasize that this paper intentionally abstracts from political-strategic considerations: I take policy makers for their word if they claim to be climate skeptic and assume that they are genuinely uncertain on the causes of climate change (rather than pretending to be climate skeptic because that might

[5] Do note that this learning process will take some time due to the lag with which the average temperature of the Earth's atmosphere responds to human activities. I will return to this issue in Sections 3 and 5.

be easier to sell to voters). Subsequently, I ask what the optimal policy for such skeptics looks like, and find that even they should reduce emissions relative to current levels. Ironically, this reduces the political attractiveness of the climate skeptic position, while it simultaneously brings consensus in the policy debate.

The remainder of this paper is structured as follows. After linking this paper to the existing literature in Section 2, Section 3 will illustrate that the learning process on the causes of global warming is facilitated by emitting either more or less greenhouse gases relative to some "confounding" (uninformative) emission level. Section 4 will then show that an optimizing agent who is faced with the irreversibilities related to the emission of greenhouse gases, prefers to experiment by *reducing* emissions. Section 5 discusses this result and its implications, after which Section 6 concludes.

2.2 Related literature

This paper applies the concept of "active learning" (also referred to as "optimal learning" or "optimal experimentation") to the climate change debate. The idea of this concept is that a decision maker optimally balances the trade-off between estimation and control of a system. In particular, active learners realize that they are learning from self-generated observations. Consequently, they optimally take the production of information into account when setting their control variable. This is what is typically referred to as the "experimentation effect".

Active learning was first developed in the engineering literature (where it is known as "dual control") and was subsequently brought to the economic sciences by MacRae (1972) and Prescott (1972). Since then, it has for example been applied to a monopolist who wants to learn his demand curve (Rothschild, 1974; Aghion *et al.* (1991)), experimental consumption of medicinal products (Grossman, Kihlstrom and Mirman, 1977), as well as to a monetary authority who wants to learn the relevant parameters of an economy it tries to control (Bertocchi and Spagat, 1993).

This paper analyzes how the learning process on the causes of global warming can be optimized. In this sense it also relates to Kelly and Kolstad (1999), as they investigate how active management of greenhouse gas emissions can teach us more on the functioning of our climate. In their paper, however, the decision maker is already convinced of the anthropogenic nature of global warming. Consequently, Kelly and Kolstad's decision maker only wants to increase the precision of his estimate of the sensitivity of global temperature to the atmospheric stock of greenhouse gases. Motivated by the observation that many policy makers in practice still question the link between global warming and human activities (recall the Introduction), this paper analyzes the problem of a policy maker who keeps an open mind to the possibility that global warming is *not* caused by the emission of greenhouse gases.

Next to the active learning literature, this paper also builds upon studies that have investigated the consequences of irreversibilities in environmental settings. Starting with the seminal contributions of Arrow and Fisher (1974) and Henry (1974), many papers have pointed out that there exists a so-called "quasi-option value" to maintaining flexibility if the quality of information increases over time (see Heal and Kriström (2002) for a survey of this voluminous literature).

The difference between the present paper and the traditional irreversibility literature is twofold. First and foremost, these earlier contributions assume that information arrives exogenously with the passage of time, while the key of this paper is that the acquisition of information is endogenized. Second, papers along the Arrow-Fisher-Henry lines typically compare the optimal decision rule under exogenous learning with the decision rule that would be optimal in the absence of learning and investigate whether the prospect of learning leads to more cautious first period decisions. This paper, on the other hand, compares a skeptic's optimal active learning rule to his passive learning rule (*i.e.* the rule that neglects the experimentation motive), as well as to the policy rule of a decision maker who is convinced that global warming is anthropogenic.

2.3 Learning the causes of climate change

To see how the active learning process on the causes of climate change evolves, this section develops a simple and tractable learning model that will be employed in the policy maker's decision problem in Section 4.

Say that the change in global temperature τ from period $t-1$ to period t

($\Delta\tau_t \equiv \tau_t - \tau_{t-1}$) is given by:[6]

$$\Delta\tau_t = \alpha + \beta c_t + \varepsilon_t \tag{2.1}$$

Here, c_t is the period t emission of greenhouse gases, while the intercept α represents the possibility that global temperature is increasing because of exogenous factors unrelated to greenhouse gas emissions. Hence, α represents potential long run trends in forces that might affect the Earth's climate. Examples are the forces that have created ice ages in the past (whatever they might have been caused by - there is no consensus on that yet) or a trend in solar activity (*cf.* Solanki *et al.* (2004) who argue that solar activity has been exceptionally high over the past 70 years). The slope parameter β on the other hand captures the possible relationship between changes in global temperature and greenhouse gas emissions. Finally, ε_t is a noise term representing other shorter-lived phenomena that temporarily affect global temperature (such as a hot summer in Latin America). It is assumed that the noise ε is i.i.d. and that it is drawn from a uniform distribution with known, bounded support, *i.e.* $\varepsilon \sim U\left[-\overline{\varepsilon}, \overline{\varepsilon}\right]$. As we will see later on, this assumption implies that learning is discrete (either nothing is learned, or the full truth is learned), which delivers analytical convenience without losing generality (*cf.* Bertocchi and Spagat (1993), from whom I copy this approach). Relaxing this assumption would not affect the experimentation incentives (see the Appendix), but doing so greatly increases analytical complexity because of the nasty form Bayesian updating then takes.

The key is that the true values of both α and β are unobserved. It is however known that there are only two possible states of the world (let us refer to the accompanying parameter values as θ_1 and θ_2). State 1 represents the IPCC scenario in which increases in global temperature are driven by the emission of greenhouse gases, while there is no upward trend due to exogenous factors (such as solar activity). Hence, $\theta_1 = \left\{\alpha = 0, \beta = \overline{\beta}\right\}$, with $\overline{\beta} > 0$.

In state 2 on the other hand, the upward trend in global temperature is completely exogenous to human behavior and the emission of greenhouse gases does

[6] Integrating (1) leads to $\tau_t = \tau_0 + t\alpha + \beta \sum_{j=1}^{t} c_j$, where $\sum_{j=1}^{t} c_j$ is the atmospheric stock of greenhouse gases at time t (potentially relative to a certain base level c_0, which I have normalized to zero for convenience). This expression for the time t temperature level is relatively standard in the literature (up to the inclusion of the potential exogenous trend α), *cf.* equation (1a) of Kelly and Kolstad (1999).

not play a role. Hence, $\theta_2 = \{\alpha = \overline{\alpha}, \beta = 0\}$, with $\overline{\alpha} > 0$. We thus have:

$$\begin{aligned} \Delta\tau_t(\theta_1) &= \overline{\beta}c_t + \varepsilon_t \\ \Delta\tau_t(\theta_2) &= \overline{\alpha} + \varepsilon_t \end{aligned}$$

The remainder of this section shows that these two cases are easier to distinguish from each other for more extreme emission levels c_t. Intuitively, in the limiting case where we stop emitting greenhouse gases altogether (*i.e.* set $c_{t+k} = 0 \ \ \forall k > 0$), the upward trend in global temperature should slow down if the IPCC is right (as illustrated in Figure 1 of Solomon *et al.* (2009)). On the other hand: if the Earth continues to heat up at a constant rate well after we have stopped emitting greenhouse gases, global warming is likely to be driven by exogenous factors.

A similar reasoning can be followed for upward experimentation: if temperature increases $\Delta\tau$ do not accelerate after pumping a lot of greenhouse gases into the atmosphere, the IPCC is probably wrong, while the IPCC's case is supported if the Earth's temperature would start to increase at a faster rate.

Do note that this learning process is likely to take a while, due to the lags associated with global temperature dynamics: typically, it takes about 10 years before changes in greenhouse gas emissions start to have noticeable effects on global temperature (Heal and Kriström, 2002: 11). Consequently, a period in the model should be thought of as about a decade in reality. As pointed out in Wigley (2005), it takes much longer (perhaps as long as a century) before the effect of a higher atmospheric stock of greenhouse gases is *fully* incorporated into the Earth's temperature level due to the fact that it takes time to heat up the oceans (the so-called "ocean thermal lag"). However, after about 10 years the initial first-round effect becomes visible.

Moreover, as argued by Solomon *et al.* (2009), it can take even as long as 1,000 years before global temperature actually starts falling significantly if man-made emissions were to cease immediately. However, this is not what we need to wait for in order to see the IPCC-case confirmed. For that, any response in temperature dynamics following a change in emission levels (such as a change in the *rate* at which temperature increases, as visible in the simulation underlying Figure 1 of Solomon *et al.* (2009)) would do.

The workings of such an "active learning by doing" strategy have already been illustrated in the past, namely in the acid rain debate:[7] during the 1960s, an increasing number of streams and lakes in Norway were reported to be acidic. Initially, it was debated whether these changes were anthropogenic or not, but nevertheless governments decided to reduce sulfur emissions in a drastic manner. After some years (again due to the time lag with which nature tends to respond to human activities), the level of acidity in Norwegian waters fell, as a result of which there is nowadays little doubt left that the changes were in fact man-made. Consequently, we are still careful with emitting sulfur today, so the aggressive sulfur reduction policies followed by governments in the 1970s and 80s did produce valuable information.[8]

To formalize this reasoning for the global warming case, we can exploit the assumption that the noise term ε has bounded support (*i.e.* that $\varepsilon \in [-\overline{\varepsilon}, \overline{\varepsilon}]$). Consequently, we know with certainty that:

$$\Delta\tau_t(\theta_1) \in \left[\overline{\beta}c_t - \overline{\varepsilon}, \overline{\beta}c_t + \overline{\varepsilon}\right] \tag{2.2}$$

$$\Delta\tau_t(\theta_2) \in \left[\overline{\alpha} - \overline{\varepsilon}, \overline{\alpha} + \overline{\varepsilon}\right] \tag{2.3}$$

Graphically, this can be visualized as in Figure 1.

Now the key is to observe that learning is complete if $c_t < c^*$ or $c_t > c^{**}$. In those cases, the regions for $\Delta\tau_t(\theta_1)$ and $\Delta\tau_t(\theta_2)$ are non-overlapping as a result of which we immediately find out which state we are in.

We can actually derive analytical expressions for these cut-off values. In particular, c^* is defined by:

$$\overline{\beta}c^* + \overline{\varepsilon} = \overline{\alpha} - \overline{\varepsilon} \Leftrightarrow c^* = \frac{\overline{\alpha} - 2\overline{\varepsilon}}{\overline{\beta}},$$

while c^{**} solves:

$$\overline{\beta}c^{**} - \overline{\varepsilon} = \overline{\alpha} + \overline{\varepsilon} \Leftrightarrow c^{**} = \frac{\overline{\alpha} + 2\overline{\varepsilon}}{\overline{\beta}}$$

[7] See Seip (2001) for an overview of this debate with a particular focus on Norway.

[8] In principle, the same information could have been produced by *increasing* sulfur emissions, but that strategy would have caused irreversible damage to Norwegian nature - thereby making the acquired information less valuable (since there would have been less nature left to preserve). As will be shown in Section 4, this line of reasoning makes the policy maker want to experiment "in the safe direction".

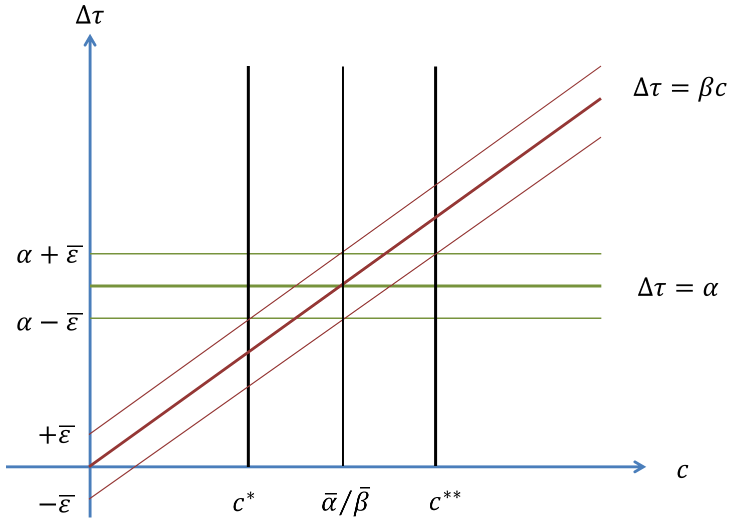

Figure 2.1: Graphical representation of the learning process

In the overlap region for which $c_t \in [c^*, c^{**}]$, learning is probabilistic: if you are lucky enough to receive a $(\Delta\tau_t, c_t)$-observation that lies outside the range of either $\Delta\tau_t(\theta_1)$ or $\Delta\tau_t(\theta_2)$ (as defined by (2) and (3)), you learn the full truth. Exploiting the uniformity on ε, the probability of learning the truth for all $c_t \in [c^*, c^{**}]$ can be shown to equal:

$$P = \frac{\left|\overline{\alpha} - \overline{\beta} c_t\right|}{2\overline{\varepsilon}} \tag{2.4}$$

Note from equation (4) that $P = 1$ if either c^* or c^{**} is chosen (so you always learn the complete truth if you choose the revealing emission levels (which is consistent with the way they are defined)), while $P = 0$ for $c = \overline{\alpha}/\overline{\beta}$. Hence, for $c_t = \overline{\alpha}/\overline{\beta}$ the endogenous and exogenous warming case are indistinguishable. The reason is that this is the so-called "confounding" emission level (*i.e.* that level of emissions where the two lines intersect). So from an informational point of view, this is the emission level you want to avoid, as it does not teach you anything about how our climate functions.

Also observe that the function P is symmetric around this confounding emission level: the production of information only depends on $\left|c_t - \overline{\alpha}/\overline{\beta}\right|$ (not on $sgn\left(c_t - \overline{\alpha}/\overline{\beta}\right)$), so positive deviations of c_t from $\overline{\alpha}/\overline{\beta}$ are as informative as negative ones.

The fact that positive deviations of c_t from $\overline{\alpha}/\overline{\beta}$ are *exactly* as informative as negative ones is due to the assumption that equation (1) is linear. Although this assumption is relatively standard in the literature (*cf.* footnote 6), assuming that (1) is non-linear could make a difference. In this respect, climate studies suggest that the temperature response to an emission reduction might be faster than that to an increase in emissions (due to the fact that oceans tend to cool down more rapidly than they warm up).[9] Since this non-linearity increases the attractiveness of downward experimentation relative to upward experimentation for reasons other than the irreversibility (namely learning speed), this modification would only strengthen the findings that this paper will arrive at.

2.4 Optimizing model

Now that we are familiar with the active learning process on the causes of climate change, we can investigate how this process affects decisions related to the emission of greenhouse gases. I do this through a simple three period model, in which the experimentation motive is going to interact with the fact that emitting greenhouse gases is irreversible. In the model, period 0 characterizes the current state of affairs. Given this current state, the present paper asks what the *optimal* policy for a climate skeptic looks like. That is: what kind of policy would be implemented by a social planner who is uncertain on the causes of global warming if he were to take over power at the beginning of period 1? In this respect, period 1 represents the learning phase, while period 2 captures the remaining future.

In particular, I consider a decision maker who derives utility from the consumption of greenhouse gases c_t in each period t. Upon consuming a unit of c, it is emitted into the Earth's atmosphere. Consumption of greenhouse gases is free, but the presence of these gases in the atmosphere may prove to be harmful in the future (if it turns out to cause global warming). Furthermore, greenhouse gas consumption is irreversible: once emitted, it is not possible to remove greenhouse

[9] See Stouffer (2004): any development that heats the ocean surface makes the oceans more stable, isolating the deeper waters from the surface. Any development that tends to cool the ocean surface on the other hand, makes the oceans more unstable - thereby promoting mixing between the surface and deeper (cooler) waters. As a result, cooling the oceans down is easier than warming them up. Given the tight link between oceanic temperatures and the average temperature of the Earth's atmosphere, climatologists believe that the atmospheric temperature response to an emission decrease is faster than that to an increase in emissions.

gases from the atmosphere again (apart from natural decay).

At the beginning of period 0, the Earth is endowed with an exhaustible stock G of greenhouse gases that resides beneath the Earth's surface (think of this as oil *in situ*). Since this stock is depleted over time, future consumption choices are constrained by the amount of gases consumed in the past. In particular:

$$c_0 \in [0, G] \tag{2.5}$$

$$c_1 \in [0, G - c_0] \tag{2.6}$$

$$c_2 \in [0, G - c_0 - c_1] \tag{2.7}$$

I assume that the pay-off function (the sum of the utility generated by the consumption of greenhouse gases and the disutility of global warming) is concave. For concreteness I will take it to be logarithmic in greenhouse gas consumption and linear in temperature changes, but as long as their sum is concave other specifications would yield the same qualitative result. Abstracting from discounting and natural decay of atmospheric greenhouse gases to lighten notation, the pay-off function (conditional on the realizations of α, β and the ε's, indicated by $\widehat{\alpha}$, $\widehat{\beta}$ and $\widehat{\varepsilon}$ respectively) is given by:[10]

$$\begin{aligned} U(c_0, c_1, c_2; \widehat{\alpha}, \widehat{\beta}, \widehat{\varepsilon}) &= \sum_{t=0}^{2} [\log(c_t) - \Delta\tau_t] \\ &= \sum_{t=0}^{2} \log(c_t) - 3\widehat{\alpha} - \widehat{\beta} \sum_{t=0}^{2} c_t - \sum_{t=0}^{2} \widehat{\varepsilon}_t \end{aligned}$$

How and whether the decision maker finds out which state we are actually in (*i.e.* what the true values of α and β are), will be differentiated in the following subsections. In particular, those subsections will show how different assumptions on the model's information structure affect the optimal decision rule.

[10] In this specification all three periods receive the same weight. Since the final period is meant to capture the infinite future, that period should actually obtain a larger weight (for example proportional to $1/(1-\delta)$, with δ representing the discount factor). However, as this would only affect the optimal *level* of experimentation (which is an issue that is not analyzed in this paper, since it would require a more realistic model to start from), I neglect this for analytical convenience.

2.4.1 Characterization of past policy

In order to be able to compare the optimal climate policy to actual policies followed by most countries in the recent past, we first need to characterize the latter in terms of our stylized model.

I believe that past policy has been characterized by two notable features. First, the possibility that future decision makers may have better information on the causes of global warming does not seem to have influenced recent policy. The reason for this is probably political: because the learning process on the causes of climate change takes a while, the benefits of better future information are likely to accrue to a different group of policy makers. Consequently, current policy makers neglect this positive informational spillover when determining current policy.

Second, three of the most important countries when it comes to emitting CO2 (China, Russia and the United States[11]), all have (or had in the recent past) climate skeptical policy makers in office (recall the Introduction to this paper). In addition, prior to the increased scientific consensus on the anthropogenic nature of global warming (which was achieved in the 1970s and communicated to a wider audience through the first IPCC report in 1990), basically all governments determined climate policy while being uncertain on the costs of emitting greenhouse gases (just like climate skeptics claim to be nowadays).

Consequently, it seems reasonable to assume that recent policy has been shaped by skeptical decision makers who neglected the positive informational spillover. If we denote their belief that global warming is endogenous by $\theta \in (0,1)$, their problem reads:

$$\begin{aligned} & \max_{c_0,c_1,c_2} \log(c_0)+\log(c_1)+\log(c_2)-3\mathbb{E}_0\{\alpha\}-\mathbb{E}_0\{\beta\}[c_0+c_1+c_2] \\ = & \max_{c_0,c_1,c_2} \log(c_0)+\log(c_1)+\log(c_2)-3(1-\theta)\overline{\alpha}-\theta\overline{\beta}[c_0+c_1+c_2], \end{aligned}$$

subject to (5)-(7).

The accompanying first-order conditions are given by:

$$\frac{1}{c_t}-\theta\overline{\beta}=0 \quad \text{for } t=0,1,2, \tag{2.8}$$

[11] According to the United States Department of Energy's Carbon Dioxide Information Analysis Center, these three countries are responsible for about 50 percent of all anthropogenic CO2 emissions.

Consequently, the solution to this problem for the current period ("period 0") reads:

$$c_0^{\Diamond} = \frac{1}{\theta\overline{\beta}} \tag{2.9}$$

Now, if climate skeptics want to claim that they are genuinely uncertain on the causes of global warming (rather than just trying to be politically cunning), it *has* to be the case that these recent emission levels are "confounding" (uninformative) in their eyes. This view is for example expressed by Todd Myers (a prominent climate skeptic). He has stated that "climate models indicate that the impact of current CO2 concentrations on the climate is slight, within the noise level in the data. In other words, according to the climate models, we are at levels in which it is hard to distinguish the CO2 impacts from natural forces".[12]

As shown in Section 3, the confounding emission level equals $\overline{\alpha}/\overline{\beta}$ so apparently the values of $\overline{\alpha}$ and $\overline{\beta}$ that climate skeptics like Todd Myers have in mind, are such that the exogenous and endogenous global warming case are difficult to distinguish from each other around current emission levels. Consequently, the following must hold (at least approximately):

$$c_0^{\Diamond} = \frac{\overline{\alpha}}{\overline{\beta}} \tag{2.10}$$

After all, if this would not be the case, climate skeptics could not claim to be genuinely skeptic.

2.4.2 Characterization of optimal policy

Current climate policy as characterized in the previous subsection is however suboptimal given the fact that learning considerations (and the associated informational spillovers) are not taken into account. Let us therefore investigate what the optimal policy looks like. That is: what policy would be implemented by a social planner (who does incorporate that the quality of information may increase over time) if he were to take over power at the beginning of period 1?

Complete certainty

First consider the extreme case of an omniscient social planner who knows at the beginning of period 1 already whether global warming is exogenous or endogenous.

[12] See "Climate Data That Sounds Meaningful...But Isn't" (published by the Washington Policy Center).

If we normalize the stock of greenhouse gases *in situ* at the beginning of the first period (at which point there is $G - c_0^{\Diamond}$ left) to 1, his problem reads:

$$\max_{c_1, c_2} \log(c_1) + \log(c_2) - 2\widehat{\alpha} - \widehat{\beta}\left[c_1 + c_2\right],$$

subject to (6) and (7).

Here, $\widehat{\alpha} = 0$ and $\widehat{\beta} = \overline{\beta}$ ($\widehat{\alpha} = \overline{\alpha}$ and $\widehat{\beta} = 0$) for a social planner who knows that global warming is endogenous (exogenous). The first-order conditions then imply:

$$c_t^{CC} = \begin{cases} \frac{1}{2} & \text{if warming is exogenous} \\ \frac{1}{\overline{\beta}} & \text{if warming is endogenous} \end{cases} \quad \text{for } t = 1, 2, \tag{2.11}$$

where the superscript "CC" indicates that this is the solution under complete certainty.[13]

Skepticism, passive learning

As set out in the Introduction to this paper, many policy makers in practice do not adhere one of the extreme positions covered by (11), but hold a skeptical attitude towards the causes of global warming instead.

Consider therefore a skeptical social planner who does not know the true values of α and β at the beginning of period 1, but who knows that he will get to observe these values with probability P at the beginning of period 2. Assuming that the costs of global warming are incurred in the final period (so that they are contained in the indirect utility functions $V(\cdot)$, which is analytically convenient), the optimization problem can be written as:

$$\max_{c_1, c_2} \log(c_1) + P\mathbb{E}_1\left\{V(c_2^{SL})\right\} + [1 - P]\,\mathbb{E}_1\left\{V(c_2^{SNL})\right\},$$

subject to (6) and (7).

Here, the fact that learning is discrete (as a result of which Bayesian updating

[13] This solution implicitly assumes that $\overline{\beta} > 2$. For $\overline{\beta} \leq 2$, the possibility of endogenous global warming is no longer "binding" and there is no value to learning the true causes of climate change anymore. After all, in that case the optimal policy does not depend on whether global warming is endogenous or exogenous (then $c_t^{CC} = 1/2$ for $t = 1, 2$ in both states) and there is nothing to argue about for policy makers. Judging from the existence of a fierce policy debate accompanying the question which state we are actually in (exogenous or endogenous global warming?), this does not seem to be the case in reality. Consequently, I assume that $\overline{\beta} > 2$.

takes a particularly simple form) pays off in terms of tractability. After all, at the beginning of period 2 there are only two possibilities: either the decision maker has learned the true state of nature (which happens with probability P), or not. In the latter case his period 2 problem reads:

$$\max_{c_2 \in [0,1-c_1]} \log(c_2) - 2(1-\theta)\overline{\alpha} - \theta\overline{\beta}[c_1 + c_2],$$

and he sets:

$$c_2^{SNL} = \frac{1}{\theta\overline{\beta}}, \tag{2.12}$$

where "SNL" indicates that this is the solution to a skeptic's problem who has not learned the causes of global warming. Following this strategy brings him indirect utility $V(c_2^{SNL})$.

For a skeptical decision maker who has learned the causes of global warming by the end of period 1 (indicated by "SL"), the period 2 problem reads:

$$\max_{c_2 \in [0,1-c_1]} \log(c_2) - 2\widehat{\alpha} - \widehat{\beta}[c_1 + c_2]$$

Hence:

$$c_2^{SL} = \begin{cases} 1 - c_1 & \text{if warming is exogenous} \\ \frac{1}{\beta} & \text{if warming is endogenous} \end{cases}, \tag{2.13}$$

and the decision maker attains utility $V(c_2^{SL})$.

Observe that we can rewrite the objective function as:

$$\max_{c_1} \log(c_1) + P\mathbb{E}_1\left\{V(c_2^{SL}) - V(c_2^{SNL})\right\} + \mathbb{E}_1\left\{V(c_2^{SNL})\right\} \tag{2.14}$$

Here $\mathbb{E}_1\left\{V(c_2^{SL}) - V(c_2^{SNL})\right\}$ represents the expected utility gain from learning whether global warming is anthropogenic or not.

By substituting the optimal decision rules (12) and (13) into the pay-off function it follows that:

$$V(c_2^{SL}) = \begin{cases} \log(1 - c_1) - 2\overline{\alpha} & \text{if warming is exogenous} \\ -\log\left(\overline{\beta}\right) - \overline{\beta}c_1 - 1 & \text{if warming is endogenous} \end{cases}, \tag{2.15}$$

and:

$$V(c_2^{SNL}) = \begin{cases} -\log\left(\theta\overline{\beta}\right) - 2\overline{\alpha} & \text{if warming is exogenous} \\ -\log\left(\theta\overline{\beta}\right) - \overline{\beta}c_1 - \frac{1}{\theta} & \text{if warming is endogenous} \end{cases} \tag{2.16}$$

Given prior belief θ, the expected value for a decision maker who does not know what causes global warming at the beginning of period 2 reads:

$$\mathbb{E}_1\left\{V(c_2^{SNL})\right\} = -\log\left(\theta\overline{\beta}\right) - \theta\overline{\beta}c_1 - 2\left(1-\theta\right)\overline{\alpha} - 1 \tag{2.17}$$

Similarly, the expected value added of learning the causes of global warming at the beginning of period 2 equals:

$$\mathbb{E}_1\left\{V(c_2^{SL}) - V(c_2^{SNL})\right\} = \log\left(\theta\right) + \left(1-\theta\right)\left[1 + \log\left(\overline{\beta}\right) + \log\left(1-c_1\right)\right] \tag{2.18}$$

Using that $1/\theta\overline{\beta} \leq (1-c_1)$ (by feasibility) and $-\log(\theta) - 1/\theta < -1$ (since $\theta \in (0,1)$), a comparison of equations (15) and (16) shows that:

$$\mathbb{E}_1\left\{V(c_2^{SL}) - V(c_2^{SNL})\right\} > 0$$

This means that knowing whether emitting greenhouse gases is damaging or not is valuable, which is intuitive: after all, once we know which state of the world we are in, we can condition our decision and implement the optimal policy for that state. If we remain uncertain on the other hand, we cannot condition. In that case, we would have to work with some kind of average rule (in this paper's context given by (12)), which works well in expectation, but is suboptimal for either state realization.

However, do observe from (18) that:

$$\frac{\partial\mathbb{E}_1\left\{V(c_2^{SL}) - V(c_2^{SNL})\right\}}{\partial c_1} = -\frac{1-\theta}{1-c_1} \leq 0 \tag{2.19}$$

This expression is key to the results in this paper. It tells us that the expected benefit from learning the truth is decreasing in first period consumption c_1. The intuition for this is that a higher choice of c_1 reduces the action space of the decision maker in the second period, as a result of which he has less room to actually use the information that he acquired during the first period. In the limit, if the decision maker decides to consume the complete resource stock in period 1, he is very likely to find out whether global warming is endogenous or not, but he has no freedom left to exploit this information: irrespective of what the outcome of his learning process is, there is nothing left to consume.

By using (17) and (18), one can rewrite the objective function (14) as:

$$\max_{c_1} \ \log(c_1) + P\left\{\log(\theta) + (1-\theta)\left[1 + \log(\overline{\beta}) + \log(1-c_1)\right]\right\}$$
$$+\left\{-\log\left(\theta\overline{\beta}\right) - \theta\overline{\beta}c_1 - 2(1-\theta)\overline{\alpha} - 1\right\}$$

Hence, the accompanying first-order condition reads:

$$\frac{1}{c_1} - \frac{P(1-\theta)}{1-c_1} - \theta\overline{\beta} = 0, \tag{2.20}$$

which can be solved to yield:

$$c_1^{SPL} = \frac{(1-\theta)P + \theta\overline{\beta} + 1 - \sqrt{\left[(1-\theta)P + \theta\overline{\beta} + 1\right]^2 - 4\theta\overline{\beta}}}{2\theta\overline{\beta}}, \tag{2.21}$$

where "SPL" indicates that this is the solution for a skeptical policy maker who learns in a passive fashion.[14]

Skepticism, active learning

Just taking into account that additional information may arrive over time, is however not enough to make a skeptical policy maker implement the optimal policy. The reason is that last section's passive learner neglects the fact that the probability of learning the truth P is actually endogenous. In particular, the passive learner fails to realize that he is learning from self-generated observations and erroneously sees the probability of learning the truth as an exogenous constant (that is: he thinks that $dP/dc_1 = 0$). The truth is however that P is given by equation (4) and is hence a function of the first period decision c_1. Consequently, the fully rational active learner takes the effect of first period consumption choices

[14] As for any quadratic equation with a non-zero discriminant, there is a second solution but that one is not feasible in this paper's setting as it implies $c_1^{SPL} > 1$.

on the probability of learning the truth into account. The first-order condition of a fully optimizing agent therefore reads:

$$\frac{1}{c_1} + \left.\frac{dP}{dc_1}\right|_{c_1=c_1^{SPL}} \mathbb{E}_1\left\{V(c_2^{SL}) - V(c_2^{SNL})\right\} - \frac{P(1-\theta)}{1-c_1} - \theta\overline{\beta} = 0, \quad (2.22)$$

$$\text{with} \quad \frac{dP}{dc_1} = \frac{\overline{\beta}c_1 - \overline{\alpha}}{|\overline{\beta}c_1 - \overline{\alpha}|}\frac{\overline{\beta}}{2\overline{\varepsilon}}$$

This first-order condition characterizes the optimal period 1 consumption choice under the active learning rule (c_1^{SAL}). Unfortunately, it is not possible to solve this equation explicitly for c_1^{SAL}, but we can determine the direction of experimentation by comparing (22) with the first-order condition of the passive learner (20).

When doing so, it should be noted that $\mathbb{E}_1\left\{V(c_2^{SL}) - V(c_2^{SNL})\right\} > 0$ (as set out in the previous section, information is valuable), while $dP/dc_1|_{c_1=c_1^{SPL}} = -\overline{\beta}/2\overline{\varepsilon} < 0$. The reason for the latter is that $c_1^{SPL} < \overline{\alpha}/\overline{\beta}$ (this follows from comparing (20) with equations (8) and (10)), which places us to the left of the confounding emission level in Figure 1. In that part of the state space the probability of learning the truth is decreasing in emission levels, so the active learning motive induces the decision maker to *reduce* emissions. A skeptical policy maker thus has an experimentation motive, while the irreversibilities associated with emitting greenhouse gases make him want to experiment "in the safe direction" (recall the discussion following (19)).

Hence, *even* climate skeptics have an incentive to reduce emissions relative to current (in their eyes uninformative) emission levels. A failure to do so can be explained by political motives (such as the fact that future generations, who would benefit if we were to introduce some policy experimentation, cannot vote), but is hard to justify upon grounds of optimality.

So once learning considerations are taken into account, the policy implications of the various positions in the climate debate turn out to be surprisingly consistent, thereby reducing the relevance of this (currently heated) debate: both "climate believers" as well as climate skeptics should argue in favor of a more cautious policy. The former, trivially, because they are convinced that emitting greenhouse gases is damaging (which is/was not taken into account by many current/recent policy makers), while the latter (who apparently find current emission levels not informative enough on the causes of climate change) should do so for learning considerations.[15]

In fact, if the value to knowing the truth (given by $\mathbb{E}_1 \left\{ V(c_2^{SL}) - V(c_2^{SNL}) \right\}$) is large, it can even be the case that $c_1^{SAL} < 1/\overline{\beta}$ (the emission level set by a social planner who is convinced that global warming is anthropogenic; *cf.* (11)). Hence, it is even possible that a skeptical policy maker should actually argue for a *lower* emission level than a policy maker who is convinced that global warming is anthropogenic. Whether this is the case for reasonable calibrations should be investigated with more sophisticated models of climate change, but it is a theoretical possibility.

2.5 Discussion

Central to this paper is the learning process on the causes of climate change. As pointed out in Section 3, this process is likely to take some time, since it takes about 10 years before changes in greenhouse gas emission levels start to have a noticeable effect on global temperature. Do note in this respect that I did not have to assume anything on the length of a period in my model. As long as *some* information is produced over time by more extreme emission levels, and as long as society attaches an epsilon-positive weight to the well-being of future generations, a certain extent of policy experimentation becomes optimal - no matter how long the learning process takes. The actual degree of experimentation would of course be affected (Should emissions be reduced by 5% or by 50%? And for how long?), but to give a serious answer to those questions takes a less stylized model (and is perhaps more a task for climatologists rather than for economists). This paper only intends to point out the direction in which current policy should move if we want to get it closer to the optimum emission level.

Section 4 also showed that calls against greenhouse gas reductions by climate skeptics are difficult to rationalize in an optimizing framework unless one refers to myopia or strategic motives. It would therefore be interesting to consider the politico-economic aspects of climate politics (such as lobbying and intergenerational issues), the analysis of which I leave for future work. In this respect, the

[15] Only policy makers who are 100 percent sure that global warming is exogenous would implement the optimal policy (given their model) by increasing emissions. However, as noted in footnote 3, this is not what policy makers typically claim in practice.

aforementioned acid rain case deserves closer study, as policy makers were able to set strategic issues and myopia aside to combat that problem in a way that looks very much like the "active learning by doing" strategy proposed by this paper.

Another issue is *how* emissions ought to be reduced. Broadly speaking, there are two options: implement a carbon tax and leave it to the market, or adopt regulation that for example requires the installation of abatement capital. Next to the general objections that exist against regulation-based solutions (such as the fact that firms may find a way to get around it), Kolstad (1996) presents a strong case against the forced installation of abatement capital. He notes that this requires irreversible investment, which may turn out to be wasted if global warming turns out to be exogenous. This introduces a second option value that makes emission reduction less attractive. In this respect, implementing a carbon tax brings two important advantages: firstly, it does not introduce any direct irreversibilities, and secondly it gives firms complete freedom as to how to respond to the changed incentives. If Kolstad's option value indeed is important in practice, firms may choose to abstract from installing irreversible abatement capital and they could pass the carbon tax on to consumers via prices (directly leading to less greenhouse gas consumption).

For policy makers who are convinced that the IPCC scenario is the truth, the next step may be to increase the precision of their estimate of the responsiveness of global temperature to greenhouse gas emissions. As argued by Kelly and Kolstad (1999), this process is facilitated by *increasing* emissions. Given the irreversible nature of this strategy, experimenting by pumping more gases into the atmosphere is pretty risky since this responsiveness may turn out to be high. (For the same reason, this paper's skeptical policy maker prefers to experiment "in the safe direction", *i.e.* by reducing emissions.) Consequently, one may want to take model uncertainty and explicit concerns for robustness into account when analyzing the question of how to learn the so-called "climate sensitivity parameter" optimally. In this respect, Cogley *et al.* (2008) have already shown that these considerations tend to decrease experimentation incentives for a monetary authority who is trying to learn the Phillips curve in an active manner and it would be interesting to see whether this result carries over to the environmental setup.

Finally, although this paper is written with the climate change debate in mind, the underlying idea could also be applied to other environmental problems. As pointed out in Section 3, the workings of the "active learning by doing" strategy have already been illustrated in practice in the acid rain debate, while this approach could also be of value to future problems that are yet to arise. More generally, this paper has developed a theory of optimal experimentation under irreversibilities, which may have applications outside environmental economics. One application that comes to mind is optimal consumer experimentation with potentially addictive goods/experiences, such as hard drugs, Facebook or Apple products.

2.6 Conclusion

Given the popularity of the climate skeptic position among policy makers, this paper has analyzed what the optimal climate policy for these skeptics looks like. Typically, climate skeptics argue that the possibility that global warming is exogenous implies that it is optimal *not* take additional action towards reducing greenhouse gas emissions. This paper has however shown that the conclusion should actually go in the other direction: if climate skeptics want to implement the optimal policy, they should argue for lower emissions relative to current levels. Hence, unless one refers to myopia or political motives, calls of climate skeptics against reductions in greenhouse gas emission levels are hard to justify upon grounds of optimality.

The reason is that if climate skeptics are genuinely uncertain on the causes of climate change, they apparently find the available emission/temperature data not informative enough on this issue. Since there is a positive value to knowing the true causes of global warming, this implies that these skeptical policy makers should argue for a policy that moves us away from current (in their eyes uninformative) emission levels. They can do this by either in- or decreasing greenhouse gas emissions, but the option of upward experimentation is unattractive since the irreversibilities associated with emitting greenhouse gases erode the value added of the information produced. Hence, uncertainty on the causes of global warming gives skeptical decision makers an incentive to experiment, while the irreversible nature of emitting greenhouse gases induces them to experiment "in the safe direction". Theoretically, it is even possible that the experimentation motive is so strong that a skeptical policy maker should argue in favor of *lower* emissions than a policy maker who is convinced of the anthropogenic nature of global warming.

So once learning considerations are taken into account, the heated question whether one is a "climate believer" or a climate skeptic becomes of subordinate importance from a policy point of view. After all, the policy implications of the different positions turn out to be surprisingly consistent: both "climate believers" as well as climate skeptics should argue in favor of a more cautious climate policy. The former, trivially, because they are convinced that emitting greenhouse gases is damaging (which is/was not taken into account by most current/recent policy makers), while the latter (who apparently find current emission levels not informative enough on the causes of climate change) should do so for learning considerations.

2.7 Appendix

In the main text it was assumed that the noise term ε has bounded support. As a result, leaning was discrete and Bayesian updating took a particularly simple form - thereby maintaining analytical tractability. This appendix shows that the fact that extremer emission levels facilitate the learning process on the causes of global warming (with negative deviations from the confounding emission level being as informative as positive ones) continues to hold in a model where the noise term ε has infinite support. Assuming that ε_t is Gaussian, the model is given by:

$$\begin{aligned} \Delta\tau_t &= \alpha + \beta c_t + \varepsilon_t \\ \varepsilon_t &\sim N(0, \sigma_\varepsilon^2) \end{aligned}$$

As before, there are two possible states: either the IPCC is right (in which case $\alpha = 0$ and $\beta = \overline{\beta} > 0$), or the skeptics are right (which implies $\alpha = \overline{\alpha} > 0$ and $\beta = 0$). Consequently, π_t (the time t belief that the IPCC is right) is equal to the relative probability of observing a particular $(\Delta\tau_t, c_t)$-observation under that regime. Bayes' rule now implies that:

$$\pi_t = \frac{\pi_{t-1} \exp\left(-0.5 \frac{\left[\Delta\tau_t - \overline{\beta} c_t\right]^2}{\sigma_\varepsilon^2}\right)}{\pi_{t-1} \exp\left(-0.5 \frac{\left[\Delta\tau_t - \overline{\beta} c_t\right]^2}{\sigma_\varepsilon^2}\right) + [1 - \pi_{t-1}] \exp\left(-0.5 \frac{[\Delta\tau_t - \overline{\alpha}]^2}{\sigma_\varepsilon^2}\right)} \tag{A1}$$

From (A1), it is easily verified that $\pi_t = \pi_{t-1}$ if the decision maker sets $c_t = \overline{\alpha}/\overline{\beta}$. Hence, at the confounding emission level $\overline{\alpha}/\overline{\beta}$, no information is produced as to how our climate functions, as a result of which beliefs cannot be updated. Any deviation of c_t from $\overline{\alpha}/\overline{\beta}$ does produce valuable information - thereby enabling agents to update their beliefs. Since the term $\left[\Delta\tau_t - \overline{\beta} c_t\right]$ only enters (A1) in a squared fashion, the direction of the deviation does not matter.