

van Wijnbergen, Sweder; Willems, Tim

Working Paper

Optimal Learning on Climate Change: Why Climate Skeptics should reduce Emissions

Tinbergen Institute Discussion Paper, No. 12-085/2

Provided in Cooperation with:

Tinbergen Institute, Amsterdam and Rotterdam

Suggested Citation: van Wijnbergen, Sweder; Willems, Tim (2012) : Optimal Learning on Climate Change: Why Climate Skeptics should reduce Emissions, Tinbergen Institute Discussion Paper, No. 12-085/2, Tinbergen Institute, Amsterdam and Rotterdam, <http://nbn-resolving.de/urn:nbn:nl:ui:29-442084>

This Version is available at:

<http://hdl.handle.net/10419/87345>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

TI 2012-085/2
Tinbergen Institute Discussion Paper



Optimal Learning on Climate Change: Why Climate Skeptics should reduce Emissions

Sweder van Wijnbergen¹

Tim Willems²

¹ Faculty of Economics and Business, University of Amsterdam, and Tinbergen Institute;

² Oxford University.

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

More TI discussion papers can be downloaded at <http://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 1600

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Duisenberg school of finance is a collaboration of the Dutch financial sector and universities, with the ambition to support innovative research and offer top quality academic education in core areas of finance.

DSF research papers can be downloaded at: <http://www.dsf.nl/>

Duisenberg school of finance
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 8579

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

Sweder van Wijnbergen[†] and Tim Willems[‡]

August 16, 2012

Abstract

Climate skeptics argue that the possibility that global warming is exogenous implies that we should not take additional action towards reducing greenhouse gas emissions until we know more. This paper however shows that even climate skeptics have an incentive to reduce emissions: such a change of direction facilitates their learning process on the causes of global warming. Since the optimal policy action depends on these causes, they are valuable to know. Although an increase in emissions would also ease learning, that option is inferior due to the fact that emitting greenhouse gases is irreversible. Consequently, the policy implications of the different positions in the global warming debate turn out to coincide - thereby diminishing the relevance of this debate from a policy perspective. Uncertainty is no reason for inaction.

JEL-classification: D83, Q54, Q58

Key words: climate policy, global warming, climate skepticism, active learning, irreversibilities

*We thank Christian Gollier, Mark Kagan, Shaun Larcom, Radek Stefanski and Cees Withagen for useful comments and discussions. Comments by John van Boxel and Bart Verheggen greatly improved the climate scientific side of this paper. Any errors are naturally ours.

[†]Tinbergen Institute and Department of Economics, University of Amsterdam, Valckenierstraat 65-67, 1018 XE Amsterdam, The Netherlands. E-mail: s.j.g.vanwijnbergen@uva.nl.

[‡]Nuffield College and Department of Economics, University of Oxford, Manor Road Building, Manor Road, Oxford OX1 3UQ, UK. E-mail: tim.willems@economics.ox.ac.uk.

1 Introduction

Although most climatologists agree that the upward trend in the average temperature of the Earth's atmosphere is mainly 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 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 provides a normative analysis of what the optimal policy for these skeptics actually looks like. We define a climate skeptic as someone who is uncertain on the causes of global warming.³ In practice, these skeptics typically propose a rather passive policy by arguing that the possibility that global warming is driven 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 (the Republican candidate for the 2012 US elections) has made similar arguments. In

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

²See "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.

³Note 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").

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 "IPCC 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 as it 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 uncertainty is not a reason for inaction, but a powerful motive for action instead. In particular, it is shown that even skeptical policy makers obtain an incentive to reduce greenhouse gas emissions relative to current levels - the reason being that current emission levels are in their eyes apparently not informative enough on the nature of climate change.

When faced with uncertainty on the relationship between a control variable (such as emissions) and a dependent variable (such as global temperature), decision makers obtain an incentive to implement a policy change (like reducing emissions). The reason is that such a change of direction facilitates the learning process on whether global warming is indeed related to the emission of greenhouse gases.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 information that they have produced over time - thereby rendering this information useless. After all, if we (after consciously emitting more greenhouse gases) learn that

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

⁵Do 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. We will return to this issue in Sections 3 and 5.

global warming indeed is caused by this channel, we cannot undo the previous policy by removing the gases from the atmosphere (although we have by then found out that they are harmful).

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 incorporating it.

Summarizing, this paper shows that uncertainty on the impact of a control variable gives agents an incentive to experiment via implementing a policy change. 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 our knowledge), which may be of independent interest.

At this stage, we wish to emphasize that this paper is normative in nature and intentionally abstracts from political-strategic considerations: we 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 be easier to sell to voters). In addition, we also assume that skeptics are open to update their beliefs in response to new data points (a refusal to do so could be motivated on political grounds, but would not be socially optimal).⁶ Subsequently, we ask what the optimal policy for these skeptics looks like, and find that even they should reduce emissions relative to current levels. This brings consensus in the policy debate, while it - somewhat ironically - also reduces the political attractiveness of the climate skeptic position. In fact, we will show that the active learning motive can be so strong that it is even possible that a climate

⁶See "The Conversion of a Climate-Change Skeptic" in *The New York Times* of July 28, 2012 for an account of a former climate skeptic (Richard A. Muller, professor of physics at UC Berkeley) who has become an IPCC believer after seeing more data points.

skeptic should actually argue for tighter emission standards than a convinced IPCC believer!

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 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.

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); Willems (2012)), 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 functioning of our climate can be optimized. In this sense it also relates to Kelly and Kolstad (1999). 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 skeptical 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 *e.g.* Epstein (1980) and Gollier, Jullien and Treich (2000); Kolstad (1996ab) and Ulph and Ulph (1997) focus explicitly at global warming).

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 by incorporating the active learning motive. Second, papers along the Arrow-Fisher-Henry lines typically compare the decision rule under exogenous, passive 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, solves for a skeptic's optimal action (the active learning rule) and compares this to his passive learning rule, as well as to the policy rule of a decision maker who is convinced that global warming is anthropogenic.

3 Learning the causes of global warming

To see how the active learning process on the causes of global warming 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:

$$\Delta\tau_t = \alpha + \beta c_t + \varepsilon_t \tag{1}$$

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.⁷ Hence, α represents potential long run trends in forces that might affect the Earth’s climate, such as 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 disturbance term representing other shorter-lived phenomena that temporarily affect global temperature (such as a cold winter in Latin America due to tropical volcanic eruptions). It is assumed that disturbances ε are i.i.d. and that they are drawn from a uniform distribution with known, bounded support, *i.e.* $\varepsilon \sim U[-\bar{\varepsilon}, \bar{\varepsilon}]$. 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 we copy this approach). Relaxing this assumption would not affect the active learning incentives (see Appendix A), 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 ϕ_1 and ϕ_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, $\alpha(\phi_1) = 0$ and $\beta(\phi_1) = \bar{\beta}$, with $\bar{\beta} > 0$.

In state 2 on the other hand, the upward trend in global temperature is completely

⁷Observe that repeated substitution on (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 we have normalized to zero for convenience). This expression for the time t temperature level is relatively standard in the literature (apart from the inclusion of the potential exogenous trend α), *cf.* equation (1a) of Kelly and Kolstad (1999).

exogenous to human behavior and the emission of greenhouse gases does not play a role. Hence, $\alpha(\phi_2) = \bar{\alpha}$ and $\beta(\phi_2) = 0$, with $\bar{\alpha} > 0$.⁸ We thus have:

$$\begin{aligned}\Delta\tau_t(\phi_1) &= \bar{\beta}c_t + \varepsilon_t \\ \Delta\tau_t(\phi_2) &= \bar{\alpha} + \varepsilon_t\end{aligned}$$

The remainder of this section shows that these two cases are easier to distinguish from each other for more extreme emission levels c_t (*i.e.* emission levels that are located further away (in an absolute value sense) from the uninformative "confounding" emission level, which will be defined below). Intuitively, in the limiting case where we stop emitting greenhouse gases altogether (*i.e.* set $c_{t+k} = 0 \quad \forall k > 0$), the upward trend in global temperature should slow down if the IPCC is right (as illustrated in Figure 1 of Solomon *et al.* (2009), which is reproduced as Figure 2 in Appendix B of this paper). 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 would be supported if the Earth's temperature would start to increase at a faster rate.

Note that learning does not require us to reduce the atmospheric *stock* of greenhouse gases (which is technically impossible): since we are currently on an upward trajectory for both temperature and atmospheric greenhouse gas concentrations, we can already learn by reducing the *flow* of these gases and by seeing whether the *rate* of temperature increases responds to that (again see Figure 2 in Appendix B).

One should keep in mind that this learning process is likely to take a while, due

⁸Our assumption that global warming is either fully due to exogenous factors or fully due to greenhouse gas emissions, is probably a stretch of reality: there, the IPCC acknowledges a (minor) role for natural factors, while some skeptics also allow for a (minor) role for greenhouse gases. What does seem true however, is that $\alpha(\phi_1) < \alpha(\phi_2)$ while $\beta(\phi_2) < \beta(\phi_1)$. This is the only element that is important for our argument. Consequently, our assumption that $\alpha(\phi_1) = \beta(\phi_2) = 0$ can just be seen as a convenient normalization that eases the analysis.

to the lags associated with global temperature dynamics: typically, it takes about 10 years before changes in greenhouse gas emissions have noticeable effects on global temperature (see *e.g.* Heal and Kriström (2002) and Figure TS.26 in IPCC (2007)). 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 an initial first-round effect becomes visible.

In addition, 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 (Solomon *et al.*, 2009). But 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 (like a change in the *rate* at which temperature increases, as visible in the simulation underlying Figure 2 in Appendix B) would do.

History moreover suggests that the time delay associated with the learning process is not an insurmountable problem: the workings of the “active learning by doing” strategy we propose have already been shown in relation to a similar problem from the past, namely that of acid rain.⁹ 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 several 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 (and will remain so in the future), so the aggressive sulfur reduction policies followed by governments in the 1970s and 80s did produce valuable information.

To formalize the learning process for the global warming case, we can exploit the assumption that the disturbance term ε has bounded support (*i.e.* that $\varepsilon \in [-\bar{\varepsilon}, \bar{\varepsilon}]$).

⁹See Seip (2001) for an overview of this debate with a particular focus on Norway.

Consequently, we know with certainty that:

$$\Delta\tau_t(\phi_1) \in [\bar{\beta}c_t - \bar{\varepsilon}, \bar{\beta}c_t + \bar{\varepsilon}] \quad (2)$$

$$\Delta\tau_t(\phi_2) \in [\bar{\alpha} - \bar{\varepsilon}, \bar{\alpha} + \bar{\varepsilon}] \quad (3)$$

Graphically, this can be visualized as in the upper panel of Figure 1.




Figure 1: Graphical representation of the learning process

Now the key is to observe that learning is complete if $c_t \leq c^*$ or $c_t \geq c^{**}$. In those cases, the regions for $\Delta\tau_t(\phi_1)$ and $\Delta\tau_t(\phi_2)$ 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 partic-

ular, c^* is defined by:

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

while c^{**} solves:

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

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 $\Delta\tau_t(\phi_1)$ or $\Delta\tau_t(\phi_2)$ (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:

$$P = \frac{|\bar{\alpha} - \bar{\beta}c_t|}{2\bar{\varepsilon}} \quad (4)$$

Note from equation (4) that $P = 1$ if either c^* or c^{**} (or something more extreme) is chosen, which is consistent with the way these revealing emission levels are defined. On the other hand, $P = 0$ for $c = \bar{\alpha}/\bar{\beta}$ (also see the lower panel of Figure 1, which depicts P as a function of c). Hence, for $c_t = \bar{\alpha}/\bar{\beta}$ 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 from equation (4) and the lower panel of Figure 1 that the function P is symmetric around this confounding emission level: the production of information only depends on $|c_t - \bar{\alpha}/\bar{\beta}|$ (not on $sgn(c_t - \bar{\alpha}/\bar{\beta})$), so positive deviations of c_t from $\bar{\alpha}/\bar{\beta}$ are just as informative as negative ones.

The fact that positive deviations of c_t from $\bar{\alpha}/\bar{\beta}$ 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 7), 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).¹⁰ Since this non-linearity increases the attractiveness of implementing a policy change in the downward direction relative to one in the upward direction for reasons other than the irreversibility (namely learning speed), this modification would only strengthen the findings that this paper will arrive at.

4 Optimizing model

Now that we are familiar with the active learning process on the causes of global warming, we can investigate how this process affects decisions related to the emission of greenhouse gases. We do this through a simple three period model, in which the active learning 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 from period 1 onwards. 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? Here, period 1 represents the learning phase, while period 2 captures the remaining future.

In particular, we 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 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

¹⁰See 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 reduction is faster than that to an increase in emissions.

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:

$$c_0 \in [0, G] \tag{5}$$

$$c_1 \in [0, G - c_0] \tag{6}$$

$$c_2 \in [0, G - c_0 - c_1] \tag{7}$$

We follow the literature in assuming that the objective function (the sum of the utility generated by the consumption of greenhouse gases and the disutility of global warming) is concave. For concreteness we 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 results. 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 $\hat{\alpha}$, $\hat{\beta}$ and $\hat{\varepsilon}$ respectively) is given by:¹¹

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

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.

¹¹In 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 less stylized model to start from), we neglect this for analytical convenience.

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.

We believe that past policy has been characterized by two 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 generation of policy makers and voters. 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 States¹²), 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 to future decision makers. If we denote their belief that global warming is endogenous by $\theta \in (0, 1)$, their problem reads:

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

subject to (5)-(7).

¹²According 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.

The accompanying first-order conditions are given by:

$$\frac{1}{c_t} - \theta\bar{\beta} = 0 \text{ for } t = 0, 1, 2, \quad (8)$$

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

$$c_0^\diamond = \frac{1}{\theta\bar{\beta}} \quad (9)$$

Now, if climate skeptics are genuinely uncertain on the causes of global warming, it has to be the case that 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".¹³

As shown in Section 3, the confounding emission level equals $\bar{\alpha}/\bar{\beta}$ so apparently the values of $\bar{\alpha}$ and $\bar{\beta}$ 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, we follow skeptics like Myers by assuming that current emission levels are uninformative, which in terms of our model implies that:

$$c_0^\diamond = \frac{\bar{\alpha}}{\bar{\beta}} \quad (10)$$

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

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

¹³See "Climate Data That Sounds Meaningful...But Isn't" (published by the Washington Policy Center).

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?

4.2.1 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. 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:

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

subject to (6) and (7).

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

$$c_t^{CC} = \begin{cases} \frac{1}{2} & \text{if warming is exogenous} \\ \frac{1}{\bar{\beta}} & \text{if warming is endogenous} \end{cases} \quad \text{for } t = 1, 2, \quad (11)$$

where the superscript "CC" indicates that this is the solution under complete certainty.¹⁴

¹⁴In the endogenous global warming case, this solution assumes that $\bar{\beta} > 2$. For $\bar{\beta} \leq 2$, the possibility of endogenous global warming is no longer "binding" and there is no value to learning the true causes of climate change anymore (also see Ulph and Ulph (1997) on this). 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, we assume that $\bar{\beta} > 2$.

4.2.2 Skepticism, no learning

As set out in the Introduction to this paper, many policy makers in practice do not adhere to one of the extreme positions covered by (11), but hold a skeptical attitude towards the causes of global warming instead. Let us therefore analyze what their optimal policy looks like. We start by considering a skeptic who neglects the possibility that he may learn more about the causes of global warming in the future. Again using θ to denote the decision maker's belief that global warming is anthropogenic, the problem for such a skeptic is characterized by:

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

subject to (6) and (7).

The first-order conditions read:

$$\frac{1}{c_t} - \theta\bar{\beta} = 0 \quad \text{for } t = 1, 2, \tag{12}$$

so optimality implies that:¹⁵

$$c_t^{SNL} = \frac{1}{\theta\bar{\beta}} \quad \text{for } t = 1, 2 \tag{13}$$

Here, the superscript "SNL" indicates that this is the solution for a skeptical decision maker who does not learn over time.

By comparing (13) with (11) one can see that climate skepticism (in particular: having $\theta < 1$) is a reason to be less careful with emitting greenhouse gases: $c_t^{SNL} > 1/\bar{\beta}$, which is the emission level chosen by a decision maker who is convinced of the anthropogenic nature of global warming. This is probably the kind of reasoning that climate skeptics have in mind when they argue that uncertainty on the causes of global warming weakens the case for emission reductions. As we will show in the remainder of this section, this conclusion is however drastically overturned once one

¹⁵For the reason set out in the previous footnote we again assume that $c_t^{SNL} < \frac{1}{2}$ (which now calls for $\theta\bar{\beta} > 2$).

takes learning considerations into account.

4.2.3 Skepticism, passive learning

First consider a skeptical social planner who again does not know the true values of α and β at the beginning of period 1, but who now takes into account that he will learn 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:

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

subject to (6) and (7).

We can solve this problem by backward induction. Here, the fact that our disturbance term ε has bounded support (as a result of which learning is discrete and Bayesian updating 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:

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

and he sets:

$$c_2^{SNL} = \frac{1}{\theta\bar{\beta}}, \tag{14}$$

where "SNL" refers to the fact 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:

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

Hence:

$$c_2^{SL} = \begin{cases} 1 - c_1 & \text{if warming is exogenous} \\ \frac{1}{\bar{\beta}} & \text{if warming is endogenous} \end{cases} \quad (15)$$

Note that expression (15) is intuitive: if global warming turns out to be exogenous, there is nothing wrong with emitting greenhouse gases and the planner decides to consume whatever there is left at the beginning of the final period (*i.e.*: $1 - c_1$). If the planner learns that global warming is anthropogenic on the other hand, he will be more cautious - and increasingly so the more responsive global temperature is to the atmospheric stock of greenhouse gases (captured by $\bar{\beta}$).

Observe that we can rewrite the objective function as:

$$\max_{c_1} \log(c_1) + P\mathbb{E}_1 \{V(c_2^{SL}) - V(c_2^{SNL})\} + \mathbb{E}_1 \{V(c_2^{SNL})\} \quad (16)$$

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

By substituting the optimal decision rules (14) and (15) into the pay-off function it follows that:

$$V(c_2^{SL}) = \begin{cases} \log(1 - c_1) - 2\bar{\alpha} & \text{if warming is exogenous} \\ -\log(\bar{\beta}) - \bar{\beta}c_1 - 1 & \text{if warming is endogenous} \end{cases}, \quad (17)$$

and:

$$V(c_2^{SNL}) = \begin{cases} -\log(\theta\bar{\beta}) - 2\bar{\alpha} & \text{if warming is exogenous} \\ -\log(\theta\bar{\beta}) - \bar{\beta}c_1 - \frac{1}{\theta} & \text{if warming is endogenous} \end{cases} \quad (18)$$

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:

$$\mathbb{E}_1 \{V(c_2^{SNL})\} = -\log(\theta\bar{\beta}) - \theta\bar{\beta}c_1 - 2(1-\theta)\bar{\alpha} - 1 \quad (19)$$

Similarly, the expected value added of learning the causes of global warming at the beginning of period 2 (indicated by $\mathbb{E}_1 \{\Delta V^L\}$) equals:

$$\begin{aligned} \mathbb{E}_1 \{\Delta V^L\} &\equiv \mathbb{E}_1 \{V(c_2^{SL}) - V(c_2^{SNL})\} \\ &= \log(\theta) + (1-\theta) [1 + \log(\bar{\beta}) + \log(1-c_1)] \end{aligned} \quad (20)$$

Using that $1/\theta\bar{\beta} \leq (1-c_1)$ (by feasibility) and $-\log(\theta) - 1/\theta < -1$ (since $\theta \in (0, 1)$), a comparison of equations (17) and (18) shows that:

$$\mathbb{E}_1 \{\Delta V^L\} > 0 \quad (21)$$

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 (13)), which works well in expectation, but is suboptimal for either state realization.

However, do observe from (20) that:

$$\frac{\partial \mathbb{E}_1 \{\Delta V^L\}}{\partial c_1} = -\frac{1-\theta}{1-c_1} \leq 0 \quad (22)$$

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 (while reducing the atmospheric stock of greenhouse gases above the natural rate of decay is not possible either).

By using (19) and (20), one can rewrite the objective function (16) as:

$$\begin{aligned} \max_{c_1} \log(c_1) + P \{ \log(\theta) + (1 - \theta) [1 + \log(\bar{\beta}) + \log(1 - c_1)] \} \\ + \{ -\log(\theta\bar{\beta}) - \theta\bar{\beta}c_1 - 2(1 - \theta)\bar{\alpha} - 1 \} \end{aligned}$$

Hence, the accompanying first-order condition reads:

$$\frac{1}{c_1} - \frac{P(1 - \theta)}{1 - c_1} - \theta\bar{\beta} = 0, \tag{23}$$

By comparing the first-order condition of this section's passive learner (23) with that of the non-learning skeptic of the previous section (12), one can see that the prospect of learning reduces a skeptic's optimal emission level, *i.e.* $c_1^{SPL} < c_1^{SNL}$.

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

reads:

$$\frac{1}{c_1} + \frac{dP}{dc_1} \Big|_{c_1=c_1^{SPL}} \mathbb{E}_1 \{ \Delta V^L \} - \frac{P(1-\theta)}{1-c_1} - \theta \bar{\beta} = 0, \quad (24)$$

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

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 (24) with the first-order condition of the passive learner (23).

When doing so, it should be noted that $\mathbb{E}_1 \{ \Delta V^L \} > 0$ (information is valuable, *cf.* equation (21)), while $dP/dc_1|_{c_1=c_1^{SPL}} = -\bar{\beta}/2\bar{\varepsilon} < 0$. The reason for the latter is that $c_1^{SPL} < \bar{\alpha}/\bar{\beta}$ (this follows from comparing (23) 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 experimentation motive induces the decision maker to *reduce* emissions relative to the passive learning solution c_1^{SPL} and even more so relative to its current level c_0^\diamond . That is: $c_1^{SAL} < c_1^{SPL} < c_0^\diamond$.

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 (22)).

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 similar, thereby reducing the relevance of this (currently heated) debate: both IPCC 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.¹⁶

4.3 Should skeptics actually emit less than IPCC believers?

Having said all this, one may wonder: is it possible that learning considerations become so dominant that a climate skeptic should actually argue for lower emissions than a policy maker who believes the IPCC analyses? That is: could optimality imply that the skeptical George W. Bush should have actually pleaded for a *more cautious* policy than Al Gore did?

To analyze this question we rewrite the first-order condition for the active learner and compare it to that of an IPCC believer (referred to as "IB" for short):

$$IB : \frac{1}{c_1} = \bar{\beta} \tag{25}$$

$$SAL : \frac{1}{c_1} = \bar{\beta} - \{(1 - \theta)\bar{\beta}\}_I + \left\{ \frac{P(1 - \theta)}{1 - c_1} \right\}_{II} - \left\{ \frac{dP}{dc_1} \Big|_{c_1^{SPL}} \right\}_{III} \mathbb{E}_1 \{ \Delta V^L \}_{IV} \tag{26}$$

In the comparison, we refer to the terms between curly brackets by the Roman subscripts attached to them ($I - IV$). II is a positive term that raises the RHS of (26) and thus lowers c_1 . This term stems from the pure impact of learning as such and works towards conservatism (recall the discussion following equation (22) for the intuition behind this).

We have already seen in Section 4.2.4 that III is negative while IV is a positive term, as a result of which their product preceded by a minus sign raises the RHS of (26) and hence also lowers c_1 . This is the active learning effect, which also leads to more conservatism.

Finally, term I is positive so preceded by a minus sign it lowers the RHS of (26) - thereby increasing c_1 . This term is unrelated to learning and simply captures the

¹⁶Only policy makers who are 100 percent sure that global warming is exogenous would implement the optimal policy (given their beliefs) by increasing emissions. However, as noted in footnote 3, this is not what policy makers typically claim to be in practice.

fact that doubts on the link between greenhouse gas emissions and global warming in isolation justify more greenhouse gas emissions. This seems to be an important argument for climate skeptics who argue that uncertainty on the causes of climate change weakens the case for emission reductions. As equation (26) nicely shows, this kind of reasoning however neglects terms *II*, *III* and *IV*. These terms are all related to the learning process and work towards more conservatism.

The net result is ambiguous: a climate skeptic who is learning in the optimal, active manner may or may not want to set tighter emission standards than a convinced IPCC believer would.¹⁷ This should not come as a surprise: for θ moving to its endpoints (0 or 1), uncertainty goes down and so does the value of learning (whether passive or active).

This is not where the discussion ends however. We can say more on the conditions under which it becomes more likely that optimality requires skeptics to set tighter emission standards than IPCC believers. First, this becomes more likely whenever the value of learning (term *IV*) goes up. This is obvious and intuitive: if learning is more valuable, you want to do more of it. This, coupled with the bias towards safe experimentation induced by the irreversibility, increases the likelihood that it is optimal for a climate skeptic to follow a more conservative climate policy than a committed IPCC believer.

Second, a higher effectiveness of experimentation (captured by *III*) also promotes a more conservative climate policy for skeptics. By recalling that

$$\left. \frac{dP}{dc_1} \right|_{c_1=c_1^{SPL}} = -\frac{\bar{\beta}}{2\bar{\varepsilon}}, \quad (27)$$

one can see that the effectiveness of experimentation is decreasing in the variance of the noise term (which is given by $\bar{\varepsilon}^2/3$ for the uniform distribution with support $[-\bar{\varepsilon}, \bar{\varepsilon}]$ underlying our model). So the noisier the climate, the less likely it becomes that a skeptic should implement a more conservative policy than an IPCC believer.

¹⁷Note that the question being asked here is different from the issue whether climate skeptics should reduce emissions relative to *current* levels. There, the answer is an unambiguous "Yes" (recall Section 4.2.4).

Although (27) also shows that experimentation becomes more effective when $\bar{\beta}$ is large (*i.e.* when global temperature is more responsive to the atmospheric stock of greenhouse gases), it is not necessarily the case that a larger $\bar{\beta}$ makes it more likely that a skeptic should set tighter emission standards than an IPCC believer should do. To see this, compare the first-order condition of an active learning with that of a passive learner:

$$\frac{1}{c_1^{SxL}} - \frac{P(1-\theta)}{1-c_1^{SxL}} - \theta\bar{\beta} + \left[\frac{dP}{dc_1} \Big|_{c_1=c_1^{SPL}} \mathbb{E}_1 \{ \Delta V^L \} \right] \mathbf{1}_{x=A} = 0, \quad (28)$$

where $\mathbf{1}_{x=A}$ is the indicator function that takes the value 1 if $x = A$ and 0 otherwise. Inserting $x = P$ into (28) and applying the implicit function theorem shows that a higher value for $\bar{\beta}$ does lead to more conservatism in the passive learning case (*i.e.* given P , the probability of learning the truth):

$$\frac{\partial c_1^{SPL}}{\partial \bar{\beta}} \Big|_P = - \left[\frac{1}{(c_1^{SPL})^2} + \frac{P(1-\theta)}{(1-c_1^{SPL})^2} \right]^{-1} \theta < 0 \quad (29)$$

When we move to the active learning solution (now $x = A$ in equation (28)), the effects on the probability of learning the truth and the value of information come into play. This makes the sign of the derivative of c_1^{SAL} with respect to $\bar{\beta}$ less clear cut. Applying the implicit function theorem now yields:

$$\frac{\partial c_1^{SAL}}{\partial \bar{\beta}} = - \left[\frac{1}{(c_1^{SAL})^2} - \frac{\bar{\beta}(1-\theta)(1-\bar{\alpha}/\bar{\beta})}{2\bar{\varepsilon}(1-c_1^{SAL})^2} \right]^{-1} \left[\theta + \frac{1}{2\bar{\varepsilon}} (\mathbb{E}_1 \{ \Delta V^L \} + (1-\theta)) \right] \quad (30)$$

The last term in (30) (the numerator) is always positive, so if the denominator is so as well, a higher value for $\bar{\beta}$ will increase the degree of conservatism for the active learner (due to the minus sign up front). The denominator will for example be positive when θ is "sufficiently close" to 1 (close enough for the first term in the denominator to dominate). So for a sufficiently strong prior belief that global warming is anthropogenic, an increase in the supposed impact of emissions on global

temperature $\bar{\beta}$ will lead to tighter policies under the active learning solution, which is intuitive. On the other hand, when θ is not "close enough" to 1 (*i.e.*: sufficiently close to 0), the opposite result obtains: then a smaller $\bar{\beta}$ leads to tighter policies.

It is thus certainly possible that a skeptical policy maker (like George W. Bush was) should argue for a *lower* emission level than a policy maker who is convinced that global warming is man-made (such as Al Gore). Whether this is the case for realistic calibrations should however be investigated with more sophisticated models of climate change. In any case, it seems an intriguing possibility.

5 Discussion

Central to this paper is the learning process on the causes of climate change. As pointed out in Section 3, this process takes time, since it takes about 10 years before changes in greenhouse gas emission levels become noticeable in global temperature. Do note that we did not have to assume anything on the length of a period in our 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 seems 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 we leave for future work. In this respect, the 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 (1996a) 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 proves to be exogenous. This introduces a second option value that actually makes emission reduction *less* attractive. One could interpret this as suggesting that we should do less to reduce greenhouse gas emissions. Another interpretation, given by Kolstad (1996b) himself, is however that we should seek for greenhouse gas reduction policies that work via *reversible* actions. In this respect, Kolstad (1996b) pleads for the installation of a temporary carbon tax, which we also see as an attractive option for two reasons: firstly, it does not introduce any direct irreversibilities, and secondly it gives firms complete freedom on how to respond to the changed incentives. If Kolstad's (1996a) 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).¹⁸

While this paper has focused at the irreversibilities related to the emission of greenhouse gases, it is also possible that global warming produces irreversible damages (like the extinction of species or a reversal of the Gulf Stream). These are not modeled in the present paper, but intuition suggests that they would only strengthen the decision maker's incentives to experiment "in the safe direction". After all, potential "tipping points" will be crossed earlier if one experiments by increasing emissions.

For policy makers who are convinced that the IPCC scenario is the truth, the

¹⁸In addition, we would like to add that not all "green" investments are fully wasted if global warming turns out to be exogenous. Although this is the case for carbon capture systems (which are useless if higher atmospheric CO₂ concentrations turn out to be harmless), this is much less so for investments in non-fossil energy sources. After all, they continue to be productive - even if global warming would turn out to be unrelated to the emission of CO₂. Moreover, given the exhaustible nature of fossil fuels, we will have to shift to non-fossil sources of energy at some point in the future anyway.

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 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 practical applicability of the "active learning by doing" strategy has already been illustrated 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 as well. One application that comes to mind is optimal consumer experimentation with potentially addictive goods/experiences, such as Apple products, hard drugs or Facebook.

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 actually looks like (although this paper can also be reinterpreted in terms of other environmental problems). Typically, climate skeptics adhere to a rather passive policy by arguing 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 uncertainty on the causes of global warming does not provide a reason for inaction, but yields a strong incentive for action instead. 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 what causes climate change, this implies that skeptical policy makers should argue for a policy change that moves us away from current (in their eyes uninformative) emission levels, as such a change of direction would produce information on the nature of global warming. They can change direction by either in- or decreasing greenhouse gas emissions, but the option of increasing emissions is unattractive because the irreversibilities associated with emitting greenhouse gases erode the value of the information that is produced.

So whereas uncertainty on the causes of global warming gives skeptical decision makers an incentive to implement a policy change, the irreversible nature of emitting greenhouse gases induces them to do so "in the safe direction". Theoretically, it is even possible that the active learning 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 IPCC 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 similar: both IPCC believers as well as climate skeptics should argue for 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.

To estimate how large a reduction in greenhouse gas emissions the active learning motive actually calls for, requires a more realistic model of climate change. We leave this issue for future work. In this paper we have focused at maintaining analytical tractability in order to gain insight into the exact mechanisms at play, which forced

us to simplify along several dimensions. A more realistic model would generalize our three-period setup as well as the learning process, while it would ideally also take the irreversibilities in abatement investments and environmental damages into account. On this latter issue, the "Stern Review" has argued that the investments in abatement capital necessary to avoid the worst effects of climate change are small relative to the potential damages. This suggests that our main conclusion is robust to the joint incorporation of these two irreversibilities that we abstract from in our current analysis.

7 Appendix A

In the main text it was assumed that the disturbance 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 disturbance 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 $\beta = \bar{\beta} > 0$), or the skeptics are right (which implies $\alpha = \bar{\alpha} > 0$ 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:

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

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

8 Appendix B

Figure 2 shows the outcome of a simulation study by Solomon *et al.* (2009).




Figure 2: Climate system responses for a ramp of CO2 emissions at a rate of 2% per year to peak CO2 values of 450, 550, 650, 750, 850, and 1200 ppmv, followed by zero emissions.

The figure shows what would happen if we were to reduce CO2 emissions (*i.e.*: the flow) at certain points in time. Solomon *et al.* (2009) take the extreme case in which

emissions are reduced all the way down to zero, but a more realistic partial emission reduction would induce similar responses (although obviously somewhat muted (see Figure S1 of Solomon *et al.* (2009))).

In both cases one can observe a clear change in the rate at which global temperature increases after emissions have been reduced. If we were to observe such a temperature response in reality after decreasing emissions, that would support the IPCC's case. Absence of such a response on the other hand, would suggest that global warming is exogenous. Once this knowledge has been acquired, environmental policy can be conditioned on this information - enabling better economic outcomes.

9 References

Aghion, P., P. Bolton, C. Harris and B. Jullien (1991), "Optimal Learning by Experimentation", *Review of Economic Studies*, 58 (4), pp. 621-654.

Arrow, K.J. and A.C. Fisher (1974), "Environmental Preservation, Uncertainty, and Irreversibility", *Quarterly Journal of Economics*, 88 (2), pp. 312-319.

Bertocchi, G. and M. Spagat (1993), "Learning, Experimentation, and Monetary Policy", *Journal of Monetary Economics*, 32 (1), pp. 169-183.

Cogley, T., R. Colacito, L.P. Hansen and T.J. Sargent (2008), "Robustness and US Monetary Policy Experimentation", *Journal of Money, Credit and Banking*, 40 (8), pp. 1599-1623.

Epstein, L.G. (1980), "Decision Making and the Temporal Resolution of Uncertainty", *International Economic Review*, 21 (2), pp. 269-283.

Gollier, C., B. Jullien and N. Treich (2000), "Scientific Progress and Irreversibility: An Economic Interpretation of the Precautionary Principle", *Journal of Public Economics*, 75 (2), pp. 229-253.

Grossman, S., R. Kihlstrom and L. Mirman (1977), "A Bayesian Approach to the Production of Information and Learning By Doing", *Review of Economic Studies*, 44 (3), pp. 533-547.

Heal, G. and B. Kriström (2002), "Uncertainty and Climate Change", *Environmental and Resource Economics*, 22 (1-2), pp. 2-29.

Henry, C. (1974), "Investment Decisions Under Uncertainty: The Irreversibility Effect", *American Economic Review*, 64 (6), pp. 1006-1012.

IPCC (2007), "IPCC Fourth Assessment Report: Climate Change 2007 (AR4) - The Physical Science Basis".

Kelly, D.L. and C.D. Kolstad (1999), "Bayesian Learning, Growth, and Pollution", *Journal of Economic Dynamics and Control*, 23 (4), pp. 491-518.

Kolstad, C.D. (1996a), "Fundamental Irreversibilities in Stock Externalities", *Journal of Public Economics*, 60 (2), pp. 221-233.

Kolstad, C.D. (1996b), "Learning and Stock Effects in Environmental Regulation: The Case of Greenhouse Gas Emissions", 31 (1), pp. 1-18.

MacRae, E.C. (1972), "Linear Decision With Experimentation", *Annals of Economic and Social Measurement*, 1, pp. 437-448.

Prescott, E.C. (1972), "The Multi-Period Control Problem under Uncertainty", *Econometrica*, 40 (6), pp. 1043-1058.

Seip, H.M. (2001), "Acid Rain and Climate Change - Do These Environmental Problems Have Anything in Common?", *Cicerone*, 6.

Solanki, S.K., I.G. Usoskin, B. Kromer, M. Schüssler and J. Beer (2004), "Unusual Activity of the Sun During Recent Decades Compared to the Previous 11,000 Years", *Nature*, 431, pp. 1084-1087.

Solomon, S., G.K. Plattner, R. Knutti and P. Friedlingstein (2009), "Irreversible Climate Change Due to Carbon Dioxide Emissions", *Proceedings of the National Academy of Sciences*, 106 (6), pp. 1704-1709.

Stouffer, R.J. (2004), "Time Scales of Climate Response", *Journal of Climate*, 17 (1), pp. 209-217.

Ulph, A. and D. Ulph (1997), "Global Warming, Irreversibility and Learning", *Economic Journal*, 107 (442), pp. 636-650.

Wigley, T.M.L. (2005), "The Climate Change Commitment", *Science*, 307, pp. 1766-1769.

Willems, T. (2012), "Actively Learning by Pricing: A Model of an Experimenting Seller", mimeo, University of Amsterdam.