UvA-DARE (Digital Academic Repository)

Optimal learning on climate change: why climate skeptics should reduce emissions

van Wijnbergen, S.; Willems, T.

DOI

[10.1016/j.jeem.2014.12.002](https://doi.org/10.1016/j.jeem.2014.12.002)

Publication date

2015

Document Version

Submitted manuscript

Published in

Journal of Environmental Economics and Management

[Link to publication](#)

Citation for published version (APA):

van Wijnbergen, S., & Willems, T. (2015). Optimal learning on climate change: why climate skeptics should reduce emissions. *Journal of Environmental Economics and Management*, 70, 17-33. <https://doi.org/10.1016/j.jeem.2014.12.002>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

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

Abstract

Climate skeptics typically argue that the possibility that global warming is exogenous, implies that we should not take additional action towards reducing emissions until we know what drives warming. This paper however shows that even climate skeptics have an incentive to reduce emissions: such a directional change generates information on the causes of global warming. Since the optimal policy depends upon these causes, they are valuable to know. Although increasing emissions would also generate information, that option is inferior due its irreversibility. We show that optimality can even imply that climate skeptics should actually argue for lower emissions than believers.

JEL-classification: D83, Q54, Q58

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

1 Introduction

Many policy makers and members of the public question the supposed link between global warming and man-made emissions of greenhouse gases: according to a 2007/8 Gallup Polls survey, 97 percent of all US adult citizens say they are aware of global warming, but only 49 percent of them believe that it is anthropogenic. Although there is more consensus among climate scientists, some remain uncertain as well: Farnsworth and Lichter (2011) for example report that 16 percent of climate scientists are not wholly convinced of

*We thank two anonymous referees, Andrea Craig, Christian Gollier, Mark Kagan, Larry Karp, Shaun Larcom, Radek Stefanski, Cees Withagen, and audiences at the LSE (Grantham Institute), University of Oxford (OxCarre), VU University Amsterdam, and the EAERE-conference in Toulouse for their input. Comments and references provided by John van Boxel, Gian-Kasper Plattner, and Bart Verheggen greatly improved the climate-scientific side of this paper. Any errors are naturally ours.

the anthropogenic view. Skepticism among policy makers seems even more widespread: 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 for example be heard among policy makers in China (Xie Zhenhua, their lead negotiator in the last three UN Climate Conferences), the Czech Republic (their previous 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 casts the underlying debate in a formal framework and provides a normative analysis of what the optimal policy for these skeptics actually looks like. We focus on skepticism on the causes of global warming (is it due to natural, or due to anthropogenic forces?) and define a climate skeptic as someone who is uncertain on these causes.² In practice, such skeptics typically propose a rather passive policy. They tend to argue that the possibility that global warming is driven by exogenous factors (like increases in solar activity) implies that we should not take additional action towards reducing greenhouse gas emissions until we know what causes our climate to change: Mitt Romney (the Republican candidate in the 2012 US elections) for example argued in October 2011 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”.³ 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 genuine uncertainty on the causes of global warming weakens the case for emission reductions may make intuitive sense at first sight, this paper shows that it is incomplete 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, we show that even climate skeptics have an

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

²This seems the most common form of climate skepticism, but there are others. Note that this is different from someone who is 100 percent certain that global warming is exogenous. This extremer position, which leaves no room for learning, is relatively rare among policy makers (*cf.* the aforementioned quote by George W. Bush, where he talks about “a debate”, while Xie Zhenhua has stated that China is keeping “an open attitude” on the causes of climate change). It is moreover also inconsistent with the standard definition of a skeptic (“a person inclined to question or doubt accepted opinions”).

³See “Mitt Romney Embraces Climate Denial: ‘We Don’t Know What’s Causing Climate Change’ ” in *The Huffington Post* of October 28, 2011.

incentive to reduce greenhouse gas emissions relative to current levels: when uncertain on the relationship between a control variable (such as emissions) and a potentially dependent variable (such as global temperature), decision makers obtain an incentive to change the control through a policy shift (like reducing emissions). The reason is that such a change of direction produces information on whether the potentially dependent variable is indeed dependent. Since the optimal policy depends on whether this is the case or not, information on the causes of global warming is valuable.

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 they have produced over time - thereby rendering this information useless. After all, information is only valuable if it is able to change behavior. If we (after consciously emitting more greenhouse gases) learn that global warming is indeed caused by this channel, we cannot use this information by undoing the previous policy via removing these gases from the atmosphere (although we have by then found out that they are harmful).

In contrast, a cautious policy leaves all options open - making it more robust: if this policy teaches us that there is a link between the stock of greenhouse gases and global temperature, the prudence was justified. Alternatively, 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 as decision makers can improve their future actions by incorporating it.

One may of course question whether outspoken characters (like Al Gore in the “IPCC camp” or Václav Klaus in the “non-believers camp”) actually want to learn and are willing to change their beliefs in response to new data.⁴ Even in the most extreme case where this is not so, more outspoken data would make it more likely that new policy makers (who have not joined a particular camp yet) will adopt the correct view - especially if more informative data also raises the proportion of voters with correct beliefs (as this increases the probability that they will elect politicians who share these beliefs).

⁴But keep in mind that agents have updated their beliefs in similar debates which had equally polarized starts, like that on acid rain (see Section 3 below). Also note that the former Russian president Dmitry Medvedev did “switch camps” by abandoning his initial skepticism (see “6 Global Warming Skeptics Who Changed Their Minds” in *The Week* of September 1, 2010). So did the Australian prime minister Tony Abbott: although he initially dismissed the man-made global warming hypothesis (famously stating in 2009 that “the argument is absolute crap”), he said in 2013 that “climate change is real” and that “humanity makes a contribution” (cf. <http://www.abc.net.au/insiders/content/2012/s3838363.htm>). This still did not prevent him from scrapping Australia’s carbon tax, though. Also see “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.

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 (descriptively, this assumption seems most appropriate for skeptical voters). Subsequently, we ask what the optimal policy for these skeptics looks like, and find that even they should argue (or vote) for emission reductions relative to current levels. This brings consensus in the policy debate, while it - ironically - also reduces the political attractiveness of the climate skeptic position. In fact, we will show that it is even possible that a climate 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 illustrates that the learning process on the causes of global warming is facilitated by either in- or decreasing greenhouse gas emissions relative to some 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 notion 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, recognizing that their actions can also produce information. 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 parameters of an economy it tries to control (Bertocchi and Spagat, 1993).

This paper analyzes how we can optimize our learning process on the causes of climate change. In this sense it also relates to Kelly and Kolstad (1999). In their paper, however,

the decision maker is convinced that global warming is anthropogenic and 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 agents in practice are stuck at an earlier stage (questioning whether there is a link between global warming and human activities in the first place), this paper abstracts from the learning process on the climate sensitivity parameter - all the more so because the exact value of that parameter does not seem to be very learnable (even in the face of more extreme data; see Kelly and Kolstad (1999) and Roe and Baker (2007)). This reduces the active learning incentive along that dimension (after all: why introduce costly deviations from the myopically optimal emission level if such deviations do not produce much information?). Here, we take this to the extreme by completely neglecting the active learning motive for the climate sensitivity parameter - focusing on the learning process about whether global warming is endogenous or exogenous instead.

Finally, this paper also builds upon studies that have investigated the consequences of irreversibilities in environmental settings. Following the seminal contributions of Arrow and Fisher (1974) and Henry (1974), many papers have analyzed the so-called “quasi-option value” to maintaining flexibility that exists if the quality of information increases over time (see *e.g.* Epstein (1980) and Gollier and Treich (2003); contributions by Kolstad (1996a,b) and Ulph and Ulph (1997) focus explicitly on global warming). These papers however assume that information arrives exogenously with the passage of time, while the key of this paper is that the acquisition of information is endogenized.

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 tractable Bayesian learning model that will be employed in the policy maker’s decision problem in Section 4.

In the model, 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:⁵

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

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 long-run exogenous factors,

⁵ “Global temperature” should be taken as a broad, encompassing metric. It should not just capture the temperature of the Earth’s atmosphere, but also that of the oceans (as they absorb heat as well; Guemas et al. (2013) argue that this is why the warming of our atmosphere has slowed down recently).

such as a trend in solar activity (see Solanki *et al.* (2004) who show that solar activity has been exceptionally high over the past 70 years).⁶ The slope parameter β on the other hand captures the possible relationship between global temperature and greenhouse gas emissions.⁷ 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 Bayes' rule renders learning discrete (either nothing is learned, or the full truth is learned), which delivers analytical convenience without losing generality (*cf.* Bertocchi and Spagat (1993)). 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 greenhouse gas emissions, 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 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:

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

⁶Although we take α to be constant, one could also allow for exogenous changes in this parameter. But in order to be able to match the observed persistent increases in global temperature (which have been going on for over 100 years already), one would have to make the cycle for α slow-moving with a very long period. Consequently, holding α constant over the time-scale of our learning model might not be such an unrealistic approximation.

⁷Observe that repeated substitution on (1) leads to $\tau_t = \underline{\tau} + (t+1)\alpha + \beta\sum_{j=0}^t c_j$, where $\sum_{j=0}^t c_j$ is the atmospheric stock of greenhouse gases at time t (potentially relative to a certain base level $\underline{c}$, which we have normalized to zero for convenience). This expression for the time t temperature level is relatively standard in the literature (*cf.* equation (1a) of Kelly and Kolstad (1999)), apart from the fact that we allow for an exogenous trend α .

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 $\alpha(\phi_1) < \alpha(\phi_2)$ and $\beta(\phi_2) < \beta(\phi_1)$. This is the essential part of our argument. Consequently, our assumption that $\alpha(\phi_1) = \beta(\phi_2) = 0$ is without loss of generality. In reality there is also a debate on the exact value of $\bar{\beta}$, but as set out in Section 2 the value of this climate sensitivity parameter is not very learnable, as a result of which we will abstract from active learning motives with respect to this term.

The remainder of this section shows that these two cases are easier to distinguish for more extreme emission levels (*i.e.* emission levels that are located further away (in absolute value) 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 3 in Appendix B to 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 applies to upward experimentation.

The idea is thus to engineer a trend break in the control (emissions) and see whether this subsequently translates into a trend break in the potentially endogenous variable (global temperature). The rise in global temperature following the post-World War II trend break in emissions already convinced many skeptics that global warming is man-made, but others attribute it to the coinciding increase in solar activity (for example documented by Solanki *et al.* (2004)). As can be seen from the so-called “Keeling Curve” (which shows a sustained linear increase in carbon concentration levels), there are no trend breaks in the post-war data. Consequently, the current data are not able to falsify either hypothesis. Introducing a trend break would change this.

Since our model is essentially concerned with hypothesis testing, agents should of course commit themselves to a particular hypothesis that can be falsified by the data. If such falsification occurs, they could always propose ex-post modifications to the theory so that it can explain the data after all. The point however is that this does the credibility of their theory little good - particularly in the climate setup. After all, it would be unlikely that an exogenous, unexpected, non-cyclical change in natural factors turns out to coincide with a man-induced trend break in emissions (especially since natural changes of this kind are extremely rare and often operate at geological time-scales).⁹

One should also keep in mind that this learning process will take time: despite the fact that changes in greenhouse gas emissions have an instantaneous impact on atmospheric temperature (Solomon *et al.*, 2009), the presence of noise (weather) obscures inference by decreasing the signal-to-noise ratio. Santer *et al.* (2011) take this into account and show that one needs about 15 years of data before temperature changes can be attributed to

⁹The required natural phenomenon should be an unexpected, non-cyclical change: after all, if the change would be expected (*e.g.* driven by a turning point in the solar cycle, which is predictable up to an error margin of several years), it would be factored into the predictions made. Moreover, the solar cycle is *cyclical* (with a period of about 11 years), as a result of which it would not be able to account for a *structural* break in the temperature data, which is what agents in our paper are looking for.

either changes in greenhouse gas emission levels or exogenous factors. Consequently, a period in the model should be thought of as about a decade and a half in reality.

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.¹⁰ During the 1960s, an increasing number of streams and lakes in Norway were reported to be acidic. Initially, it was fiercely debated whether these changes were anthropogenic or not, but nevertheless sulfur emissions were cut in a drastic manner. After several years (again due to the existence of short-run fluctuations), observers began to notice that the level of acidity in Norwegian waters had fallen, as a result of which there is nowadays little doubt left that the changes were in fact man-made (even former skeptics have updated their beliefs by now). Consequently, we are still careful with emitting sulfur today, and will remain so in the future, so the aggressive sulfur reduction policies from the 1970s and 80s did produce valuable information. Along similar lines, the EU recently decided to ban pesticides linked to bee deaths, despite the fact that the supposed link is debated. The European Commission however said that it wanted to “err on the side of safety”.¹¹ The ban is in any case going to be in place for two years (to allow for the learning lag), unless any new scientific evidence emerges before that time.

To formalize the learning process for the global warming case, we can exploit our assumption that the disturbance term ε has bounded support (*i.e.* that $\varepsilon \in [-\bar{\varepsilon}, \bar{\varepsilon}]$). 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. There, the key is to observe that learning is complete when $c_t \leq c^*$ 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 particular, c^*

¹⁰See Seip (2001) for an overview of this debate with a particular focus on Norway.

¹¹See <http://money.cnn.com/2013/04/29/news/world/bees-ban-pesticide-europe/index.html>. This is an expression of the so-called “precautionary principle”, which prescribes that if there is a risk that a certain activity produces irreversible damage, lack of consensus should *not* be used as an argument to postpone measures that can prevent such damage from occurring (see e.g. Bodansky (1991) and Immordino (2003) on this principle).

Figure 1: Graphical representation of the learning process

and c^{**} are defined by:

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

We will assume that $\bar{\varepsilon} < \bar{\alpha}/2$, such that $c^* > 0$. This restricts the power of the active learning motive in such a way that it will never call for a negative greenhouse gas emission level (this will become clear in our discussion around equation (23) below). Given the magnitude of current emissions, we see this as a realistic restriction. After all, it is hard to imagine that the experimentation motive by itself becomes so strong that, starting from current greenhouse gas emission levels of about 50,000 megatonnes per annum, it would actually call for reducing emissions all the way down to zero.

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 (2) or (3), you learn the full truth. When the $(\Delta\tau_t, c_t)$ -observation can occur under both scenarios, the uniformity assumption on the noise term implies that nothing is learned (the posterior belief is just equal to the prior belief). Exploiting this uniformity of 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 when either c^* or c^{**} (or something more extreme) is chosen, $P = 1$ - 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, where the two lines intersect and where we do not learn anything about what drives global warming.

Also observe from (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 $\text{sgn}(c_t - \bar{\alpha}/\bar{\beta})$), 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 common 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 is faster than that to an increase in emissions.¹² 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 greenhouse gas emissions. We do this through a simple two-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 1 represents the learning phase, while period 2 captures the remaining

¹²See Stouffer (2004): any development that cools the ocean surface makes the oceans less stable, 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, this implies that the atmospheric temperature response to an emission reduction is faster than that to an increase in emissions.

future. Over time, learning occurs by applying Bayes' rule in the face of new data points.

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 atmospheric presence of these gases may prove to be harmful in the future (if it causes global warming). Furthermore, greenhouse gas consumption is irreversible: once emitted, it is not possible to remove greenhouse gases from the atmosphere again, so $c_t \geq 0$.

Following Ulph and Ulph (1997, p. 640), we assume that damages resulting from global warming are going to materialize at the end of period 2. If global warming is man-made, these damages will be proportional to $\sum_{t=1}^2 c_t$, the atmospheric greenhouse gas stock at that point (as those then determine $(\tau_2 - \underline{\tau})$, which is the total amount of warming relative to the base temperature level $\underline{\tau}$).

At the beginning of period 1, the Earth is endowed with an exhaustible stock of greenhouse gases that resides beneath the Earth's surface (think of this as subsoil oil). We normalize this stock to 1. Since it is depleted over time, future consumption choices are constrained by the amount of gases consumed in the past. In particular:

$$c_1 \in [0, 1] \quad (5)$$

$$c_2 \in [0, 1 - c_1] \quad (6)$$

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 linear in temperature increases and logarithmic in greenhouse gas consumption (but our results are robust to other common specifications, such as the more general CRRA-utility function). Abstracting from discounting and natural decay of atmospheric greenhouse gases to lighten notation, the pay-off function (conditional on the actual values of α and β and the realizations of the ε 's, indicated by $\hat{\alpha}$, $\hat{\beta}$, and $\hat{\varepsilon}$ respectively) is given by:¹³

$$U(c_1, c_2; \hat{\alpha}, \hat{\beta}, \hat{\varepsilon}) = \sum_{t=1}^2 \log(c_t) - (\tau_2 - \underline{\tau}) = \sum_{t=1}^2 \log(c_t) - 2\hat{\alpha} - \hat{\beta} \sum_{t=1}^2 c_t - \sum_{t=1}^2 \hat{\varepsilon}_t \quad (7)$$

¹³In this specification both periods receive equal weight. Since the final period is meant to capture the infinite future, that period should actually obtain a larger weight (proportional to $1/(1-d)$, where d is the discount factor). However, as this would only affect the optimal *level* of experimentation (which is not analyzed in this paper, since it would require a less stylized model to start from), we neglect this for analytical convenience. Moreover, as we will point out in footnote 17, the effect of this assumption on the crucial first-order condition (23) is actually not that unrealistic.

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.

We will use our model to answer two questions:

1. Given the current state, what does the *optimal* policy for a climate skeptic look like from period 1 onwards? That is: what kind of policy would be implemented by a climate skeptical social planner if he were to take over power at the beginning of period 1? Would he in- or decrease emissions relative to current levels? This is dealt with in Section 4.1.
2. Can optimality imply that climate skeptics should actually argue for lower emission levels than IPCC believers? This is dealt with in Section 4.2.

To answer Question 1, we need to compare the optimal climate policy to actual policies followed by most countries in the recent past. Consequently, we first need to characterize the latter in terms of our model. To obtain such a “benchmark” emission level (henceforth referred to as c_b), we assume that it is based upon an optimization process in our two-period framework when the policy maker is climate skeptic, but when he does not take the possibility of learning into account (which seems a reasonable description of current practices in reality). In this respect, it is straightforward to show that maximizing (7) while ignoring the possibility of learning leads to:

$$c_b = \frac{1}{\theta\bar{\beta}}, \quad (8)$$

where $\theta \in (0, 1)$ denotes the policy maker’s prior belief that global warming is endogenous.¹⁴

According to climate skeptics, this emission level c_b is non-informative on the causes of climate change. Todd Myers, a prominent climate skeptic, has for example 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”.¹⁵

¹⁴Strictly speaking the solution is $c_b = \min \left[\frac{1}{\theta\bar{\beta}}, \frac{1}{2} \right]$, but the problem becomes trivial if the exhaustibility constraint binds in the absence of learning considerations. In that case, the optimality of downward experimentation follows immediately from concavity of the utility function: when the resource constraint is binding, an increase in emissions by X in period 1 should be met by a corresponding decrease of at least X in period 2. But when marginal utility decreases with consumption, the latter correction induces a larger utility loss than the utility gain enjoyed in period 1. Consequently, downward experimentation becomes optimal for reasons that are not the focus of this paper. We therefore assume that we are at an interior optimum such that $c_b = \frac{1}{\theta\bar{\beta}}$. We thank an anonymous referee for pointing this out.

¹⁵See “Climate Data That Sounds Meaningful...But Isn’t” (by the Washington Policy Center).

This description fits closely with our Figure 1, which shows that 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 the current emission level c_b is uninformative on the causes of global warming, which in terms of our model implies that:

$$c_b = \frac{\bar{\alpha}}{\bar{\beta}} \quad (9)$$

4.1 Characterization of optimal policy

Climate policy as encapsulated in the choice of c_b is clearly suboptimal given the fact that informational spillovers (and the associated learning considerations) are not incorporated. Let us therefore investigate what the *optimal* policy looks like. That is: what policy would be implemented by a social planner who takes into account that the quality of information may increase over time?

In order to answer this question, we start by considering the problem under complete certainty in Section 4.1.1. In this case, the decision maker knows the actual values of α and β (indicated by $\hat{\alpha}$ and $\hat{\beta}$) at the start of period 1 (before setting c_1) already. Section 4.1.2 analyzes what happens if he does not have this information at the start of period 1, but may learn it (in a passive manner) at the start of period 2 (before setting c_2). Section 4.1.3 subsequently analyzes the case of active learning, in which the decision maker realizes that the probability of learning is actually endogenous. In all specifications the pay-off function is given by (7), while c_t is set sequentially at the beginning of each period t .

4.1.1 Complete certainty

First consider the extreme case where the decision maker knows at the beginning of period 1 already whether global warming is exogenous or endogenous (so that uncertainty and learning do not play a role). 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 (5) and (6).

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

$$c_t^{CC} = \begin{cases} \frac{1}{2} & \text{if warming is exogenous} \\ \min\left[\frac{1}{\bar{\beta}}, \frac{1}{2}\right] & \text{if warming is endogenous} \end{cases}, \text{ for } t = 1, 2 \quad (10)$$

where “CC” indicates that this is the solution under complete certainty. Equation (10) shows that if global warming is known to be exogenous, the exhaustible resource will be fully exploited; the exhaustibility constraint is then always binding and $c_t = \frac{1}{2}$ for $t = 1, 2$. If global warming is known to be endogenous, (10) implies that either the resource will be exploited fully as well (if $\bar{\beta} < 2$), or greenhouse gas consumption will be reduced compared to the exogenous warming case (if $\bar{\beta} > 2$).

When $\bar{\beta} < 2$, climate sensitivity is so low that emissions are not a problem worth worrying about. As a result, the exhaustibility constraint will be binding even it is known that global warming is endogenous and there is no value to learning the true causes of climate change anymore: the optimal policy does not depend on the nature of global warming and $c_t^{CC} = 1/2$ in both states. Consequently, there is nothing to argue about for policy makers. In the second case ($\bar{\beta} > 2$), warming is serious enough such that knowing that global warming is anthropogenic *would* lead to a change in policy compared to the exogenous warming case. In particular, emissions would be reduced.

Judging from the existence of a fierce policy debate on which state we are actually in (exogenous or endogenous global warming?), climate policy being irrelevant because inconsequential does not seem to be the case in reality. We therefore assume that $\bar{\beta} > 2$ henceforward, as a result of which there *does exist* a meaningful climate-debate. Under this condition, a decision maker living in a world where global warming is known to be anthropogenic would set a lower emission level than a decision maker living in a world where global warming is known to be exogenous.

4.1.2 Skepticism, passive learning

As set out in the Introduction to this paper, many policy makers (and voters) in practice do not adhere to one of the extreme positions covered by (10), but hold a skeptical attitude towards the causes of global warming instead.

Let us therefore consider 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 learn these values with probability P at the beginning of period 2 (before he sets c_2). Exploiting our assumption that the costs of global warming are incurred in the final period (so that they are contained in the indirect utility functions $V(\cdot)$), 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 (5) and (6).

We can solve the problem by backward induction. Here, the fact that our disturbance term ε is drawn from a uniform distribution with bounded support (as a result of which learning is discrete) pays off in terms of tractability. Thanks to this assumption, there are only two possibilities at the beginning of period 2: either the decision maker has learned the true state of nature (which happens with probability P), or he has not. Consequently, those are the only two cases that we need to consider at the start of the second period, after which we can move back and analyze the problem at the beginning of period 1.

Period 2, no learning First consider the beginning-of-period-2 problem of a decision maker who has *not* learned the true causes of global warming by the end of period 1. Consequently, he continues to use his prior belief θ in period 2 (this follows from combining Bayes' rule with the uniformity assumption on ε ; *cf.* Bertocchi and Spagat (1993)). The problem is then given by:

$$\max_{c_2 \in [0, 1-c_1]} \log(c_2) - 2(1-\theta)\bar{\alpha} - \theta\bar{\beta}[c_1 + c_2], \quad (11)$$

and the decision maker sets:

$$c_2^{SNL} = \min \left[\frac{1}{\theta\bar{\beta}}, \frac{1}{2} \right], \quad (12)$$

where “*SNL*” indicates that this is the solution to a skeptic's problem who has not learned the causes of global warming. One can furthermore show that $c_1^{SNL} = c_2^{SNL}$, which we will use later on.

Going forward, we first assume $\theta\bar{\beta} > 2$ such that (12) simplifies to $c_2^{SNL} = \frac{1}{\theta\bar{\beta}}$ (but as we will argue at the end of this section, the case where $\theta\bar{\beta} < 2$ has similar implications). Using this in the pay-off function (11), shows that indirect utility to this policy maker is:

$$V(c_2^{SNL}) = \begin{cases} -\log(\theta\bar{\beta}) - 2\bar{\alpha} & \text{if warming is exogenous} \\ -\log(\theta\bar{\beta}) - \bar{\beta}c_1 - \frac{1}{\theta} & \text{if warming is endogenous} \end{cases} \quad (13)$$

Period 2, learning Next, consider a skeptical decision maker who has learned the causes of global warming by the end of period 1 (indicated by “*SL*”). He solves:

$$\max_{c_2 \in [0, 1-c_1]} \log(c_2) - 2\hat{\alpha} - \hat{\beta} [c_1 + c_2] \quad (14)$$

So he sets:

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

Expression (15) is intuitive: if it is learned that global warming is exogenous, there is nothing wrong with emitting greenhouse gases and it is optimal to consume whatever there is left at the beginning of the final period ($1 - c_1$). If the decision maker 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 $\bar{\beta}$). Plugging (15) into the pay-off function (14) gives:

$$V(c_2^{SL}) = \begin{cases} \log(1 - c_1) - 2\bar{\alpha} & \text{if warming is exogenous} \\ -\log(\bar{\beta}) - \bar{\beta}c_1 - 1 & \text{if warming is endogenous} \end{cases} \quad (16)$$

Period 1 Now we can move back to the start of period 1. At this stage our decision maker is uncertain on the causes of global warming (as expressed by his prior belief $\theta \in (0, 1)$). He does however realize that there is a probability P that he will learn these true causes before setting c_2 . We can rewrite the first-period maximization problem for this decision maker as:

$$\max_{c_1 \in [0, 1]} \log(c_1) + P\mathbb{E}_1 \{V(c_2^{SL}) - V(c_2^{SNL})\} + \mathbb{E}_1 \{V(c_2^{SNL})\} \quad (17)$$

By applying the prior belief θ to equation (13), one can show that the beginning-of-period-1 expectation of the indirect utility that will be obtained by a policy maker who will *not* learn the causes of global warming before the start of period 2, equals:

$$\mathbb{E}_1 \{V(c_2^{SNL})\} = -\log(\theta\bar{\beta}) - \theta\bar{\beta}c_1 - 2(1 - \theta)\bar{\alpha} - 1 \quad (18)$$

Similarly, the expected value added from learning the causes of global warming at the beginning of period 2 (indicated by $\mathbb{E}_1 \{V(c_2^{SL}) - V(c_2^{SNL})\} \equiv \mathbb{E}_1 \{\Delta V^L\}$) equals:

$$\mathbb{E}_1 \{\Delta V^L\} = \log(\theta) + (1 - \theta) [1 + \log(\bar{\beta}) + \log(1 - c_1)] \quad (19)$$

Using the result 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 (13) and (16) shows that:

$$\mathbb{E}_1 \{ \Delta V^L \} > 0 \quad (20)$$

So 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 (19) that:

$$\frac{\partial \mathbb{E}_1 \{ \Delta V^L \}}{\partial c_1} = -\frac{1 - \theta}{1 - c_1} \leq 0 \quad (21)$$

This expression is important to the results in this paper (while a similar effect typically arises in other papers along the Arrow-Fisher-Henry lines). 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 produced during the first period (also see Gollier and Treich (2003) on this). 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 the outcome of his learning process, there is nothing left to consume (while reducing the atmospheric stock of greenhouse gases above the natural rate of decay is not feasible because of the restriction that $c_t \geq 0$).

Using (18) and (19), one can rewrite the problem in (17) as:

$$\begin{aligned} \max_{c_1 \in [0, 1]} & \log(c_1) + P \{ \log(\theta) + (1 - \theta) [1 + \log(\bar{\beta}) + \log(1 - c_1)] \} \\ & + \{ -\log(\theta\bar{\beta}) - \theta\bar{\beta}c_1 - 2(1 - \theta)\bar{\alpha} - 1 \} \end{aligned}$$

Hence, the accompanying first-order condition reads:

$$\frac{1}{c_1} - \theta\bar{\beta} - \frac{P(1 - \theta)}{1 - c_1} = 0 \quad (22)$$

By comparing the passive learner's first-order condition (22) with that of a non-learning skeptic (given by $\frac{1}{c_t} - \theta\bar{\beta} = 0$), one can see that the prospect of learning reduces a skeptic's optimal emission level, i.e. $c_1^{SPL} < c_1^{SNL}$. As noted by O'Neill *et al.* (2006, p.

585), this conclusion is supported by a wide variety of richer climate models.

So far, we have restricted our discussion to the case where $\theta\bar{\beta} > 2$. One should however note that $c_1^{SPL} < c_1^{SNL}$ also holds in the alternative case where $\theta\bar{\beta} < 2$ (which implies that $c_1^{SNL} = \frac{1}{2}$), provided that $\theta\bar{\beta} + \frac{P(1-\theta)}{1-c_1} > 2$ (which makes $c_1^{SPL} < \frac{1}{2}$ by (22)). When $\theta\bar{\beta} + \frac{P(1-\theta)}{1-c_1} < 2$, we have $c_1^{SPL} = c_1^{SNL}$. Crucially, however, c_1^{SPL} will *never exceed* c_1^{SNL} - so in no parameter-configuration will the possibility of future learning lead to *higher* emissions in the current period.

4.1.3 Skepticism, active learning

Just taking into account that additional information may arrive over time, is however not enough to make a skeptic implement the optimal policy. The reason is that last section's passive learner neglects the fact that the signal-to-noise ratio, and thereby the probability of learning the truth P , are 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 optimizing active learner takes the effect of first period consumption on the probability of learning the truth into account. The first-order condition of such an agent therefore reads:¹⁶

$$\begin{aligned} \frac{1}{c_1} - \theta\bar{\beta} - \frac{P(1-\theta)}{1-c_1} + \frac{dP}{dc_1} \Big|_{c_1=c_1^{SPL}} \mathbb{E}_1 \{ \Delta V^L \} &= 0, \\ \text{with } \frac{dP}{dc_1} &= \frac{\bar{\beta}c_1 - \bar{\alpha}}{|\bar{\beta}c_1 - \bar{\alpha}|} \frac{\bar{\beta}}{2\bar{\varepsilon}} \end{aligned} \quad (23)$$

This first-order condition characterizes the active learning rule (c_1^{SAL}), which is the optimal period 1 consumption choice. Unfortunately, it is not possible to solve this equation explicitly for c_1^{SAL} , but we can determine the direction of experimentation by comparing (23) with the first-order condition of the passive learner (22).

When doing so, it should be noted that $\mathbb{E}_1 \{ \Delta V^L \} > 0$ (information is valuable, *cf.*

¹⁶Formally, (23) assumes that $\bar{\varepsilon} < \bar{\alpha}/2$. This implies that $c^* > 0$, such that there exists a region $(0, c^*)$ wherein learning is guaranteed with probability 1 (recall Section 3). If this weren't the case, there is the theoretical possibility that the active learning motive becomes so strong that it actually calls for setting $c_1^{SAL} < 0$ (even if $c_1^{SPL} > 0$), which is not feasible due to the non-negativity constraint on c_t . But given the magnitude of current emissions (about 50,000 megatonnes per year), it strikes us as extremely implausible that the experimentation motive by itself would call for reducing emissions by such a large amount. Consequently, we are comfortable with ruling this extreme scenario out by imposing $\bar{\varepsilon} < \bar{\alpha}/2$.

equation (20)), while $dP/dc_1|_{c_1=c_1^{SPL}} = -\bar{\beta}/2\bar{\varepsilon} < 0$. The reason for the latter is that $c_1^{SPL} < \bar{\alpha}/\bar{\beta}$ (compare (22) with equations (8) and (9)), which follows from the fact that taking the prospect of learning into account reduces first-period emissions. This 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 active learning 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_b). That is: $c_1^{SAL} < c_1^{SPL} < c_b$.¹⁷

A climate skeptic thus has an incentive to change course, while the irreversibilities associated with emitting greenhouse gases make him want to change “in the safe direction” (recall the discussion following (21)). Hence, *even* climate skeptics have an incentive to argue (or vote) for emission reductions relative to current (in their eyes uninformative) levels. For politicians, a failure to do so could be explained by political motives, but is hard to justify upon grounds of optimality.

So far, we have focused on the case where $\theta\bar{\beta} > 2$. In the alternative scenario where $\theta\bar{\beta} < 2$, it follows from (23) that $c_1^{SAL} < c_1^{SPL}$ still holds when $c_1^{SPL} < c_1^{SNL} (= c_b)$, i.e. the case when $\theta\bar{\beta} + \frac{P(1-\theta)}{1-c_1} > 2$ (recall the end of Section 4.1.2). If the latter condition is not met, we again have that $c_1^{SPL} = c_1^{SNL} (= c_b)$ in which case the derivative $dP/dc_1|_{c_1=c_1^{SPL}}$ is no longer defined as its denominator $|\bar{\beta}c_1 - \bar{\alpha}|$ then collapses to zero (since then $c_1 = c_b = \bar{\alpha}/\bar{\beta}$, at which point the P -function given by equation (4) is non-differentiable). Intuitively, even though the decision maker still wishes to move away from the confounding emission level, he is indifferent as to whether he wants to deviate in the upward or downward direction. But because $\theta\bar{\beta} < 2$ implies that the resource constraint will be binding in the no learning case, it immediately follows that downward experimentation is the only *possible* option left to him (due to the binding constraint, there is simply no room left to experiment by emitting more).

4.2 Should skeptics actually emit less than IPCC believers?

Having said all this, one may wonder: is it actually possible that learning considerations become so strong that a climate skeptic should in fact argue (or vote) for *lower* emissions than a believer of the IPCC analyses?

¹⁷Under explicit incorporation of a discount factor d , the last two terms of (23) would be pre-multiplied by $d/(1-d)$. Note that when $d = 0.5$, this expression equals unity (as is implicitly the case in equation (23)). Since a model period should be thought of as about 15 years in reality, $d = 0.5$ corresponds to an annual discount rate of approximately 4.5 percent, which is at the upper end of the rates used in most climate studies (Nordhaus tends to use 4.3 percent). Using a lower discount rate (like the “Stern Review” did), would only strengthen the active learning motive.

To answer 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):

$$IB : \frac{1}{c_1} = \bar{\beta} \quad (24)$$

$$SAL : \frac{1}{c_1} = \bar{\beta} - \underbrace{(1 - \theta)\bar{\beta}}_I + \underbrace{\frac{P(1 - \theta)}{1 - c_1}}_{II} - \underbrace{\frac{dP}{dc_1}\bigg|_{c_1=c_1^{SPL}}}_{III} \underbrace{\mathbb{E}_1\{\Delta V^L\}}_{IV} \quad (25)$$

In the comparison, we refer to the terms by the Roman subscripts attached to them. *II* is a positive term that raises the RHS of (25) 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 (21) for the intuition behind this).

We have already seen in Section 4.1.3 that *III* is negative while *IV* is a positive term, as a result of which their product is negative. Because of the minus sign preceding it in (25), the product term *III* * *IV* thus raises the RHS of (25) 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 (25) - thereby increasing first period emissions c_1 . This term is unrelated to learning and simply captures the fact that doubts on the link between greenhouse gas emissions and global warming in isolation, justify more consumption of greenhouse gases. This seems to be an argument for climate skeptics who argue that uncertainty on the causes of climate change weakens the case for emission reductions. But as equation (25) nicely shows, that kind of reasoning neglects terms *II*, *III* and *IV*. These terms are all related to the learning process and work towards more conservatism.¹⁸

Whether the “skeptical” or the “believing” first-order condition prescribes the lower emission level, is therefore 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

¹⁸If we would incorporate a discount factor (d), terms *II*, *III*, and *IV* would all be multiplied by $d/(1 - d)$. This makes it more likely that skeptics should argue for lower emissions than believers when d is close to 1. Since the benefits of experimentation lie in the future (while the costs are incurred today), discounting decreases experimentation incentives (Keller and Rady (1999) prove this in a general setup). However, because the information produced by experimentation is not subject to depreciation (the mere presence or absence of a link between greenhouse gas emissions and global temperature is unlikely to change in the future (only the exact intensity may vary)), there is a strong benefit to learning whenever some weight is attached to future generations (this is what the $1/(1 - d)$ -term captures): once we have learned whether the atmospheric presence of greenhouse gases is harmful or not, we can exploit that information “until the end of the world” (*cf.* how the information produced by aggressive sulfur reduction policies in the 1970s and 80s is still useful to us today in preventing acid rain).

IPCC believer (or another skeptic, but one with a higher θ) would.¹⁹ To shed light on that question, consider the derivative of c_1^{SAL} with respect to θ which can be obtained by differentiating equation (23) and using the implicit function theorem (see Appendix C for the derivation):²⁰

$$\frac{\partial c_1^{SAL}}{\partial \theta} = - \left[\frac{1}{(c_1^{SAL})^2} + \frac{\bar{\beta}(1-\theta)(\bar{\alpha}/\bar{\beta} - c_1^{SAL})}{2\bar{\varepsilon}(1 - c_1^{SAL})^2} \right]^{-1} \frac{\bar{\beta}}{2\bar{\varepsilon}} \left[\frac{1}{\theta} - (1 + \log \bar{\beta} + \log(1 - c_1^{SAL})) \right] \quad (26)$$

Since $c_1^{SAL} \leq \frac{1}{2}$ and $c_1^{SAL} < \bar{\alpha}/\bar{\beta}$ (recall Section 4.1.3), the denominator of (26) is always positive. So if we define θ_2 as the value for which the last term in (26) equals zero:

$$\theta_2 \equiv \frac{1}{1 + \log \bar{\beta} + \log(1 - c_1^{SAL})}, \quad (27)$$

we get the following result for the derivative of c_1^{SAL} with respect to θ :

$$\frac{\partial c_1^{SAL}}{\partial \theta} = \begin{cases} < 0 & \text{for } \theta < \theta_2 \\ > 0 & \text{for } \theta > \theta_2 \end{cases} \quad (28)$$

The logic behind the results summarized in equation (28) is as follows. Starting from $\theta = 0$, an increase in θ initially makes the decision maker more cautious (*i.e.* c_1^{SAL} falls). The reason is twofold. Firstly, a decision maker with a higher θ becomes more convinced of the IPCC-hypothesis, which reduces the optimal emission level as the expected damage caused by emissions ($\theta\bar{\beta}$) goes up. Stated in terms of the relevant first-order condition (25), an increase in θ diminishes the importance of term I . Simultaneously, starting from a value of $\theta < \frac{1}{2}$, an increase in θ makes the decision maker also more uncertain: the variance of his belief, given by $\theta(1-\theta)$, increases. This strengthens his learning motive (by increasing the importance of the term $III * IV$), which makes c_1^{SAL} fall even further. Hence, as long as $\theta < \theta_2$, it follows that $\partial c_1^{SAL} / \partial \theta < 0$.

However, as θ continues to rise, it will at some stage cross the point at which the expected value of learning (given by (19)) is maximized, θ_2 .²¹ Beyond that point, the active learning term $III * IV$ becomes less important and first-period consumption under

¹⁹Note: the question being asked here is very different from the issue whether skeptics should reduce emissions relative to *current* levels (which are uninformative in the eyes of skeptics). There, the answer is an unambiguous “Yes” (recall Section 4.1).

²⁰We thank an anonymous referee for suggesting this analysis to us.

²¹Here, the expected value of learning reaches its maximum at a point that is somewhat higher than the maximum variance point $\theta = \frac{1}{2}$ because of the asymmetry in the underlying pay-offs in the different states of the world.

the active learning rule comes back up again with θ . So for $\theta \in (\theta_2, 1)$, we have that $\partial c_1^{SAL}/\partial\theta > 0$ - implying that, in that range, a decision maker with a lower θ (*i.e.* a decision maker who is relatively *less* convinced of the IPCC-hypothesis), should emit *less*.

The implications of all this for the optimal emission level c_1^{SAL} are illustrated in Figure 2. The figure shows that skeptics with $\theta \in (\theta_1, 1)$ should emit less than believers in the IPCC-hypothesis (those with $\theta = 1$).

Figure 2: Impact of changes in θ on c_1^{SAL} for $\bar{\alpha} = 2$, $\bar{\beta} = 7$ and $\bar{\varepsilon} = 0.6$.

Evaluating the impact of θ on the passive learning emission level by analyzing $\partial c_1^{SPL}/\partial\theta$ sharpens this result (again see Appendix C for the derivation):

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

In (29), the denominator is clearly positive, while the numerator is so as well (because $P \leq 1$, $c_1^{SPL} \leq \frac{1}{2}$ and $\bar{\beta} > 2$ (by our assumption that there is a meaningful climate debate)). Consequently, $\partial c_1^{SPL}/\partial\theta|_{P=\bar{P}} < 0$ (note the minus-sign up front).

The negativity of (29) over the entire range for θ implies that our result that, possibly, climate skeptics should be more conservative than IPCC believers *can never occur under passive learning*. The reason is that $\lim_{\theta \rightarrow 0} c_1^{SPL}(\theta) > \lim_{\theta \rightarrow 1} c_1^{SPL}(\theta) = 1/\bar{\beta}$, with $\partial c_1^{SPL}/\partial\theta < 0$ for all $\theta \in (0, 1)$. It is thus really the active learning motive that generates the potential to turn the climate-debate upside-down.

Similarly, one can also analyze how the strength of the active learning channel varies with the other relevant parameters. It is for example possible to show that a higher effectiveness of experimentation (captured by term *III*) promotes a more conservative climate policy for skeptics. By recalling that:

$$\left. \frac{dP}{dc_1} \right|_{c_1=c_1^{SPL}} = -\frac{\bar{\beta}}{2\bar{\varepsilon}}, \quad (30)$$

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. In contrast, a high potential climate sensitivity parameter $\bar{\beta}$ increases the attractiveness of experimentation along this dimension.

However, since $\bar{\beta}$ affects the problem through several channels, this does not yet establish that a larger $\bar{\beta}$ decreases first-period emissions under the optimal active learning strategy. To show this, one must obtain $\partial c_1^{SAL}/\partial \bar{\beta}$ which can be done by applying the implicit function theorem to the first-order condition of an active learner (23). Doing so yields (the derivation can once more be found in Appendix C):

$$\frac{\partial c_1^{SAL}}{\partial \bar{\beta}} = - \left[\frac{1}{(c_1^{SAL})^2} + \frac{\bar{\beta}(1-\theta)(\bar{\alpha}/\bar{\beta} - c_1^{SAL})}{2\bar{\varepsilon}(1 - c_1^{SAL})^2} \right]^{-1} \left[\theta + \frac{1}{2\bar{\varepsilon}} \left(\mathbb{E}_1 \{ \Delta V^L \} + \frac{(1-\theta)(1 - 2c_1^{SAL})}{1 - c_1^{SAL}} \right) \right] \quad (31)$$

Again since $c_1^{SAL} \leq \frac{1}{2}$ and $c_1^{SAL} < \bar{\alpha}/\bar{\beta}$, it follows that $\partial c_1^{SAL}/\partial \bar{\beta} < 0$. This implies that an increase in the supposed impact of emissions on global temperature $\bar{\beta}$, prescribes tighter policies.

It is thus certainly possible that a skeptic (like George W. Bush) should argue (or vote) for *lower* emissions than someone who is convinced that global warming is man-made (such as Al Gore). Our analysis shows that this is more likely to be the case when θ and $\bar{\beta}$ are large and $\bar{\varepsilon}$ is small. Whether the relevant conditions are met in reality should be investigated with more sophisticated models of climate change, but 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 15 years before temperature changes can be attributed to either changes in greenhouse gas emission levels or exogenous factors. 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 extremer 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 since it produces valuable information - no matter how long the learning process takes. The actual degree of experimentation would of course be affected (Should emissions be reduced by 10% or by 50%? And for how long?), but to answer those questions takes a less stylized model. This paper only intends to point out the mechanism and the direction in which current policy should move if we want to bring 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 political aspects of climate policy (such as lobbying and intergenerational issues), the analysis of which we leave for future work. In this respect, the aforementioned cases of “acid rain” and “bee deaths due to pesticides” deserve closer study, as policy makers were able to set strategic issues and myopia aside to combat those problems in a way that is consistent with the “active learning by doing” strategy proposed by this paper. More generally, voters (who also have a say in democratic systems) do not have political-strategic motives, so their electoral preferences should not be affected by them.

With respect to the decision making process, we have followed the lion’s share of the literature in assuming that decisions are made by a single agent. In reality, however, decisions come about after negotiations between the various countries involved, which introduces additional mechanisms (*cf.* Ulph and Ulph (1996)). Strategic considerations may furthermore affect the willingness of individual countries to experiment, but as shown in a general framework by Bolton and Harris (1999), it is not clear in which direction this effect works: on the one hand, the willingness to experiment will be reduced due to a free-rider effect (after all, the benefits from experimentation are public while the costs are private), but the strategic setup also brings an “encouragement effect” which *increases* the incentives for each individual agent to experiment. Which effect dominates cannot be established unambiguously.

Another issue is the fact that emission reductions could require irreversible investment,

which may turn out to be wasted if global warming proves to be exogenous. As shown by Kolstad (1996a), this introduces a second option value that actually makes emission reductions *less* attractive. One could interpret this as suggesting that we should do less to reduce greenhouse gas emissions. Another interpretation, given by Kolstad (1996b), 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 is indeed 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). In that case, the cost of emission reductions is just the utility loss associated with lower consumption of carbon-intensive goods, which is what is captured by our model.²²

While this paper has focused on irreversibilities related to the emission of greenhouse gases, it is also possible that global warming produces irreversible damages (like the extinction of species). These are not modeled in the present paper, but straightforward 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.²³

Finally, although this paper is phrased in terms of the climate debate, 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 we propose has already been illustrated in the debates surrounding acid rain and bee deaths, 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

²²In addition, not all “green” investments are fully wasted if emitting CO₂ turns out to be harmless. Although this is the case for carbon capture systems, 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 exogenous. Moreover, we will have to shift to non-fossil energy sources at some point in the future anyway, and this is likely to be costlier if we wait until this becomes a last-minute operation (Ha-Duong, Grubb, and Hourcade, 1997). In the run up to this transition, increasing the use of gas and oil (instead of coal) would be an attractive option, as oil and gas emit less CO₂ per generated unit of energy than coal does (oil is roughly 10% more efficient, while gas increases efficiency by about 40%). Since almost half of our electricity is currently generated by coal, there lies a huge potential there. Finally, the fact that emission-reducing investments produce valuable information on the causes of global warming, always justifies a certain degree of investment.

²³Keller *et al.* (2004) confirm this intuition in a formal model.

comes to mind is optimal consumer experimentation with addictive goods, such as Apple products, hard drugs, or fatty food (for example: when I am uncertain on the relationship between my weight and my calorie intake, should I experiment by consuming more calories or by consuming less?).

6 Conclusion

Given the popularity of the climate skeptic position, this paper has cast the accompanying debate in a formal framework which enables a normative analysis of what the optimal policy for climate skeptics actually looks like. Typically, such skeptics adhere to a passive policy and argue that the possibility that global warming is exogenous implies that it is optimal *not* to take additional action towards reducing greenhouse gas emissions until we know what causes our climate to change.

This paper has however shown that uncertainty on the causes of global warming does not provide a case for inaction, but yields a particular incentive for action instead. The reason is that the learning process and speed are endogenous, since we are learning from self-generated observations. 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 drives global warming, this implies that these skeptics should argue (or vote) for a policy change that moves us away from current (in their eyes uninformative) emission levels, as such a change of direction produces information on the nature of climate change. One can change direction by either in- or decreasing emissions, but the option of increasing emissions is inferior because the irreversibilities associated with emitting greenhouse gases erode the value of the information that is produced. So where 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”.

Once learning considerations are taken into account, the heated question whether one is an IPCC believer or a climate skeptic thus 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, although for different reasons. 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 should do so for learning considerations. In fact, we have shown that it is even possible that the

debate should actually take the exact opposite form compared to what it currently looks like (since optimality could very well imply that climate skeptics should argue for tighter emission standards than IPCC believers).

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 so as to gain insight into the exact mechanisms at play, which forced us to simplify along several dimensions. A more realistic model would generalize our two-period setup as well as the learning process, while it would also take the irreversibilities in abatement investments (but see footnote 22) and environmental damages into account. On this 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 (while this paper has pointed out that emission-reducing investments also carry an “informational return”, which was not taken into account by the Review). This suggests that our main conclusion is robust to the joint incorporation of these two additional irreversibilities, although this of course remains to be verified through a formal analysis.²⁴

7 Appendix A

In the main text it was assumed that the disturbance term ε has bounded support. As a result, learning was discrete and Bayesian updating took a particularly simple form - thereby maintaining analytical tractability. This appendix shows that the idea that extremem emission levels facilitate the learning process on the causes of global warming (with negative deviations from the confounding emission level being equally 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

²⁴Also see the commentary by O'Neill *et al.* (2006) on this. Drawing from various studies (many of which taking the aforementioned issues lacking from our analysis into account), they conclude that “a clear consensus on a central point is that the prospect of learning does not support the postponement of emissions reductions today”.

$\alpha = 0$ 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 (which still equals $\bar{\alpha}/\bar{\beta}$), 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 3 shows the outcome of a simulation study by Solomon *et al.* (2009). The figure shows what would happen under the IPCC-scenario 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 as they note on p. 1705, a more realistic partial emission reduction would induce similar responses (although obviously somewhat muted).

In both cases one can observe a clear change in the rate at which global temperature increases after emissions have been reduced. So under the IPCC-scenario, a trend break in the control (emissions) translates into a trend break for the endogenous variable (global temperature). 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.

Figure 3: Climate system responses for a ramp of CO2 emissions at a rate of 2% per year to peak CO2 values of 450 to 1200 ppmv, followed by zero emissions.

9 Appendix C

In this appendix, we derive (26), (29) and (31) from the main text. We start by deriving (31). To do so, first consider the first-order condition determining c_1^{SAL} :

$$\frac{1}{c_1} - \theta\bar{\beta} - \frac{P(1-\theta)}{1-c_1} + \frac{dP}{dc_1}\bigg|_{c_1=c_1^{SPL}} \mathbb{E}_1\{\Delta V^L\} = 0, \quad (C1)$$

with $\frac{dP}{dc_1}\bigg|_{c_1=c_1^{SPL}} = -\frac{\bar{\beta}}{2\bar{\epsilon}}$ and $P = \frac{|\bar{\alpha} - \bar{\beta}c_1|}{2\bar{\epsilon}}$

Under passive learning, the endogeneity of P is ignored and we get:

$$\frac{1}{c_1^{SPL}} = \theta\bar{\beta} + \frac{P(1-\theta)}{1-c_1^{SPL}},$$

so given P (*i.e.* on the differentiable submanifold in parameter space where $P = \bar{P}$):

$$\begin{aligned}\frac{-1}{(c_1^{SPL})^2} dc_1^{SPL} &= \theta d\bar{\beta} + \frac{P(1-\theta)}{(1-c_1^{SPL})^2} dc_1^{SPL} \\ \Rightarrow \frac{dc_1^{SPL}}{d\bar{\beta}} \Big|_{P=\bar{P}} &= - \left[\frac{1}{(c_1^{SPL})^2} + \frac{P(1-\theta)}{(1-c_1^{SPL})^2} \right]^{-1} \theta < 0\end{aligned}$$

Incorporating changes in P yields:

$$\begin{aligned}\frac{-1}{(c_1^{SPL})^2} dc_1^{SPL} &= \theta d\bar{\beta} + \frac{P(1-\theta)}{(1-c_1^{SPL})^2} dc_1^{SPL} + \frac{(1-\theta)}{(1-c_1^{SPL})} \left[-\frac{\beta}{2\bar{\varepsilon}} dc_1^{SPL} - \frac{c_1^{SPL}}{2\bar{\varepsilon}} d\bar{\beta} \right] \\ \Rightarrow \frac{dc_1^{SPL}}{d\bar{\beta}} &= - \left[\frac{1}{(c_1^{SPL})^2} + \frac{P(1-\theta)}{(1-c_1^{SPL})^2} - \frac{(1-\theta)}{(1-c_1^{SPL})} \frac{\beta}{2\bar{\varepsilon}} \right]^{-1} \left[\theta - \frac{(1-\theta)}{(1-c_1^{SPL})} \frac{c_1^{SPL}}{2\bar{\varepsilon}} \right]\end{aligned}$$

Next consider active learning. Then terms involving $\frac{dP}{dc_1} \Big|_{c_1=c_1^{SPL}} \mathbb{E}_1 \{ \Delta V^L \}$ come into play. Recall that this term is given by:

$$\frac{dP}{dc_1} \Big|_{c_1=c_1^{SPL}} \mathbb{E}_1 \{ \Delta V^L \} = -\frac{\bar{\beta}}{2\bar{\varepsilon}} (\log(\theta) + (1-\theta) [1 + \log(\bar{\beta}) + \log(1-c_1)]) \quad (\text{C2})$$

Again, let's start by considering the case where we ignore the impact of $\bar{\beta}$ and c_1 on P , so we once more differentiate on the submanifold where $P = \bar{P}$:

$$\begin{aligned}\frac{-1}{(c_1^{SAL})^2} dc_1^{SAL} &= \theta d\bar{\beta} + \frac{P(1-\theta)}{(1-c_1^{SAL})^2} dc_1^{SAL} + \frac{\bar{\beta}}{2\bar{\varepsilon}} \left[\frac{1-\theta}{(1-c_1^{SAL})} (-dc_1^{SAL}) + (1-\theta) \frac{1}{\bar{\beta}} \right] d\bar{\beta} + \frac{\mathbb{E}_1 \{ \Delta V^L \}}{2\bar{\varepsilon}} d\bar{\beta} \\ \Rightarrow \frac{dc_1^{SAL}}{d\bar{\beta}} \Big|_{P=\bar{P}} &= - \left[\frac{1}{(c_1^{SAL})^2} + \frac{P(1-\theta)}{(1-c_1^{SAL})^2} - \frac{\bar{\beta}}{2\bar{\varepsilon}} \frac{1-\theta}{1-c_1^{SAL}} \right]^{-1} \left[\theta + \frac{1}{2\bar{\varepsilon}} (\mathbb{E}_1 \{ \Delta V^L \} + 1 - \theta) \right]\end{aligned}$$

Note that $P = \frac{|\bar{\alpha} - \bar{\beta} c_1|}{2\bar{\varepsilon}}$, so by using that $\frac{1-\theta}{(1-c_1^{SAL})^2} \left[P - \frac{\bar{\beta}}{2\bar{\varepsilon}} (1-c_1^{SAL}) \right] = \frac{\bar{\beta}(1-\theta)}{2\bar{\varepsilon}(1-c_1^{SAL})^2} \left(\frac{\bar{\alpha}}{\bar{\beta}} - 1 \right)$ we can rewrite the denominator as:

$$\frac{dc_1^{SAL}}{d\bar{\beta}} \Big|_{P=\bar{P}} = - \left[\frac{1}{(c_1^{SAL})^2} - \frac{\bar{\beta}(1-\theta)}{2\bar{\varepsilon}(1-c_1^{SAL})^2} \left(1 - \frac{\bar{\alpha}}{\bar{\beta}} \right) \right]^{-1} \left[\theta + \frac{1}{2\bar{\varepsilon}} (\mathbb{E}_1 \{ \Delta V^L \} + 1 - \theta) \right]$$

Now consider the derivative without the restriction $P = \bar{P}$:

$$\begin{aligned}\frac{-1}{(c_1^{SAL})^2} dc_1^{SAL} &= \theta d\bar{\beta} + \frac{P(1-\theta)}{(1-c_1^{SAL})^2} dc_1^{SAL} + \frac{\bar{\beta}}{2\bar{\varepsilon}} \left[\frac{1-\theta}{(1-c_1^{SAL})} (-dc_1^{SAL}) + (1-\theta) \frac{1}{\bar{\beta}} \right] d\bar{\beta} \\ &\quad + \frac{\mathbb{E}_1 \{ \Delta V^L \}}{2\bar{\varepsilon}} d\bar{\beta} + \frac{1-\theta}{1-c_1^{SAL}} \left[-\frac{\bar{\beta}}{2\bar{\varepsilon}} dc_1^{SAL} - \frac{c_1}{2\bar{\varepsilon}} d\bar{\beta} \right]\end{aligned}$$

$$\Rightarrow \frac{dc_1^{SAL}}{d\bar{\beta}} = - \left[\frac{1}{(c_1^{SAL})^2} + \frac{\bar{\beta}(1-\theta)(\bar{\alpha}/\bar{\beta} - c_1^{SAL})}{2\bar{\varepsilon}(1 - c_1^{SAL})^2} \right]^{-1} \left[\theta + \frac{1}{2\bar{\varepsilon}} \left(\mathbb{E}_1 \{ \Delta V^L \} + \frac{(1-\theta)(1 - 2c_1^{SAL})}{1 - c_1^{SAL}} \right) \right]$$

This is equation (31) in the paper.

Finally, we analyze the derivative of c_1^{SAL} with respect to θ . Start out again with (C1).

First assuming $P = \bar{P}$ provides us with:

$$\begin{aligned} \frac{-1}{(c_1^{SPL})^2} dc_1^{SPL} &= \left(\bar{\beta} - \frac{P}{1 - c_1^{SPL}} \right) d\theta + \frac{P(1-\theta)}{(1 - c_1^{SPL})^2} dc_1^{SPL} + \frac{(1-\theta)}{1 - c_1^{SPL}} dP \\ \Rightarrow \frac{dc_1^{SPL}}{d\theta} \Big|_{P=\bar{P}} &= - \left[\frac{1}{(c_1^{SPL})^2} + \frac{P(1-\theta)}{(1 - c_1^{SPL})^2} \right]^{-1} \left[\bar{\beta} - \frac{P}{1 - c_1^{SPL}} \right] \end{aligned}$$

This is equation (29) in the paper.

Incorporating the expression for P and dropping the assumption of fixed P (as a preparation for our analysis of the active learning case) does not change much, since the expression for P does not contain θ :

$$\begin{aligned} \frac{-1}{(c_1^{SPL})^2} dc_1^{SPL} &= \left(\bar{\beta} - \frac{P}{1 - c_1^{SPL}} \right) d\theta + \frac{P(1-\theta)}{(1 - c_1^{SPL})^2} dc_1^{SPL} + \frac{(1-\theta)}{1 - c_1^{SPL}} \left(-\frac{\bar{\beta}}{2\bar{\varepsilon}} \right) dc_1^{SPL} \\ \Rightarrow \frac{dc_1^{SPL}}{d\theta} &= - \left[\frac{1}{(c_1^{SPL})^2} - \frac{(1-\theta)\bar{\beta}}{(1 - c_1^{SPL})^2} \left(1 - \frac{\bar{\alpha}}{\bar{\beta}} \right) \right]^{-1} \left[\bar{\beta} - \frac{P}{1 - c_1^{SPL}} \right] \end{aligned}$$

Next, consider the active learning term (C2). Incorporating it leads to the following additional terms in the expression for the derivative with respect to θ :

$$\frac{\bar{\beta}}{2\bar{\varepsilon}}(1-\theta) \frac{dc_1^{SAL}}{1 - c_1^{SAL}} + \frac{\bar{\beta}}{2\bar{\varepsilon}} \left[(1 + \log(\bar{\beta})) + \log(1 - c_1^{SAL}) - \frac{1}{\theta} \right] d\theta$$

By using the substitutions derived earlier for the terms followed by dc_1^{SAL} we get:

$$\frac{dc_1^{SAL}}{d\theta} = - \left[\frac{1}{(c_1^{SAL})^2} + \frac{\bar{\beta}(1-\theta)(\bar{\alpha}/\bar{\beta} - c_1^{SAL})}{2\bar{\varepsilon}(1 - c_1^{SAL})^2} \right]^{-1} \frac{\bar{\beta}}{2\bar{\varepsilon}} \left[\frac{1}{\theta} - (1 + \log \bar{\beta} + \log(1 - c_1^{SAL})) \right],$$

which is equation (26) of the paper.

10 References

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

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

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

D. Bodansky, Scientific Uncertainty and the Precautionary Principle, *Environment*, 33 (1991) 4-44.

P. Bolton and C. Harris, Strategic Experimentation, *Econometrica* 67 (1999) 349-374.

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

S.J. Farnsworth and S.R. Lichter, The Structure of Scientific Opinion on Climate Change, *International Journal of Public Opinion Research* 24 (2011) 93-103.

C. Gollier and N. Treich, Decision-Making under Scientific Uncertainty: The Economics of the Precautionary Principle, *Journal of Risk and Uncertainty* 27 (2003) 77-103.

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

V. Guemas, F.J. Doblas-Reyes, I. Andreu-Burillo and M. Asif, Retrospective Prediction of the Global Warming Slowdown in the Past Decade, *Nature Climate Change* 3 (2013) 649-653.

M. Ha-Duong, M.J. Grubb and J.C. Hourcade, Influence of Socio-economic Inertia and Uncertainty on Optimal CO₂-emission Abatement, *Nature* 390 (1997) 270-273.

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

G. Immordino, Looking for a Guide to Protect the Environment: The Development of the Precautionary Principle, *Journal of Economic Surveys* 17 (2003) 629-643.

K. Keller, B.M. Bolker, and D.F. Bradford, Uncertain Climate Thresholds and Optimal Economic Growth, *Journal of Environmental Economics and Management* 48 (2004) 723-741.

G. Keller and S. Rady, Optimal Experimentation in a Changing Environment, *Review of Economic Studies* 66 (1999) 475-507.

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

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

C.D. Kolstad, Learning and Stock Effects in Environmental Regulation: The Case of Greenhouse Gas Emissions, *Journal of Environmental Economics and Management* 31

(1996b) 1-18.

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

B.C. O'Neill, P. Crutzen, A. Grübler, M. Ha Duong, K. Keller, C. Kolstad, J. Koomey, A. Lange, M. Obersteiner, M. Oppenheimer, W. Pepper, W. Sanderson, M. Schlesinger, N. Treich, A. Ulph, M. Webster and C. Wilson, Learning and Climate Change, *Climate Policy* 6 (2006) 585-589.

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

G.H. Roe and M.B. Baker, Why Is Climate Sensitivity So Unpredictable?, *Science* 318 (2007) 629-632.

B.D. Santer et al., Separating Signal and Noise in Atmospheric Temperature Changes: The Importance of Timescale, *Journal of Geophysical Research* 116 (2011) 1-19.

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

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

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

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

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

A. Ulph and D. Ulph, Who Gains from Learning about Global Warming?, in: E.C. van Ierland and K. Gorka (Eds.), *Economics of Atmospheric Pollution*, Springer-Verlag, Berlin, 1996, pp. 31-68.

T. Willems, Actively Learning by Pricing: A Model of an Experimenting Seller, mimeo, University of Oxford, 2012.