VU Research Portal

Active Learning and Optimal Climate Policy

Hwang, In Chang; Tol, Richard S.J.; Hofkes, Marjan W.

published in

Environmental and Resource Economics
2019

DOI (link to publisher)

[10.1007/s10640-018-0297-x](https://doi.org/10.1007/s10640-018-0297-x)

document version

Publisher's PDF, also known as Version of record

document license

Article 25fa Dutch Copyright Act

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Hwang, I. C., Tol, R. S. J., & Hofkes, M. W. (2019). Active Learning and Optimal Climate Policy. *Environmental and Resource Economics*, 73(4), 1237-1264. <https://doi.org/10.1007/s10640-018-0297-x>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Active Learning and Optimal Climate Policy

In Chang Hwang¹ · Richard S. J. Tol^{2,3,4,5,6} · Marjan W. Hofkes^{3,4}

Accepted: 26 October 2018 / Published online: 1 November 2018
© Springer Nature B.V. 2018

Abstract

This paper develops a climate-economy model with uncertainty, irreversibility, and active learning. Whereas previous papers assume learning from one observation per period, or experiment with control variables to gain additional information, this paper considers active learning from investment in monitoring, specifically in improved observations of the global mean temperature. We find that the decision maker invests a significant amount of money in climate research, far more than the current level, in order to increase the rate of learning about climate change. This helps the decision maker make improved decisions. The level of uncertainty decreases more rapidly in the active learning model than in the passive learning model with only temperature observations. As the uncertainty about climate change is smaller, active learning reduces the optimal carbon tax. The greater the risk, the larger is the effect of learning. The method proposed here is applicable to any dynamic control problem where the quality of monitoring is a choice variable, for instance, the precision at which we observe GDP, unemployment, or the quality of education.

Keywords Climate policy · Irreversibility · Uncertainty · Learning · Active learning

JEL Classification Q54 · O3 · C63

1 Introduction

It has long been known that expectations on future information affect current decisions. This is typically specified as *passive learning*, where the decision maker updates information about state variables as model time progresses. In some cases, it might be optimal to *learn*

✉ In Chang Hwang
ichwang@si.re.kr

¹ The Seoul Institute, Room 418, 57 Nambusunhwan-ro, 340-gil, Seocho-gu, Seoul 06756, South Korea

² Department of Economics, University of Sussex, Brighton, UK

³ Institute for Environmental Studies, VU University Amsterdam, Amsterdam, The Netherlands

⁴ Department of Spatial Economics, VU University Amsterdam, Amsterdam, The Netherlands

⁵ Tinbergen Institute, Amsterdam, The Netherlands

⁶ CESifo, Munich, Germany

by *experimenting* with control variables. In this paper, we consider *active learning*, i.e., investment in improved monitoring, as an alternative route to increase information, and investigate the implications for climate policy.

There is considerable uncertainty about all aspects of climate change and climate policy. As the consequences of emission abatement decisions are irreversible or at least long-lived, the prospect of future learning, materially affects optimal climate policy in the short-run. The previous literature assumes that new knowledge either arrives exogenously or arrives without cost in the form of new observations about the climate system.¹ However, learning requires investment and can be decelerated or accelerated by monitoring and research. In this paper, we explore simultaneous decisions about learning and abatement.

Uncertainty becomes smaller as the decision maker observes the system. As the state of the system, and thus the information it contains, depends on past decisions, the planner controls, to a certain extent, what is learned. However, in the literature (e.g., Kelly and Kolstad 1999; Kelly and Tan 2015; Hwang et al. 2017), learning is a by-product of decisions on the carbon tax. In such a representation, there are two state variables (e.g., the carbon stock and the stock of knowledge) but only one control variable (the carbon tax). Decisions are therefore second best optimal (Tinbergen 1954). We introduce a second control variable (i.e. research investment) in order to control learning separately from emissions.

In a seminal paper, Prescott (1972) shows that there is a tradeoff between control and information when there is uncertainty about the effect of a policy instrument. He finds it is optimal to sacrifice part of the current benefits of control in order to obtain information that improves future decisions—we refer to this as (*learning by*) *experimentation*. Especially when uncertainty is large or when the time horizon is long, experimentation becomes more important (MacRae 1972; Kendrick 1982).² Later developments include multiple uncertainties (e.g., Wieland 2000a), time-varying parameters (e.g., Beck and Wieland 2002), alternative utility functions (e.g., Johnson 2007), and correlated information (e.g., Marcoul and Weninger 2008).³

In the Bayesian learning literature the decision maker has a prior belief on an uncertain parameter and she updates her belief with new information using Bayes' Rule. If the decision maker has no control over the precision of the new information or the rate at which it arrives, we refer to this as *passive learning*. For instance, Kelly and Kolstad (1999), Leach (2007), Kelly and Tan (2015), and Hwang et al. (2017) investigate the impact of passive learning about the climate system on optimal greenhouse gas emission control. In their models, learning is a by-product of control.

This paper differs from the literature. First and foremost, one of the control variables is used exclusively to increase the speed of learning—we refer to this as *active learning*. Specifically, we manipulate the precision of new observations, or the sharpness of the likelihood function in Bayes' Rule. In previous papers, learning occurs as time progresses (passive learning)

¹ Manne and Richels (1992), Peck and Teisberg (1993), Kolstad (1996a, b), Nordhaus and Popp (1997), Ulph and Ulph (1997), Gollier et al. (2000), and Webster (2002) incorporate exogenous learning. In these studies, information arrives exogenously at some points in time, and thus learning is independent of actions of the decision maker. Kelly and Kolstad (1999), Leach (2007), Webster et al. (2008), Cai et al. (2012), Jensen and Traeger (2013), Lemoine and Traeger (2014), Kelly and Tan (2015), Fitzpatrick and Kelly (2017), Rudik (2016), and Hwang et al. (2017) construct endogenous learning models. The cost of learning is not explicitly considered in their papers, and learning happens as new instances of state variables are realized. Karp and Zhang (2006), Van Wijnbergen and Willem (2015) investigate the effect of such passive learning with theoretical models.

² In the literature, experimentation has been mainly investigated for monetary policy (e.g., Kendrick 1982; Bertocchi and Spagat 1993; Wieland 2000b; Beck and Wieland 2002; Yetman 2003).

³ See Bar-Shalom and Tse (1976), Kendrick (1982, 2005) for non-Bayesian approaches.

or learning is a by-product of control (learning by experimentation). Because learning is a control variable, we have to take the cost of learning into account. Most other papers do not consider the cost of learning explicitly.⁴ The gains from learning are improved decisions, whereas the costs of learning are losses. By balancing the gains and losses, the rate of learning is determined, together with the optimal control rate for emissions.

The implementation of active learning in a climate economy model is worthwhile since public as well as private funders subsidize research activities to learn about climate processes. The rate of learning depends on such efforts. For instance, WMO and UNEP (2010) estimate that global annual expenditures on climate observations are about \$4–6 billion. As a result, observational errors on temperatures have substantially decreased (Kennedy et al. 2011). In addition, our approach highlights the channel of learning. We cannot experiment with the climate system. For climate sensitivity, defined as the equilibrium global warming in response to a doubling of the atmospheric concentration of carbon dioxide, we get only one observation a year (Knutti et al. 2002) from a natural, uncontrolled experiment. Considering the fact that the quality of the data can be improved when the observation is based on a more accurate measurement, our channel of learning (i.e., improving the quality of monitoring) is a useful way to increase the speed of learning, despite the limitations of the data.

This paper proceeds as follows. Section 2 describes the model and methods. The revised DICE model (Nordhaus 2008) is used in this paper. Main revisions include the introduction of a fat-tailed distribution of climate sensitivity and the implementation of active learning about the climate sensitivity. Main results and sensitivity analyses are given in Sect. 3. Section 4 provides conclusions.

2 The Model and Methods

2.1 The Revised DICE Model

We revise the well-known DICE model (Nordhaus 2008) to introduce uncertainty and learning about climate change. Here we focus only on the main revisions. For the full model see “Appendix A”. Unless otherwise noted, initial values for the state variables and parameter values follow Nordhaus (2008). The decision maker in our model chooses the rate of emissions control and the level of research investment for each time period so as to maximize social welfare defined as in Eq. (1). Gross output, net of damage costs and abatement costs, is allocated to consumption, research investment, and gross investment (other than climate research) as in Eq. (2).

$$\max_{\mu_t, R_t} \mathbb{E} \sum_{t=0}^{\infty} L_t \beta^t U(C_t, L_t) \quad (1)$$

$$C_t = \left(1 - \theta_1 \mu_t^{\theta_2}\right) \Omega_t Q_t - I_t - R_t \quad (2)$$

where $\mathbb{E}$ is the expectation operator, U is the utility function, C is consumption, L is labor force, μ is the rate of emissions control, I is gross investment (other than climate research), Ω is the damage function, Q is gross output, R is investment in climate research, β is the discount factor, θ_1 and θ_2 are parameters.

⁴ Strictly speaking, the cost of learning is implicit in their papers in that experimentation with higher carbon emissions may lead to higher damages in the future. The magnitude of this effect has not been formally investigated yet.

The research capital stock accumulates as follows:

$$K_{R_{t+1}} = (1 - \delta_R)K_{R_t} + R_t \quad (3)$$

where K_R is the research capital stock, $\delta_R = 0.15$ is the depreciation rate of the research capital stock (Hall et al. 2010).⁵

Unlike DICE, we apply annual time steps with an infinite time horizon. In order to consider the effect of uncertainty more properly, we set the lower bound of consumption as low as possible and remove the upper bound of temperature increases. Reducing the computational burden, the savings rate is assumed to be exogenous ($I = sQ$).⁶ As shown in Hwang et al. (2017), this assumption does not materially affect the results of DICE and its variants, since the savings rate does not change much for many plausible model specifications. In addition, backstop technology is not considered in the model. With backstop technology climate risk (e.g., catastrophic consumption loss) can be more easily eliminated by setting emissions to be zero at a finite cost.⁷

The temperature response model of the DICE model is:

$$T_{AT_{t+1}} = T_{AT_t} + \xi_1 \{ RF_{t+1} - (\eta/\lambda)T_{AT_t} - \xi_3(T_{AT_t} - T_{LO_t}) \} \quad (4)$$

$$T_{LO_{t+1}} = T_{LO_t} + \xi_4 \{ T_{AT_t} - T_{LO_t} \} \quad (5)$$

where T_{AT} and T_{LO} are atmospheric and oceanic temperature changes, respectively, from 1900, $RF_{t+1} = \eta \ln(M_{AT_t}/M_b)/\ln(2) + RF_{N_t}$ is radiative forcing, M_{AT} is the carbon stock in the atmosphere, M_b is the pre-industrial carbon stock in the atmosphere, RF_{N_t} is radiative forcing from other than greenhouse gas, λ is the equilibrium climate sensitivity, η , ξ_1 , ξ_3 , and ξ_4 are parameters.

This paper follows the tradition of Bayesian statistical decision theory which requires that uncertainty or partial ignorance can be represented as a probability distribution (De Groot 1970). We apply the framework of feedback analysis (Hansen et al. 1984) in order

⁵ $\delta_R = 0.15$ is generally assumed in the literature, although some recent studies find that δ_R is higher than 0.15 (Nadiri and Prucha 1996; Hall et al. 2010; Li and Hall 2016; Cristini et al. 2016). We find that the parameter value does not affect the main results of this paper from some sensitivity analyses (results not shown). The general finding is that research investment stabilizes at a low level, which varies according to δ_R , so as to compensate for the depreciated research capital stock.

⁶ The computational method of this paper is based on Hwang (2017). The method obtains solutions from optimality conditions, and therefore we need to derive optimal policy rules from the first order conditions. This becomes a problem when a model has many control and state variables. Since our model has several nonlinear functions with 3 control variables and 10 endogenous state variables, deriving optimal policy rules is not an easy task. In addition, endogenous savings rate increases the computational burden. For a passive learning model of Hwang et al. (2017), even a single run (with a maximum tolerance level of $1e-4$) was not solved within an hour (with a system of 2.5 GHz Intel Core i7, 16 GB 1600 MHz DDR3), which means it would take several weeks for 1000 Monte Carlo simulations. Regarding the constant savings rate, Hwang (2017) compares the results of the two versions of the DICE model (one with a constant savings rate, another with a flexible savings rate) and finds that there is no significant difference in the results between the two versions. One of the reasons is that the optimal savings rate of the DICE model does not change much over time and it gradually approaches a constant value as the economy approaches the equilibrium. For further discussions and accuracy tests see Hwang (2017). We defer the application of a flexible savings rate to future research.

⁷ As Hwang et al. (2013, 2016) show, the application of fat-tailed distribution of the climate sensitivity with CRRA utility function into the DICE model (as in our model) may lead to a catastrophic consumption loss. It would be optimal for the decision maker to reduce carbon emissions totally (i.e. zero emission) if a backstop technology is available with a finite cost. This is an implication of the Dismal Theorem (Weitzman 2009). The backstop technology in the original DICE model comes at a finite cost.

to introduce a fat-tailed climate sensitivity distribution.⁸ Risk is fat-tailed if the probability density of an uncertain variable falls more slowly than exponentially in the tail (Weitzman 2009). The probability distribution of the climate sensitivity is derived from the distribution of the total feedback factors (Roe and Baker 2007), as

$$\lambda = \lambda_0 / (1 - f) \quad (6)$$

where f is the total feedback factors, which is assumed to be strictly less than 1, and λ_0 is the equilibrium climate sensitivity in a black body planet without any feedbacks.⁹ The total feedback factors denote the aggregate impacts of physical factors such as water vapor, cloud, and albedo on radiative forcing in a way to magnifying the response of the climate system (Stocker et al. 2013).

Substituting Eq. (6) for λ in Eq. (4), replacing radiative forcing with its components, rearranging, and introducing natural variability (stochastic shocks) result in:

$$T_{AT_{t+1}} = (\zeta_1 f + \zeta_2) T_{AT_t} + \zeta_3 \ln(M_{AT_t} / M_b) + \zeta_4 T_{LOt} + \xi_1 RF_{N_t} + \varepsilon_{t+1}^{nv} \quad (7)$$

where ε^{nv} is natural variability normally distributed with mean zero (Brohan et al. 2006), ζ_i ($i = 1, 2, 3, 4$) are adjusted parameters ($\zeta_1 = \xi_1 \eta / \lambda_0$, $\zeta_2 = 1 - \zeta_1 - \zeta_4$, $\zeta_3 = \xi_1 \eta / \ln(2)$, and $\zeta_4 = \xi_1 \xi_3$).

Observed atmospheric temperature is governed by the following equation.

$$T_{AT_{t+1}}^{obs} = T_{AT_{t+1}} + \varepsilon_{t+1}^{obs} \quad (8)$$

where $T_{AT_t}^{obs}$ is the observed temperature change, ε^{obs} is observational errors normally distributed with mean zero including measurement errors and data coverage bias (Webster et al. 2008).

Notice that the climate system is not affected by observed temperature but by actual temperature. Actual temperature is not known to the decision maker and thus is not used for learning. Observed temperature affects the decision maker's belief on the climate sensitivity, resulting in the choice for the rate of emissions control (or carbon tax). The climate system is indirectly affected by observed temperature in this way. We simulate both actual temperature and observed temperature and apply each for its own purpose: actual temperature for the climate system and observed temperature for the learning process.¹⁰

⁸ The framework of feedback analysis is useful in the following reasons: (1) the total feedback factors are observable, unlike the climate sensitivity; (2) it is easy to apply the Bayes' Theorem since the total feedback factors are usually assumed to be normally distributed (Roe and Baker 2007); (3) the resulting climate sensitivity distribution has fat tails.

⁹ f should be strictly bounded above on a physical science basis. We know that λ is not defined for $f = 1$ from Eq. (6) and the climate system cannot be reached an equilibrium if $f > 1$ (Baker and Roe 2009).

¹⁰ $T_{AT_t}^{obs}$ is a function of T_{AT_t} and observational errors, which is changing according to the level of research capital (see Eqs. 8, 11). Actual temperature is determined by the law of physics and natural variability, which do not change over time. Although actual temperature is not known to the decision maker, it affects future temperature as in Eq. (7). The probability of high temperature increases at a certain point in time should be the same for both the passive learning model and the active learning model, if the choices of the decision maker are the same. But since the belief of the decision maker on the state of the world is different, her choices become different between the passive learning model and the active learning model. The change in expectation leads to the different decisions on the stringency of climate policy. If uncertainty is large (or fat-tailed), learning induces large changes in the choices of the decision maker.

Fig. 1 Uncertainty about global mean temperature (Left): The variance of global mean land air temperature (LAT) 1850–2006 (CRUTEM3, Brohan et al. 2006) as a function of the number of weather stations used to estimate the global mean temperature. (Right): The variance of global mean sea surface temperature (SST) 1925–2006 (HadSST3, Kennedy et al. 2011) as a function of the number of observations used to estimate the global mean temperature. The data were obtained from John Kennedy (personal communication)

2.2 Implementation of Active Learning

The literature on decision making under uncertainty and learning about the climate sensitivity assumes that knowledge grows by one observation per year with constant precision (or observational error) (e.g., Kelly and Kolstad 1999; Leach 2007; Webster et al. 2008; Kelly and Tan 2015; Hwang et al. 2017). Instead, this paper considers additional learning through improved monitoring, retaining annual observations. Research investment in the global climate observational system increases the precision not so much of individual temperature observations, but rather in the estimate of the global mean temperature. This in turn lowers estimation errors for the climate sensitivity (see Eq. 12 below). As the standard error in the mean of temperature increases (decreases), the signal to noise ratio falls (grows), making it more difficult (easier) to detect the true state of the world.

As the global climate observational system improves, the standard error in the mean of temperature falls. For instance, as illustrated in Fig. 1, the variance of global mean temperature has decreased over time as the number of weather stations has increased.¹¹

From Eq. (8), the standard error in the mean of observed temperature is decomposed into two elements as follows. We assume that the variance of natural variability is constant over time since natural variability is not controllable by the decision maker. In addition, we assume that natural variability and observational errors are independent.

$$\sigma_{\varepsilon_t}^2 = \sigma_{obs_t}^2 + \sigma_{nv}^2 \quad (9)$$

where σ_{ε} , σ_{ob} , and σ_{nv} are the standard error in the mean of observed temperature, observational errors, and natural variability, respectively.

Broadly speaking, observational errors are linearly related to the reciprocal of the number of observational instruments (Jones et al. 1997; Brohan et al. 2006), at least in the relevant domain (see Fig. 1).¹² Assuming independence between sea surface temperature (SST) observational errors and land air temperature (LAT) observational errors, the total observational

¹¹ The quality of observations as well as the number of observations is important for the standard error in the mean of temperature. For instance, increasing the number of observational system does not necessarily improve the precision of the measurement of climate change if the quality of observations is limited. However, for simplicity we refer the consideration of the quality of observations to future researches.

¹² The other non-linear relationships between uncertainty and research investment would be interesting to study, which is referred to future research.

errors of global mean temperature can be calculated as follows. For simplicity, we assume that observational errors approach zero as investment in the global temperature observational system becomes arbitrarily large.

$$\sigma_{obs,t}^2 = \omega_l \sigma_{obs_{l,t}}^2 + \omega_s \sigma_{obs_{s,t}}^2 = \sum_j \omega_j (\alpha_j / N_{oi,j,t}) \quad (10)$$

where $j \in \{l, s\}$ refers to each observation (l for LAT and s for SST), ω is the respective area of the land or the sea ($\omega_s + \omega_l = 1$), N_{oi} is the number of observational instruments, α_j is a parameter.

Equation (10) leads to Eq. (11), the channel through which research investment affects the uncertainty about temperature shocks.

$$\sigma_{\varepsilon_t}^2 = \omega_l c_l \alpha_l / (p K_{R_t}) + \omega_s c_s \alpha_s / \{(1-p) K_{R_t}\} + \sigma_{nv}^2 = a_R / K_{R_t} + \sigma_{nv}^2 \quad (11)$$

where K_R is the research capital stock for the global temperature observational system, $0 \leq p \leq 1$ is the proportion of money spent on land observations, $c_l \equiv p K_R / N_{o,l}$ and $c_s \equiv (1-p) K_R / N_{o,s}$ are the unit cost of LAT and SST observation, respectively, and $a_R \equiv \omega_l c_l \alpha_l / p + \omega_s c_s \alpha_s / (1-p)$.

The decision maker updates her beliefs on the total feedback factors using Bayes' Rule as follows (Cyert and De Groot 1974; Lemoine 2010):

$$p(f | T_{AT}^{obs}) \propto p(T_{AT}^{obs} | f) \times p(f) \quad (12)$$

where $p(f)$ is the prior distribution, $p(T_{AT}^{obs} | f)$ is the likelihood function (the likelihood of T_{AT}^{obs} observation given f), and $p(f | T_{AT}^{obs})$ is the posterior distribution. Since the temperature shocks are assumed to be normally distributed with mean zero, more precise temperature observations lead to an improved decision on carbon tax.

The normal distribution of the total feedback factors with parameters $\bar{f}_t$ and σ_{f_t} is used as the initial prior for the year 2005 (Roe and Baker 2007). The resulting posterior mean and the posterior variance of the total feedback factors are:

$$\bar{f}_{t+1} = \frac{\bar{f}_t + \zeta_1 T_{AT_t}^{obs} H_{t+1} v_t / v_{\varepsilon_t}}{1 + \left(\zeta_1 T_{AT_t}^{obs} \right)^2 v_t / v_{\varepsilon_t}} \quad (13)$$

$$v_{t+1} = \frac{v_t}{1 + \left(\zeta_1 T_{AT_t}^{obs} \right)^2 v_t / v_{\varepsilon_t}} \quad (14)$$

where $\bar{f}_t$ and $v_t = \sigma_{f_t}^2$ are the mean and the variance of the total feedback factors, $v_{\varepsilon_t} = \sigma_{\varepsilon_t}^2$ is the variance of temperature shocks, and $H_{t+1} \equiv T_{AT_{t+1}}^{obs} - \zeta_2 T_{AT_t}^{obs} - \zeta_3 \ln(M_{AT_t} / M_b) - \zeta_4 T_{LO_t} - \zeta_5 R F_{N_t}$.

The posterior distribution with parameters $\bar{f}_{t+1}$ and v_{t+1} of Eqs. (13) and (14) serves as the prior for the next time period. In this way the decision maker learns about the true value of the total feedback factors for each time period. Note that the parameters of the posterior distribution are affected by research investment through Eqs. (11) and (14) and the parameters become endogenous state variables. The higher research investment, the lower is the variance of the total feedback factors. From Eq. (6) the probability distribution of the climate sensitivity can be derived from the probability distribution of the total feedback factors.

2.3 Calibration

In order to calibrate a_R and σ_{nv}^2 in Eq. (11) we use global expenditures on temperature observations. We take a bottom-up approach. Global mean LAT is calculated from the records of each country's weather stations and global mean SST is calculated from the reports of observational platforms such as ships, drifting buoys, and moored buoys (Kennedy et al. 2011). With this in mind, we multiply the number of each observational platform and the unit cost of each instrument (see Table 1). Annual operational costs for temperature observational instruments are estimated to be about \$450 million in 2005.¹³ The total installation costs for all the existing instruments are about \$500 million.¹⁴ Second, σ_{nv}^2 is calculated as the difference between the total variance of temperature shocks ($=0.10^2$) estimated by Tol and De Vos (1998) and the variance of observational errors ($=0.06^2$) obtained from the HadCRUT4 dataset (Morice et al. 2012). Then we can postulate that the research capital stock ($K_{R_0} = \$950$ million) results in the variance of temperature shocks ($\sigma_{\varepsilon_0}^2 = 0.10^2$) (Eq. 11). So, $a_R = \$3.42$ million.¹⁵

As shown in Fig. 2, our parameterization implies that the variance of temperature shocks decreases (increases) as the research capital stock increases (decreases). If there is no change in the research capital stock, the variance of observational errors (in turn, the variance of temperature shocks) remains the same.¹⁶

We set the level of satisfaction of the decision maker about the level of learning, mainly for computational reasons. This assumption also reflects the point that the decision maker does not want not make an effort to reduce the variance of observational errors if she thinks there has been enough learning. We can think of this as a budget constraint. This puts an upper bound on the research investment (in turn, on the research capital stock).¹⁷ More specifically, we set $R_t = 0$ if $\sigma_{\varepsilon_t}^2 - \sigma_{nv}^2 < \omega_c$, where ω_c reflects the level of satisfaction of the decision maker on the level of learning. According to our parameterization ($\sigma_{nv}^2 = 0.064$, $\omega_c = 10^{-6}$), if the variance of the temperature shocks $\sigma_{\varepsilon_t}^2$ (always higher than σ_{nv}^2 , by assumption) becomes close

¹³ For comparison, the United States spent \$140 million on in situ climate observations in 2010 (submission of USA to UNFCCC/SBI 35). WMO and UNEP (2010) estimate that annual global expenditure on climate observations is about \$4–6 billion. Douglas-Westwood (2006) estimates that the total costs of ocean observations are \$402 million in 2005. Their estimates are not directly comparable to this paper, however, because their estimates include all kinds of observations besides temperature, such as precipitation, wind, ice, as well as satellite observations.

¹⁴ This number is fairly small compared to the world economy. The initial value for the global capital stock (in 2005) is \$137 trillion in the original DICE model. Thus the research investment in climate observations has a negligible effect on the growth path of the world economy. In this case, cost–benefit analysis rather than a complex dynamic stochastic model can be used to calculate the optimal research investment. Nevertheless, our model is useful for the investigation of the effect of learning on optimal policy simultaneously with research investment. In addition, the literature on learning about climate change generally applies dynamic stochastic learning models (e.g., Kelly and Kolstad 1999; Kelly and Tan 2015).

¹⁵ We do not separate the operational costs and the installation costs for these calibrations, for simplicity. That is, we assume that the research capital stock is the sum of the operational costs and the installation costs. An alternative is to explicitly represent operational costs in the model, but this does not affect the main results of this paper (results not shown).

¹⁶ Our calibration may overestimate the benefit of the research investment because temperature observations are not just used for learning about the climate sensitivity, but also for other various activities including weather forecast, agriculture, and transportation (AATSE 2006). For the decomposition of the cost of temperature observations according to the sector of use, we refer to future research.

¹⁷ From Eqs. (3) and (11) this assumption serves as an upper bound of research investment ($R_t \leq a_R/\omega_c - (1 - \delta_R)K_{R_t}$). For instance, with $\omega_c = 10^{-6}$ and the above parameterizations (i.e., $a_R = \$3.42$ million, $K_{R_0} = \$950$ million, $\delta_R = 0.15$), the upper bound of the initial research investment is about \$3.42 trillion.

Table 1 Global temperature observational system in 2005

	Number of instruments/observations (thousands)	Unit cost (1000 US\$)			
		Installation		Operation (per year)	
		Low	High	Low	High
LAT					
Weather station	3455	40		60	
SST					
Number of instruments					
VOS	5429			4	55
Drifting Buoy	1267			4.5	7.8
Moored Buoy	194	1150	2700	200	500
Number of observations					
VOS	1169			0.00023	
Drifting Buoy	1632				
Moored Buoy	179				
Sum	2980				

The number of land weather stations is the one used for building the database CRUTEM4 (Kennedy et al. 2011). The number of voluntary observing ships, drifting buoys, and moored buoys are available at www.bom.gov.au/jcomm/vos and www.aoml.noaa.gov/phod/dac. The unit cost for land weather station is drawn from Mburu (2006). The unit costs for voluntary observing ships, drifting buoys, and moored buoys follow Kent et al. (2000), Meldrum et al. (2010), and Detrick et al. (2000), respectively. The unit cost for data transmission using satellite communication systems is about \$0.23 per observation (North 2007). The number of SST observations is drawn from Kennedy et al. (2011)

Fig. 2 Hypothetical learning dynamics No change refers to the case where the research capita stock remains the same as in the initial year. +\$X/yr (respectively, − \$X/yr) refers to the case where the research capital stock increases (resp., decreases) \$X every year from the initial level

to (but never equals to) the variance of natural variability σ_{nv}^2 through research investment, say $0.064 < \sigma_{\epsilon_t}^2 < 0.064001$ (less than 0.002 percent higher than $\sigma_{nv}^2 = 0.064$), the decision maker does not invest money for learning.

This upper bound is binding after a certain year. For instance, when we set $\omega_c = 10^{-6}$, the upper bound of the research investment is binding after 2045. The problem is that the solution time increases a lot if there is no such an upper bound. For instance, a single run without the upper bound takes more than 30 min. With the upper bound, however, 1000 Monte Carlo simulations take about 6 h. The results between the two cases are a bit different as shown in Fig. 11 (“Appendix B”).

2.4 Computational Methods

In order to solve our large model, we apply the computational method of Hwang (2017), who refines the earlier work by Maliar and Maliar (2005) and Judd et al. (2011).¹⁸ The problem is reformulated in a recursive way as:

$$W(s_t, \theta_t) = \max_{c_t} [U(s_t, c_t, \theta_t) + \beta \mathbb{E}_t W(s_{t+1}, \theta_{t+1})] \quad (15)$$

$$W(s_t, \theta_t) \approx \sum_{n=1}^N \psi(s_t, \theta_t; \mathbf{b}_n) \quad (16)$$

where W is the value function starting from period t , $\mathbf{c}$ is the vector of control variables (μ , R), $\mathbf{s}$ is the vector of state variables (K , K_R , M_{AT} , M_U , M_L , T_{AT}^{obs} , T_{LO} , $\bar{f}$, v , L , A , σ), M_U and M_L are the carbon stocks in the upper ocean and the lower ocean, respectively, σ is the emissions-output ratio, θ is the vector of uncertain variables (f , ε), ψ is the basis function, and $\mathbf{b}$ is the vector of coefficients for the basis function.

The solution algorithm is summarized as follows. First, approximate the value function with a flexible basis function. Second, derive the first order conditions for optimal policy rules. Third, choose an initial guess on the coefficients $\mathbf{b}$ of the basis function: $\mathbf{b}^{(0)}$. Fourth, simulate a time series of variables satisfying the first order conditions, transitional equations, and boundary conditions with the initial guess $\mathbf{b}^{(0)}$.¹⁹ Fifth, calculate the left hand side and the right hand side of the Eq. (15) using the simulated time series, and then find $\mathbf{b}$ that minimizes the difference between them: $\hat{\mathbf{b}}$.²⁰ Sixth, update the initial guess $\mathbf{b}^{(0)}$ using a pre-specified updating rule: $\mathbf{b}^{(1)}$. Seventh, iterate the above process with the new guess $\mathbf{b}^{(1)}$ until the value function converges.²¹

¹⁸ Hwang (2017) takes advantage of the fact that the first order conditions of DICE result in global solutions (Solak et al. 2015). One of the advantages of the simulation-based method of Hwang (2017) is that it reduces the computational burden compared to the grid-based methods in the literature because it searches for solutions on a set satisfying the first order conditions. See Hwang (2017) for more discussions on the solution method.

¹⁹ The simulation length is set at 1000 years. Longer horizons do not affect the main results of this paper.

²⁰ The Gauss-Hermite integration is applied for the expectation in Eq. (15) with 10 integration nodes. Higher number of nodes does not affect the main results of this paper.

²¹ The maximum tolerance level is set at 1.0×10^{-4} , which is small relative to net present welfare. Existing numerical models generally apply a similar tolerance level (e.g., Kelly and Kolstad 1999; Cai et al. 2012; Lemoine and Traeger 2014). The more stringent tolerance level gains a bit more accuracy at the expense of the speed of convergence. Actually, changing the tolerance level from 1.0×10^{-4} to 1.0×10^{-5} does not change our results significantly. When the stopping criterion holds, the relative difference of two consecutive policy functions (i.e., $\max \left[\left| \frac{\mu_t^{(p+1)} - \mu_t^{(p)}}{\mu_t^{(p)}} \right|, \left| \frac{R_t^{(p+1)} - R_t^{(p)}}{R_t^{(p)}} \right| \right]$, where p refers to the p th iteration) is 6.2×10^{-4} . An alternative is to compare two consecutive policy functions for stopping criterion (i.e., $\max \left[\left| \frac{\mu_t^{(p+1)} - \mu_t^{(p)}}{\mu_t^{(p)}} \right|, \left| \frac{R_t^{(p+1)} - R_t^{(p)}}{R_t^{(p)}} \right| \right] \leq \omega$, where ω is the maximum tolerance level) and compare the relative difference of value functions. We find that this does not affect our main results. If the relative difference of two consecutive policy functions was set at 1.0×10^{-3} for the tolerance level, we

Accounting for random realizations of the uncertain variables, the model is run 1000 Monte Carlo simulations and the average of all simulations is presented. For additional results, see “Appendix B”. The true value of the total feedback factors is set at 0.6 (which corresponds to the equilibrium climate sensitivity $3\text{ }^{\circ}\text{C}/2 \times \text{CO}_2$, the most likely value according to the current scientific knowledge) (Stocker et al. 2013) for the reference case. The other true values are considered in Sect. 3.4.

For simulations, the initial values for $\bar{f}_t$ and v_t are assumed to be 0.65 and 0.13^2 , respectively, following the current scientific knowledge (Roe and Baker 2007). The models are also simulated with different initial beliefs but the general implications of these simulations do not change (results not shown). Since the total feedback factors are bounded above, the posterior distribution is derived first with the conjugate normal prior, and then an upper bound ($\bar{f}_t \leq 0.999$) is set for simulations following Hwang et al. (2017). The upper bound corresponds to the climate sensitivity of $1200\text{ }^{\circ}\text{C}/2 \times \text{CO}_2$, which is far higher than any admitted values. Note that the most likely value for the equilibrium climate sensitivity is $3\text{ }^{\circ}\text{C}/2 \times \text{CO}_2$ (Stocker et al. 2013). Hwang et al. (2017) find that higher upper bounds than this do not affect the main results. Since the atmospheric CO_2 concentration is expected to be doubled around 2100, roughly speaking, the climate sensitivity is a proxy for temperature increases around 2100.

3 Results

3.1 The Rate of Learning

Figure 3 shows the evolution of the climate sensitivity distribution. For comparison, the results of the model with passive learning are also presented. As expected, the mean parameter $\bar{f}$ converges to the pre-specified true value and the variance parameter v approaches—but never reaches—zero over time. The rate of learning, measured as the reduction in the (simulated) coefficient of variation of the climate sensitivity (Webster et al. 2008; Hwang et al. 2017), is higher under active learning than under learning only from temperature observations. It takes 44 years for the coefficient of variation to be reduced by 50% for improved observations, whereas it takes 57 years in the passive learning model. For comparison, the learning time for 50% reduction in the coefficient of variation of the climate sensitivity is about 60–70 years in Webster et al. (2008) when the prior similar to the current paper is used (see Fig. 10 of their paper).²² This is because by construction, learning in this paper constitutes an additional way to produce information.

Footnote 21 continued

get the same optimal solutions and the relative difference of two consecutive value functions is 1.0×10^{-4} . For more error analyses see “Appendix C”.

²² The rates of learning in Kelly and Kolstad (1999), Leach (2007), and Kelly and Tan (2015) are not directly comparable to the current paper since they define learning differently from ours: learning takes place in their models when the mean of the uncertain variable becomes statistically close (e.g., at a significance level of 0.05) to the pre-specified true value.

Fig. 3 Climate sensitivity distribution (Top left): The mean of the total feedback factors (Top right): The (simulated) coefficient of variation of the climate sensitivity. (Bottom left): The climate sensitivity distribution in 2055 ($0\text{--}10\text{ }^{\circ}\text{C}/2 \times \text{CO}_2$). (Bottom right): The climate sensitivity distribution in 2055 ($10\text{--}20\text{ }^{\circ}\text{C}/2 \times \text{CO}_2$)

Table 2 shows the corresponding probability of high temperature increases. The probability density in the upper tail of the climate sensitivity distribution shrinks faster for active learning than for passive learning. Therefore the probability of high temperature increases is lower in the active learning model than in the passive learning model. For instance, the probability of temperature increases higher than $6\text{ }^{\circ}\text{C}$ in 2105 for active learning is more than 5 times lower than the one for passive learning case. This is the reason why there is a substantial need for research investment, and the reason why there is a difference in carbon tax between the active learning case and the passive learning case (Sect. 3.3).

3.2 Research Investment

The optimal level of investment in climate research is much higher than the current level. For instance, the initial level of investment in the global climate observational system is about \$1.2 trillion.²³ These are the reasons. First, the benefit of learning is far greater than the cost of learning (Keller et al. 2007a, b; Baehr et al. 2008). For comparison, with a discrete uncertainty representation and exogenous learning, Peck and Teisberg (1993) estimate that the value of information on the climate sensitivity is \$148 billion, and Nordhaus and Popp (1997) estimate that the value of information on the climate sensitivity is \$6.9–11.7 billion.²⁴

²³ We conduct sensitivity analyses on the level of satisfaction of the decision maker on learning (ω_c). If the decision maker wants less precise observations (see Sect. 2.3), the amount of money spent on the global observational system should be decreased. For instance, if the decision maker sets a criterion that $\omega_c = 10^{-5}$ (10^{-4} , respectively), instead of 10^{-6} in the reference case, the level of investment is \$341 billion (\$34 billion, resp.). For those cases, however, the boundary condition (or the budget constraint) binds from the first period.

²⁴ Peck and Teisberg (1993) introduce 3 states of the world: 1, 3, and $5\text{ }^{\circ}\text{C}/2 \times \text{CO}_2$. The value of information in their model is calculated as the difference in (monetized) expected utility between instant learning and

Table 2 The probability of high temperature increase

	2055			2105		
	Passive learning, true $\lambda = 3$		Active learning	Passive Learning, true $\lambda = 3$		Active learning
	True $\lambda = 2$		True $\lambda = 3$	True $\lambda = 4,5$		True $\lambda = 2$
						True $\lambda = 3$
Prob. of $\lambda > 4,5$	0.158	0.045	0.115	0.215	0.050	3.511E-04
Prob. of $\lambda > 6$	0.053	0.010	0.031	0.068	0.006	1.089E-05
Prob. of $\lambda > 10$	0.009	0.001	0.004	0.010	1.665E-04	6.419E-08
						7.766E-06
						4.152E-03

λ is the equilibrium climate sensitivity (unit: $^{\circ}\text{C}/2 \times \text{CO}_2$). For all simulations, $\mathbb{E}_0 f = 0.65, \sigma_f = 0.13$

Keller et al. (2007b) estimate that the value of information on the climate sensitivity is about \$10 billion.²⁵ Keller et al. (2007a) estimate that the value of information associated with early detection of changes in the North Atlantic meridional overturning circulation (MOC) is in the order of tens of billions of dollars, which is far higher than the cost of MOC observation systems (tens of millions of dollars, see Baehr et al. 2008). Baker and Solak (2010) estimate that the optimal level of investment in energy technology is on the order of tens of billions of dollars.²⁶ Since the current paper deals with the fat-tailed distribution of the climate sensitivity it is not surprising that the benefit of learning is greater in our model than in the literature (Weitzman 2009).

Second, the research investment is small compared to the world economy. It is less than 1% of the global capital stock (see footnote 14). Of course, no one country would agree to pay the cost but we are concerned with a global model. In addition, the total cost of climate change (damage cost + abatement cost) can be largely reduced (see Fig. 4) when we account for (active) learning because the decision maker can rule out, to some extent, the possibility of tail events (extreme temperature increases) from learning, and thus she spends less money for emissions control. For instance, the net present value of the avoided cost (the total cost of the uncertainty case minus the total cost of the active learning case) for the next two centuries is about \$ 9.2 trillion (applying the discount rate of DICE), which is 2.1% higher than the net present value of the total research investment during the next two centuries.²⁷

Third, we model R&D as an unrestricted flow. Future research should consider capacity and expansion limits on R&D.

After the large initial investment, as shown in Fig. 4, the research investment decreases sharply. The research investment from the second year is largely due to the depreciation of the research capital stock.²⁸ The research capital stock reaches an upper bound around 2045, implying that the decision maker satisfies with the level of learning after 2045, and thus does not spend money for improved observations (see Sect. 2.3) except the compensation for the depreciated research capital stock. This reflects the point that early investment to reduce uncertainty is more beneficial because (1) it benefits a longer future and (2) knowledge saturates in our model specifications. Note that research investment is here mostly in equipment rather than specialist personnel, so that a rapid scaling-up and -down is feasible. These results imply that the optimal decision is to make uncertainty as low as possible since the cost of learning is much lower than the benefit of learning. Figure 4 also shows some sensitivity analyses on the cost of learning. The research capital stock increases as the cost of learning increases, as implied by Eq. (11).

Footnote 24 continued

no-learning. If learning in 40 years is considered, the value of information is \$24 billion. Nordhaus and Popp (1997) consider 5 discrete climate sensitivities: mean, ± 1 standard deviation, and ± 2 standard deviation. The value of information in their model is calculated as the difference in (monetized) expected utility between instant learning and learning in 50 years.

²⁵ They consider 3 climate sensitivities. They also find that if there is a temperature limit of 2.5 °C the value of learning about the climate sensitivity is \$800 billion.

²⁶ AATSE (2006) and Cristini et al. (2016) find that the benefit of ocean observations is far larger than the cost of observations, although they do not account for learning about the climate sensitivity.

²⁷ For a more thorough discussion, we need to consider the opportunity cost of research investment and the general equilibrium effect. Since our model is not well suited for such an analysis, we refer to future research.

²⁸ About 90% of the new investment after the initial year is to compensate for the depreciated research capital stock (results not shown).

Fig. 4 Research investment and the total cost of climate change (Top Left): The research investment (Top Right): The research capital stock (Bottom): The total cost of climate change (the gross output net of consumption, investment, and research investment) for the reference case and the uncertainty case

3.3 Carbon Tax

It is well known that fat-tailed risk substantially increases the stringency of climate policy. Even in the theoretical framework of Weitzman (2009), learning does not matter because the marginal damage cost (or willingness to pay for a reduction of carbon emission) becomes arbitrarily large under fat-tailed risk. In this case, the fat tail is not thinned by learning. However, as argued in the literature (Costello et al. 2010; Millner 2013; Horowitz and Lange 2014; Hwang et al. 2013, 2016), in a realistic setting such as bounded radiative forcing, the marginal damage cost can be bounded. Kelly and Tan (2015) and Hwang et al. (2017) apply this framework and find that learning substantially reduces the effect of fat-tailed risk because learning is faster in the tail. Policy recommendations such as carbon tax are substantially different from the no-learning case. The current paper follows this framework but extends the approach: active learning. The optimal carbon tax is calculated as a Pigovian tax (Nordhaus 2008).²⁹ As expected, the optimal carbon tax is highest for the uncertainty model and is

²⁹ The carbon tax is calculated as follows: $carbontax_t = -\tau(\partial W / \partial E_t) / (\partial W / \partial K_t)$, where W is social welfare defined as the discounted sum of population-weighted utility of consumption over the whole time period (Eq. 1), E is carbon emission, K is the capital stock, and τ is a constant. In our specification, with a climate sensitivity that is known to the analyst but not to the decision maker, social welfare differs with the true value of the climate sensitivity. The effect of the marginal change in carbon emission (and also in the capital stock) on social welfare also varies with the true value of the climate sensitivity.

Fig. 5 The optimal carbon tax

lowest for the deterministic model (see Fig. 5, Table 3). Active learning substantially reduces the carbon tax because it takes away so much uncertainty.³⁰

As mentioned in Sect. 2, there are limits to learning. The variance of natural variability remains fixed at 0.08^2 for the reference case. This serves as the lower bound for the variance of temperature shocks. The sensitivity of the optimal carbon tax to this lower bound is shown in Fig. 6. The optimal carbon tax decreases if the lower bound of temperature shocks decreases. The higher the magnitude of learning, the lower is the optimal carbon tax.

Table 3 presents the results when the damage function of Weitzman (2012), a highly reactive damage function, is applied.³¹

$$\Omega_t = 1 / \left[1 + \pi_1 T_{AT_t} + \pi_2 T_{AT_t}^{\pi_2} + \pi_3 T_{AT_t}^{\pi_4} \right] \quad (17)$$

where $\pi_1 = 0$, $\pi_2 = 0.0028388$, $\pi_3 = 0.0000050703$, and $\pi_4 = 6.754$. For the DICE damage function, $\pi_1 = 0$, $\pi_2 = 0.0028388$, and $\pi_3 = \pi_4 = 0$.

The effect of the fat tail on climate policy is substantial when the damage function of Weitzman (2012) is applied (see the results of the uncertainty model), but is greatly reduced when learning is introduced. Research investment further enhances the effect of learning (compare the cases where the true value of the climate sensitivity is $3^\circ\text{C}/2 \times \text{CO}_2$). This signifies the importance of learning, especially when the damage function is highly responsive to higher temperature increases. Because we can rule out, to some extent, the impact of tail events (say, temperature increases of 4.5°C or more) from learning, the intensity of climate policy can be largely reduced. The value of research investment is that the decision maker can effectively reduce the probability of substantially high temperature increases as shown in Table 2. Thereby active learning helps to avoid an unnecessarily large amount of expenditure on emission reductions. Notice that the optimal carbon tax is a measure of the intensity of climate policy in this paper.

³⁰ Considering the amount of carbon emissions and the use of carbon tax as a reference value for public policy making (Nordhaus 2017), a seemingly small difference in carbon tax rates between the cases in Table 3 is, de facto, a very important difference. Annual CO_2 emissions from the combustion of fossil fuels and cement production are about 9.5 GtC in 2011 (Stocker et al. 2013). Simply put, one-dollar difference in carbon tax rate (\$1/tC) leads to a difference of about \$10 billion a year.

³¹ The difference between the two damage functions becomes significant when temperature increases are higher than 3°C or more (Hwang et al. 2013).

Table 3 The optimal carbon tax in 2015 (US\$/tC)

	Deterministic true $\lambda = 3$	Uncertainty $\bar{f} = 0.65, \sigma_f = 0.13,$ true $\lambda = 3$	Passive learning $\mathbb{E}_0 f = 0.65, \sigma_f = 0.13$		Active learning $\mathbb{E}_0 f = 0.65, \sigma_f = 0.13$	
			True $\lambda = 2$	True $\lambda = 3$	True $\lambda = 2$	True $\lambda = 3$
			True $\lambda = 4.5$	True $\lambda = 4.5$	True $\lambda = 4.5$	True $\lambda = 4.5$
DICE damage function	32.0	39.0	31.1	37.6	29.1	35.5
Weitzman's damage function	37.7	201.2	43.9	56.4	33.9	53.2

λ is the equilibrium climate sensitivity (unit: $^{\circ}\text{C}/2 \times \text{CO}_2$)

Fig. 6 Limits to learning and the carbon tax (Left): The evolution of the learning parameters (Right): The optimal carbon tax

Fig. 7 Sensitivity analysis (the true value of the climate sensitivity) CS refers to the climate sensitivity. (Top panels): The climate sensitivity distribution in 2055 (Bottom left): Atmospheric temperature (Bottom right): The optimal carbon tax

3.4 The True Value of the Climate Sensitivity

In the real world, the true value of the climate sensitivity is not known with certainty. This subsection investigates how sensitive our main results to the different true values of the climate sensitivity. Figure 7 and Table 2 show some results. Roughly speaking, future temperature evolves according to the true value of the climate sensitivity (although not known to the decision maker) and the decisions made with regard to the states resulting from the current decisions (e.g., emissions control rate, research investment). Other things being equal, future temperature will be higher if the true value of the climate sensitivity is higher as shown in Fig. 7. We also observe that the higher the true value of the climate sensitivity, the fatter

Fig. 8 Sensitivity analysis ($CS=4.5^{\circ}C/2 \times CO_2$) CS refers to the climate sensitivity (Left): The climate sensitivity distribution (the right segment higher than $10^{\circ}C/2 \times CO_2$) (Right): The optimal carbon tax

is the right tail of the climate sensitivity distribution. Since the decision maker considers future damages for the current decisions (Bellman's principle of optimality) (Bellman and Dreyfus 1962), the rate of emissions control or the optimal carbon tax should be higher (lower, respectively) for the higher (lower, resp.) true value cases.

Table 3 and Fig. 8 show the cases where the true value of the climate sensitivity (λ) is different from the initial guess. The initial belief on λ is the same as the reference case. Other things being equal, as expected, the optimal carbon tax is higher if the true λ is higher. The optimal carbon tax for the active learning model is lower than the one for the passive learning model, as in the reference case. Applying the damage function of Weitzman, the effect of active learning becomes more prominent.³²

4 Conclusions

We define active learning as investment in knowledge acquisition. Active learning is contrasted to passive learning where new information arrives without cost at a fixed rate and precision. We investigate the effect of active learning on greenhouse gas emission reduction, noting that climate policy is typified by (deep) uncertainty and irreversibility. The decision maker chooses to reduce uncertainty through significant investment in climate research, more than three orders of magnitude greater than the current level of expenditure. This helps the social planner make better decisions on climate policy. The level of uncertainty decreases more rapidly with active learning than with passive learning. As a result, the optimal carbon tax is lower for active learning than for passive learning, which in turn is lower than the carbon tax without learning. The effect of learning is more pronounced as the tail risk increases. Earlier investment in climate research is more beneficial than later investment.

This paper is the first to introduce active learning into an integrated assessment model of climate and the economy. Applying alternative ways of learning would help to understand the role of learning further including reconstructions of past temperature which would increase the number of observations in the likelihood, and investment in climate research which would sharpen the prior.

³² These results may be sensitive to the model specifications including the probability distribution, the true value of the climate sensitivity, damage function, preference parameters, etc. Further analysis is deferred to future research.

This paper applies the solution method demonstrated in Hwang (2017). His method allows us to solve multiple iterations of the model with many state variables quickly for the Monte Carlo analysis. Hwang (2017) suggests that this should come at little cost of model accuracy. If it does reduce accuracy, the likely impacts on the results are higher savings rates as implied by Hwang (2017). However, this issue is minor in our model because we apply a constant savings rate. As shown by Hwang (2017), the difference of the results obtained from the two solution methods is minimal when a constant savings rate is applied. Even if that were the case, our paper demonstrates new methodology that can be applied to richer models as computer resources improve.

Active learning apart, this paper closely follows Nordhaus' DICE model. Other specifications should be explored, including alternative utility functions (e.g., Sterner and Persson 2008), different mitigation cost and climate impact functions, and other economic growth models. Perhaps more importantly, we here posit a true value of the climate sensitivity, rather than a probability density function, which exaggerates the effects of learning. We only consider one uncertain parameter, which suppresses both the effect of uncertainty and the cost of learning. We study the case of a global planner. Kolstad and Ulph (2008, 2011) show that uncertainty enhances cooperation. This would imply that active learning fosters free-riding. All these matters are deferred to future research.

Active learning by improved monitoring also applies to other areas of public policy. Learning by experimentation with policy variables is informative for issues with a short characteristic life time—monetary policy, for instance—but less so for issues that span long periods—besides climate change, pensions, education and structural unemployment come to mind. The method proposed here applies to any area in which knowledge of the response to policy is imperfect partly due to imperfect monitoring. Although not always acknowledged, most economic quantities are observed imperfectly or imprecisely, including population size, quality of education, occupation, prices, international trade and income. The importance of active learning through improved monitoring is therefore probably not limited to climate policy.

Acknowledgements An earlier version of this paper was presented at the 22nd annual conference of the European Association of Environmental and Resource Economists (EAERE 2016, Zurich, Switzerland), and at the ECOMOD 2016 conference (Lisbon, Portugal). The authors are grateful to Reyer Gerlagh, Andreas Lange, Ulrike Lehr, Thomas Lontzek, Antony Millner, David Popp, Hans-Peter Weikard, Cees Withagen, and two anonymous reviewers for their valuable comments and suggestions. All remaining errors are the authors'.

Appendix A: The Full Model

The list of variables and parameters are given in Tables 4 and 5.

$$\max_{\mu_t, R_t} \mathbb{E} \sum_{t=0}^{\infty} L_t \beta^t U(C_t, L_t) \quad (\text{A.1})$$

$$C_t = \left(1 - \theta_1 \mu_t^{\theta_2}\right) \Omega_t Q_t - I_t - R_t \quad (\text{A.2})$$

$$K_{R,t+1} = (1 - \delta_R) K_{R,t} + R_t \quad (\text{A.3})$$

$$K_{t+1} = (1 - \delta_k) K_t + I_t \quad (\text{A.4})$$

$$M_{AT,t+1} = (1 - \mu_t) \sigma_t Q_t + E_{LAND,t} + \delta_{AA} M_{AT,t} + \delta_{UA} M_{U,t} \quad (\text{A.5})$$

$$M_{U,t+1} = \delta_{AU} M_{AT,t} + \delta_{UU} M_{U,t} \quad (\text{A.6})$$

Table 4 Variables

U	Utility function	$= (C_t/L_t)^{1-\alpha}/(1-\alpha)$
C_t	Consumption	$= \left(1 - \theta_1 \mu_t^{\theta_2}\right) \Omega_t Q_t - I_t - R_t$
μ_t	Emissions control rate	Control variable
R_t	Investment in climate research	Control variable
K_t	Capital stock	$K_0 = \$137$ trillion
K_{Rt}	Research capital stock	$K_{R0} = \$950$ million
M_{ATt}	Carbon stocks in the atmosphere	$M_{AT0} = 808.9$ GtC
M_{Ut}	Carbon stocks in the upper ocean	$M_{U0} = 18,365$ GtC
M_{Lt}	Carbon stocks in the lower ocean	$M_{L0} = 1255$ GtC
T_{ATt}	Atmospheric temperature deviations	$T_{AT0} = 0.7307$ °C
T_{LOt}	Ocean temperature deviations	$T_{LO0} = 0.0068$ °C
$\bar{f}_t$	Mean of the total feedback factors	$\bar{f}_0 = 0.65$
v_t	Variance of the total feedback factors	$v_0 = 0.13^2$
Ω_t	Damage function	$= 1/(1 + \kappa_1 T_{ATt} + \kappa_2 T_{ATt}^{\kappa_3} + \kappa_4 T_{ATt}^{\kappa_5})$
Q_t	Gross output	$= A_t K_t^\gamma L_t^{1-\gamma}$
I_t	Investment in general	$= s Q_t \Omega_t$
A_t	Total factor productivity	Exogenous
L_t	Labor force	Exogenous
σ_t	Emission-output ratio	Exogenous
RF_{Nt}	Radiative forcing from non-CO ₂ gases	Exogenous
E_{LANDt}	GHG emissions from the sources other than energy consumption	Exogenous
ε_t	Temperature shocks	Stochastic
$v_{\varepsilon t}$	Variance of observed temperature shocks	$v_{\varepsilon 0} = 0.1^2$

The initial values for the state variables and the evolutions of the exogenous variables are from Cai et al. (2012), except for the research capital stock. The initial research capital stock does not affect the main results of this paper unless it is far higher than the default values. The lower bounds of the economic variables such as consumption, the capital stock, and gross world output are set to \$0.001 per person per year in this paper. In addition, there are no upper bounds for temperature increases

$$M_{Lt+1} = \delta_{UL} M_{Ut} + \delta_{LL} M_{Lt} \quad (\text{A.7})$$

$$T_{AT_{t+1}} = T_{AT_t} + \xi_1 \left\{ \eta \ln(M_{AT_t}/M_b) / \ln(2) + RF_{N_t} - \eta T_{AT_t} / \lambda - \xi_3 (T_{AT_t} - T_{LO_t}) \right\} + \varepsilon_{t+1}^{nv} \quad (\text{A.8})$$

$$T_{AT_{t+1}}^{obs} = T_{AT_{t+1}} + \varepsilon_t^{obs} \quad (\text{A.9})$$

$$T_{LO_{t+1}} = T_{LO_t} + \xi_4 \{ T_{AT_t} - T_{LO_t} \} \quad (\text{A.10})$$

$$\bar{f}_{t+1} = \frac{\bar{f}_t + \zeta_1 T_{AT_t}^{obs} H_{t+1} v_t / v_{\varepsilon_t}}{1 + \zeta_1^2 T_{AT_t}^{obs2} v_t / v_{\varepsilon_t}} \quad (\text{A.11})$$

Table 5 Parameters

λ	Equilibrium climate sensitivity	$=\lambda_0/(1-f)$
f	True value of the total feedback factors	0.6
λ_0	Reference climate sensitivity	$1.2\text{ }^{\circ}\text{C}/2\times\text{CO}_2$
s	Savings rate	0.245
α	Elasticity of marginal utility	2
ρ	Pure rate of time preference	0.015
β	Discount factor	$=1/(1+\rho)$
γ	Elasticity of output with respect to capital	0.3
δ_k	Depreciation rate of the capital stock	0.1
δ_R	Depreciation rate of research investment	0.15
$\kappa_1, \kappa_2, \kappa_3, \kappa_4$	Damage function parameters	$\kappa_1=0, \kappa_2=0.0028388, \kappa_3=2, \kappa_4=\kappa_5=0$
θ_1, θ_2	Abatement cost function parameters	$\theta_1=0.0561, \theta_2=2.887$
$\delta_{AA}, \delta_{UA}, \delta_{AU}, \delta_{UU}, \delta_{UL}, \delta_{LL}, \xi_1, \xi_3, \xi_4, \eta$	Climate parameters	$\delta_{AA}=0.9810712, \delta_{UA}=0.0189288, \delta_{AU}=0.0097213, \delta_{UU}=0.005, \delta_{UL}=0.0003119, \delta_{LL}=0.9996881, \xi_1=0.022, \xi_3=0.3, \xi_4=0.005, \eta=3.8$
α_R	Learning parameters	\$3.42 million
σ_{nv}^2	Learning parameters	0.0064
M_b	Pre-industrial carbon stock	596.4 GtC
ω_c	Parameter reflecting satisfaction of the decision maker with the magnitude of learning	10^{-5}
v_{ε_0}	Initial value of the variance of observed temperature shocks	0.10^2

The parameter values for climate parameters are from Cai et al. (2012). The parameter values for λ_0 and δ_R are from Roe and Baker (2007), Hall et al. (2010), respectively. The other parameters are from Nordhaus (1994, 2008) except for the learning parameters

$$v_{t+1} = \frac{v_t}{1 + \zeta_1^2 T_{AT_t}^{obs2} v_t / v_{\varepsilon_t}} \quad (\text{A.12})$$

$$v_{\varepsilon_t} = \alpha_{R_1} / K_{R_t} + \sigma_{nv}^2 \quad (\text{A.13})$$

where $\mathbb{E}$ is the expectation operator given information at point in time t (annual).

Appendix B: Additional Results

Figure 9 is the results for improved observations (the reference case in Sect. 3). Each figure is the average of 1000 Monte Carlo simulations. This figure shows how each variable evolves over time. For instance the rate of emissions control gradually increases during the first 2–3 centuries and then reaches the value 1 (full abatement). The carbon stock gradually decreases after the rate of emissions control becomes equal to 1. Atmospheric temperature follows the same pattern with a time lag. The maximum temperature increases (from 1900) are less than 4 °C (in the early twenty-second century) for all the cases. There is an initial peak in research investment and then the level of research investment becomes trivial. Consumption and gross investment (other than research investment) grows continuously since our model is based on the DICE model, which represents continuous economic growth.

Figure 10 shows the results of all runs.

Figure 11 shows the research capital stock with and without the upper bound.

Fig. 9 Additional results. The units for investment, research investment, the carbon stock, temperature increases, and consumption are \$1000 per person, trillion dollars, GtC, °C, and \$1000 per person, respectively

Fig. 10 Additional results (all simulations) (Top): The mean of the total feedback factors (Upper middle): The variance of the total feedback factors (Lower middle): Temperature increases (relative to 1900) (Bottom): The optimal carbon tax (US\$/tC)

Fig. 11 Research capital stock with and without the upper bound. Because of the computational burden, we present the average of 10 Monte Carlo simulations for the case with the upper bound (see Sect. 2.2). The initial research capital stock (\$950 million) is so small that is close to zero in the figure

Appendix C: Error Analysis

We compare social welfare obtained from our solution method and maximum social welfare obtained from grid-based simulations. Since there are 12 state variables, it is not easy to implement simulations using Chebyshev nodes for each state variable (e.g., 5 nodes for each state variable lead to more than 244 million nodes). Instead, we construct grid points based on control variables. We construct 1001 (equally spaced) grid points for the rate of emission control (from 0 to 1) and 51 (equally spaced) grid points for research investment (from 0.5 times the optimal level to 1.5 times the optimal level) each time period and calculate the corresponding social welfare (51,051 grid points in total each year). Then the maximum social welfare is calculated as in Eq. (C.1).

$$W_g(s_t, \theta_t) = \max_{c_{t,i}} [U(s_t, c_{t,i}, \theta_t) + \beta \mathbb{E}_t W(s_{t+1,i}, \theta_{t+1})] \quad (\text{C.1})$$

where i refers to grid points, W_g is the value function obtained from grid-based simulations, c is the vector of control variables, s is the vector of state variables, θ is the vector of uncertain variables.

Fig. 12 Error analysis (Upper) Grid-based simulation results for the initial year. We do not present the value of z-axis since the welfare function has an arbitrary unit. (Lower) The relative difference of the value functions between the dynamic programming method and grid-based simulations for the initial 100 years

The error is defined as the maximum relative difference between the two levels of social welfare obtained from our solution method and from grid-based simulations.

$$Error = \max_t [1 - W(s_t, \theta_t) / W_g(s_t, \theta_t)] \quad (C.2)$$

where W is the value function obtained from our solution method (Eq. 15).

The error over the first 100 years is 1.6×10^{-4} . Figure 12 shows some results. A wider domain may lead to different accuracy results, but this is not the case here because control variables are bounded in our specification (see Sect. 2.3). We do not find a wider domain (up to the upper bound of the variables) leads to less accuracy.

References

- AATSE (2006) Economics of Australia's sustained ocean observation system, benefits and rationale for public funding. Australian Academy of Technological Sciences and Engineering, Melbourne
- Baehr J, Keller K, Marotzke J (2008) Detecting potential changes in the meridional overturning circulation at 26°N in the Atlantic. *Clim Change* 91:11–27
- Baker MB, Roe GH (2009) The shape of things to come: Why is climate change so predictable? *J Clim* 22:4574–4589
- Baker E, Solak S (2010) Climate change and optimal energy technology R&D policy. *Eur J Oper Res* 213(2):442–454
- Bar-Shalom Y, Tse E (1976) Caution, probing, and the value of information in the control of uncertain systems. *Ann Econ Soc Meas* 5(3):323–337
- Beck GW, Wieland V (2002) Learning and control in a changing economic environment. *J Econ Dyn Control* 26:1359–1377
- Bellman R, Dreyfus SE (1962) Applied dynamic programming. The RAND Corporations, Santa Monica
- Bertocchi G, Spagat M (1993) Learning, experimentation, and monetary policy. *J Monet Econ* 32:169–183
- Brohan P, Kennedy JJ, Harris I, Tett SFB, Jones PD (2006) Uncertainty estimates in regional and global observed temperature changes: a new data set from 1850. *J Geophys Res* 111:D12106
- Cai Y, Judd KL, Lontzek TS (2012) DSICE: a dynamic stochastic integrated model of climate and economy. The center for robust decision making on climate and energy policy working paper no. 12-02
- Costello CJ, Neubert MG, Polasky SA, Solow AR (2010) Bounded uncertainty and climate change economics. *Proc Natl Acad Sci* 107(18):8108–8110
- Cristini L, Lampitt RS, Cardin V, Delory E, Haugan P, O'Neil N, Petihakis G, Ruhl HA (2016) Cost and value of multidisciplinary fixed-point ocean observatories. *Mar Policy* 71:138–146
- Cyert RM, DeGroot MH (1974) Rational expectations and Bayesian analysis. *J Polit Econ* 82:521–536
- De Groot MH (1970) Optimal statistical decisions. McGraw-Hill Inc., New York
- Detrick R, Frye D, Collins J, Gobat J, Grosenbaugh M, Pettitt R, Plueddeman A, der Heydt K, Wooding FB, Orcutt J (2000) DEOS moored buoy observatory design study. Woods Hole Oceanographic Institution technical report
- Douglas-Westwood (2006) Global markets for ocean observation systems. Final OOS report
- Fitzpatrick LG, Kelly DL (2017) Probabilistic stabilization targets. *J Assoc Environ Resour Econ* 4(2):611–657
- Gollier C, Jullien B, Treich N (2000) Scientific progress and irreversibility: an economic interpretation of the precautionary principle. *J Public Econ* 75:229–253
- Hall BH, Mairesse J, Mohnen P (2010) Measuring the returns to R&D. *Handb Econ Innov* 2:1033–1082
- Hansen J, Lacis A, Rind D, Russell G, Stone P, Fung I, Ruedy R, Lerner J (1984) Climate sensitivity: analysis of feedback mechanisms. *Geophys Monogr Ser* 29:130–163
- Horowitz J, Lange A (2014) Cost-benefit analysis under uncertainty: a note on Weitzman's dismal theorem. *Energy Econ* 42:201–203
- Hwang IC (2017) A recursive method for solving a climate-economy model: value function iterations with logarithmic approximations. *Comput Econ* 50(1):95–110
- Hwang IC, Reynès F, Tol RSJ (2013) Climate policy under fat-tailed risk: an application of DICE. *Environ Resour Econ* 56(3):415–436
- Hwang IC, Tol RSJ, Hofkes M (2016) Fat-tailed risk about climate change and climate policy. *Energy Policy* 89:25–35
- Hwang IC, Reynès F, Tol RSJ (2017) The effect of learning on climate policy under fat-tailed risk. *Resour Energy Econ* 48:1–18
- Jensen S, Traeger CP (2013) Optimally climate sensitive policy under uncertainty and learning. Department of Agricultural and Resource Economics, UC Berkeley
- Johnson TC (2007) Optimal learning and new technology bubbles. *J Monet Econ* 54:2486–2511
- Jones PD, Osborn TJ, Briffa KR (1997) Estimating sampling errors in large-scale temperature averages. *J Clim* 10:2548–2568
- Judd KL, Maliar L, Maliar S (2011) Numerically stable and accurate stochastic simulation approaches for solving dynamic economic models. *Quant Economics* 2:173–210
- Karp L, Zhang J (2006) Regulation with anticipated learning about environmental damages. *J Environ Econ Manag* 51:259–279
- Keller K, Deutsch C, Hall MG, Bradford DF (2007a) Early detection of changes in the North Atlantic meridional overturning circulation: implications for the design of ocean observation systems. *J Clim* 20:145–157
- Keller K, Kim SR, Baehr J, Bradford DF, Oppenheimer M (2007b) In: Schlesinger M, Khesghi H, Smith J, De La Chesnaye F, Reilly JM, Wilson T, Kolstad C (eds) What is the economic. Cambridge University Press, Cambridge, pp 343–354

- Kelly DL, Kolstad CD (1999) Bayesian learning, growth, and pollution. *J Econ Dyn Control* 23:491–518
- Kelly DL, Tan Z (2015) Learning and climate feedbacks: optimal climate insurance and fat tails. *J Environ Econ Manag* 72:98–122
- Kendrick D (1982) Caution and probing in a macroeconomic model. *J Econ Dyn Control* 4:149–170
- Kendrick DA (2005) Stochastic control for economic models: past, present and the paths ahead. *J Econ Dyn Control* 29:3–30
- Kennedy JJ, Rayner NA, Smith RO, Parker DE, Saunby M (2011) Reassessing biases and other uncertainties in sea surface temperature observations measured in situ since 1850: 2. Biases and homogenization. *J Geophys Res* 116:D14104
- Kent E, Hall AD, Leader VOSclim Task Team (2000) The voluntary observing ship (VOS) scheme. American Geophysical Union, Washington
- Kent EC, Ball G, Berry DI, Fletcher J, Hall A, North S, Woodruff SD (2010) The voluntary observing ship (VOS) scheme. In: Hall J, Harrison DE, Stammer D (eds) *Proceedings of oceanObs'09: sustained ocean observations and information for society*, vol 2. European Space Agency, pp 551–561
- Knutti R, Stocker TF, Joos F, Plattner G (2002) Constraints on radiative forcing and future climate change from observations and climate model ensembles. *Nature* 46(18):719–723
- Kolstad CD (1996a) Learning and stock effects in environmental regulation: the case of greenhouse gas emissions. *J Environ Econ Manag* 31:1–18
- Kolstad CD (1996b) Fundamental irreversibilities in stock externalities. *J Public Econ* 60:221–233
- Kolstad CD, Ulph A (2008) Learning and international environmental agreements. *Clim Change* 89:125–141
- Kolstad CD, Ulph A (2011) Uncertainty, learning, and heterogeneity in international environmental agreements. *Environ Resour Econ* 50:289–403
- Leach AJ (2007) The climate change learning curve. *J Econ Dyn Control* 31:1728–1752
- Lemoine DM (2010) Climate sensitivity distributions depend on the possibility that models share biases. *J Clim* 23(16):4395–4415
- Lemoine DM, Traeger C (2014) Watch your step: optimal policy in a tipping climate. *AEJ Econ Policy* 6(1):137–166
- Li WC, Hall BH (2016) Depreciation of business R&D capital (No. w22473). National Bureau of Economic Research
- MacRae EC (1972) Linear decision with experimentation. *Ann Econ Soc Meas* 1(4):437–447
- Maliar L, Maliar S (2005) Solving nonlinear stochastic growth models: iterating on value function by simulations. *Econ Lett* 87:135–140
- Manne A, Richels RG (1992) *Buying greenhouse insurance: the economic costs of carbon dioxide emission limits*. The MIT Press, Cambridge
- Marcoul P, Weninger Q (2008) Search and active learning with correlated information: empirical evidence from mid-Atlantic clam fishermen. *J Econ Dyn Control* 32:1921–1948
- Mburu DN (2006) Possibilities for expansion of surface weather observing systems in east Africa. WMO publications IOM-94-TECO2006 1(6)
- Meldrum D, Wallace A, Rolland J, Burnett W, Lumpkin R, Niller P, Viola H, Charpentier E, Fedak M (2010) Data buoy observations: the status quo and anticipated developments over the next decade. In: *Proceedings of OceanObs 9*
- Millner A (2013) On welfare frameworks and catastrophic climate risks. *J Environ Econ Manag* 65:310–325
- Morice CP, Kennedy JJ, Rayner NA, Jones PD (2012) Quantifying uncertainties in global and regional temperature change using an ensemble of observational estimates: the HadCRUT4 data set. *J Geophys Res* 117:D08101
- Nadiri MI, Prucha IR (1996) Estimation of the depreciation rate of physical and R&D capital in the U.S. total manufacturing sector. *Econ Inq* 34(1):43–56
- Nordhaus WD (1994) *Managing the global commons: the economics of climate change*. The MIT Press, Cambridge
- Nordhaus WD (2008) *A question of balance: weighing the options on global warming policies*. Yale University Press, New Haven
- Nordhaus WD (2017) Revisiting the social cost of carbon. *Proc Natl Acad Sci* 114(7):1518–1523
- Nordhaus WD, Popp D (1997) What is the value of scientific knowledge? An application to global warming using the PRICE model. *Energy J* 18:1–46
- North S (2007) Report of the task team on satellite communications costs. SOT-IV. 16–21 April 2007, Geneva, Switzerland
- Peck SC, Teisberg TJ (1993) Global warming uncertainties and value of information: an analysis using CETA. *Resour Energy Econ* 15:71–97
- Prescott EC (1972) The multi-period control problem under uncertainty. *Econometrica* 40(6):1043–1058
- Roe GH (2009) Feedbacks, timescales, and seeing red. *Annu Rev Earth Planet Sci* 37:93–115

- Roe GH, Baker MB (2007) Why is climate sensitivity so unpredictable? *Science* 318(5850):629–632
- Rudik I (2016) Optimal climate policy when damages are unknown. Iowa State University working paper
- Solak S, Baker E, Chen H (2015) Convexity analysis of the dynamic integrated model of climate and the economy (DICE). *Environ Model Assess* 20:443–451
- Stern T, Persson UM (2008) An even sterner review: introducing relative prices into the discounting debate. *Rev Environ Econ Policy* 2:61–76
- Stocker TF, Qin D, Plattner G, Tignor MMB, Allen SK, Boschung J, Nauels A, Xia Y, Bex V, Midgley PM (2013) Summary for policymakers. In: *Climate change 2013: the physical science basis*. Working Group I contribution to the fifth assessment report of the intergovernmental panel on climate change
- Tinbergen J (1954) Centralization and decentralization in economic policy. North-Holland, Amsterdam
- Tol RSJ, De Vos AF (1998) A Bayesian statistical analysis of the enhanced greenhouse effect. *Clim Change* 38(1):87–112
- Ulph A, Ulph D (1997) Global warming, irreversibility and learning. *Econ J* 107:636–650
- Van Wijnbergen S, Willems T (2015) Optimal learning on climate change: why climate skeptics should reduce emissions. *J Environ Econ Manag* 70:17–33
- Webster M (2002) The curious role of learning in climate policy: Should we wait for more data? *Energy J* 23(2):97–119
- Webster M, Jakobovits L, Norton J (2008) Learning about climate change and implications for near-term policy. *Clim Change* 89:67–85
- Weitzman ML (2009) On modeling and interpreting the economics of catastrophic climate change. *Rev Econ Stat* 91(1):1–19
- Weitzman ML (2012) GHG targets as insurance against catastrophic climate damages. *J Public Econ Theory* 14(2):221–244
- Wieland V (2000a) Learning by doing and the value of optimal experimentation. *J Econ Dyn Control* 24:501–534
- Wieland V (2000b) Monetary policy, parameter uncertainty and optimal learning. *J Monet Econ* 46:199–228
- WMO, UNEP (2010) Implementation plan for the global observing system for climate in support of the UNFCCC (2010 update). GOOS-184 GTOS-76 WMO-TD No. 1523
- Yetman J (2003) Probing potential output: monetary policy, credibility, and optimal learning under uncertainty. *J Macroecon* 25:311–330