

Regulation with anticipated learning about environmental damages[☆]

Larry Karp*, Jiangfeng Zhang¹

Department of Agricultural and Resource Economics, University of California, Berkeley, USA

Received 26 August 2004

Available online 4 January 2006

Abstract

A regulator anticipates learning about the relation between environmental stocks and economic damages. For a model with linear-quadratic abatement costs and environmental damages, and a general learning process, we show analytically that anticipated learning decreases the optimal level of abatement *at a given information set*. If learning causes the regulator to eventually decide that damages are higher than previously thought, learning eventually increases abatement. Learning also favors the use of taxes rather than quotas. Using a model that is calibrated to describe the problem of global warming, we show numerically that anticipated learning causes a significant reduction in first period abatement and a small increase in the preference for taxes rather than quotas. Even if the regulator's initial priors about environmental damages are much too optimistic, he is able to learn quickly enough to keep the expected stock trajectory near the optimal trajectory.

© 2005 Elsevier Inc. All rights reserved.

JEL classification: D83; L50

Keywords: Global warming; Anticipated learning; Stock pollution; Damage uncertainty; Taxes versus quotas

1. Introduction

The belief that we will eventually obtain better information about the effects of greenhouse gases is central to the current debate over efforts to reduce carbon emissions. If we were convinced that current uncertainty would persist indefinitely, we could model it like any other form of randomness. The anticipation that we will learn about the relation between greenhouse gases and global warming—together with disagreement about how this anticipated learning should affect current policies—complicates the debate. If we incur large abatement costs now and later learn that global warming is not a serious problem, we will have wasted resources. If we delay cutting emissions and later learn that global warming is a serious problem, we will suffer avoidable damages. Both sides of the debate claim that the prospect of learning supports their

[☆]This paper benefited from comments by Alain de Janvry, Tony Fisher, and Leo Simon, Aart de Zeeuw and two anonymous referees. The usual disclaimer applies.

*Corresponding author. Fax: +1 5106438911.

E-mail addresses: karp@are.berkeley.edu (L. Karp), jzhang@adb.org (J. Zhang).

¹Asian Development Bank. The contents of this paper do not necessarily reflect the views of the Asian Development Bank.

recommendations. This paper contributes to understanding the role of anticipated learning on optimal greenhouse gas policy.

We construct a dynamic model of a stock externality in which a regulator anticipates learning about the stochastic relation between the pollution stock and economic damages. This model nests as special cases the situations where the regulator expects to obtain information zero times, a finite number of times, or infinitely often in the future. Our primary result is that anticipated learning increases the optimal level of emissions, i.e. it reduces abatement, *at a given information set*. Under learning, the information about damages changes. If the regulator learns that damages are higher than previously thought, abatement is eventually higher with learning than without learning (The converse also holds.).

An additional feature of this model is that the regulator and firms have asymmetric information about abatement costs, and the regulator might use either taxes or quotas to control the externality. Our secondary result is that anticipated learning about environmental damages favors the use of taxes.

The paper's third contribution is its generalization of the linear-quadratic control problem. This model is a work-horse in applied economic dynamics, and the discovery that there is still something to be learned about it is noteworthy—and potentially useful for other applications.

We calibrate the model in order to assess the likely magnitude of the effects of learning on the level of abatement and the choice between taxes and quotas. We find that anticipated learning causes a significant reduction in the optimal level of abatement, and causes a small increase in the preference for taxes rather than quotas. Even if the regulator begins with priors that are much too optimistic, learning occurs quickly enough that the expected stock trajectory remains close to the full-information optimal trajectory.

Arrow and Fischer [1] and Henry [9] analyze the effect of learning on optimal decisions with irreversibilities. Epstein [4] provides a more general treatment of this problem; his results have recently been extended by Gollier et al. [6]. Ulph and Ulph [25] use Epstein's results to show that in a two-period model of global warming the effect of learning on first-period emissions is ambiguous in general. Chichilnisky and Heal [2] explain why anticipated learning may lead to greater initial abatement when irreversibilities are important. Heal and Kristrom [8] review the role of uncertainty in climate change.

Much of the existing literature concerning climate change uncertainty assumes that information decreases and eventually resolves uncertainty. Nordhaus and Popp [21] and Peck and Teisberg [22] consider the difference between “act and learn” and “learn and act”. Learning can occur all at once as in Kennedy [15] and Kolstad [16], or more gradually as a function of time as in Kolstad [17]. Fisher and Narain [5] study the effect of irreversibilities in the stock of gasses and of abatement capital, holding fixed the amount of learning. Kelly and Kolstad [14] consider active learning about the relation between greenhouse gas levels and global mean temperature changes; Leach [18] studies a generalization of their model.

Most of these papers rely on two-period analytic models or complex models that require numerical solutions. The numerical models permit a rich description of the environment, but their complexity sometimes makes it difficult to understand the relation between outcomes and specific features of the model. Two-period models take as exogenous the second-period maximand. When learning can occur over many periods and where the stock is persistent, as with global warming, the value of being in a particular state in the next period—the value function—depends on future decisions and on future learning.

We compromise between the two previous approaches by using a model that is linear-quadratic in emissions and stocks, but which has a very general learning component. Since the effect of learning is ambiguous even in two-period models, it will also be ambiguous in a more general dynamic model. The fact that we obtain unambiguous results for the linear-quadratic model does not, of course, mean that learning has the same effect under other functional forms. Nevertheless, the linear-quadratic model is helpful in understanding the general problem. The model is simple enough to produce analytic results in a genuinely dynamic context, i.e. one in which a regulator controls a stock externality and has many opportunities to learn about environmental damages. The generality of the learning component is important, because it allows for the possibility that the regulator discovers that environmental damages are extremely high or negligible. The model is also simple to calibrate and easy to interpret, making it possible to understand the effect of assumptions about parameters.

Several papers, [10,11,20,13], compare taxes and quotas for the control of stock externalities when firms and the regulator have asymmetric information about abatement costs. The main result from Weitzman's [26] static model (where damages are associated with a flow rather than a stock) continues to hold: a flatter

marginal environmental damage curve, or a steeper marginal abatement cost curve favors the use of taxes. These models assume that the regulator knows the parameters of the damage function. We extend these models by including anticipated learning about an uncertain damage parameter.

The intuition for our two analytic results—anticipated learning decreases abatement and favors taxes—is simple, and is likely to apply in more general settings. If the regulator never learns about the unknown damage parameter, it is appropriate to solve the problem by maximizing the expectation of the present discounted stream of utility, using the subjective distribution of the unknown parameter. In the absence of learning, this distribution is constant. With learning (and a feedback control rule) the regulator knows that future decisions will be based on the most recent information. If, for example, the regulator begins to believe that damages are more serious than previously thought, he can reduce future emissions. The ability to adapt makes the bad news about the damage parameter less bad. Similarly, good news is more valuable when the regulator can change his future decisions. Thus, anticipated learning has an effect that is similar to that of a more optimistic subjective distribution about the damage parameter. Consequently, at a given information state the optimal emissions are higher in the current period, relative to the case without learning.

The same kind of logic explains the effect of learning on the comparison between taxes and quotas: anticipated learning is similar to a more optimistic prior on the slope of marginal damages. In this context, greater optimism is equivalent to the belief that marginal damages are flatter. Flatter marginal damages favor the use of taxes rather than quotas, just as in the static and dynamic linear-quadratic models mentioned above.

Section 2 presents the linear-quadratic model of abatement costs and environmental damages with the general model of learning. Section 3 establishes the results described above. Section 4 presents a specific model of learning. Section 5 calibrates the resulting model and assesses the magnitude of the effect of learning. Section 6 discusses how our qualitative results might change under different assumptions, and relates our results to the previous literature. Section 7 concludes.

2. The model

We first specify the abatement cost and then environmental damages. Our two analytic results—the effect of anticipated learning on abatement and on the ranking of taxes and quotas—require a model with two different types of uncertainty: about damages and about abatement costs. Anticipated learning distinguishes our model from the models in previous papers that compare taxes and quotas for a stock pollutant.

2.1. Abatement costs

Abatement equals the difference between the actual level of emissions and the Business as Usual (BAU) level. We assume that the abatement costs are quadratic in abatement, so the benefit of emissions is a quadratic function of emissions. We also assume that the intercept of the marginal benefit function equals a constant a plus a mean-zero random variable θ_t with a constant and known variance σ_θ^2 . The slope of marginal benefits is a known constant b . In period t the firm, but not the regulator, knows the value of θ_t . The benefit function in period t is

$$\tilde{f} + (a + \theta_t)x_t - \frac{b}{2}x_t^2. \quad (1)$$

When the regulator sets a tax p_t per unit of emissions, the firm maximizes the benefit of emissions minus the cost of tax. Its problem is

$$\max_x \tilde{f} + (a + \theta_t)x_t - \frac{b}{2}x_t^2 - p_t x_t.$$

The first order condition to this problem implies that the level of emissions is

$$x_t^* = \frac{a - p_t}{b} + \frac{\theta_t}{b} \equiv z_t + \frac{\theta_t}{b}. \quad (2)$$

Hereafter, we assume that the tax-setting regulator chooses z_t , the expected level of emissions under a tax. Substituting x_t^* into the firm's benefit function (1) and taking expectations, gives the expected benefit of

emissions under the tax policy z_t :

$$\tilde{f} + az_t + \frac{\sigma_\theta^2}{2b} - \frac{b}{2} z_t^2. \quad (3)$$

The quota-setting regulator chooses x_t , which by assumption is binding with probability 1. Thus, the expected benefit of emissions under the quota policy x_t is simply

$$\tilde{f} + ax_t - \frac{b}{2} x_t^2. \quad (4)$$

The tax-setting regulator determines only the expected level of emissions, whereas the quota-setting regulator chooses emissions.

2.2. Environmental damages and learning

Let S_t be the stock of pollutants, and x_t be the flow of emissions in period t . All time dependent variables are constant within a period. The fraction $0 \leq \Delta \leq 1$ of the pollutant stock lasts into the next period, so the growth equation for S_t is

$$S_{t+1} = \Delta S_t + x_t. \quad (5)$$

With taxes, the flow of emissions and thus the next period pollutant stock, S_{t+1} , is stochastic since it depends on the cost shock. With quotas, the regulator is able to exactly determine the change in pollution stock.

The environmental damage in period t is

$$D(S_t, \omega_t; G^*) = \frac{G^*}{2} (S_t - \bar{S})^2 \omega_t, \quad (6)$$

where ω_t is an i.i.d. non-negative random variable with mean 1, and G^* is the true but unknown non-negative value of the damage parameter. \bar{S} is a known non-negative constant at which environmental damages are minimized. The presence of the damage shock (ω_t) means that the regulator might not learn the true value G^* in finite time.

The functional form of damages implies that the regulator is not able to influence the amount of learning by manipulating the level of stocks. That is, learning is passive rather than active in this model. To see this, use Eq. (6) to write the “data” (or signal) at time t as

$$\text{data}_t \equiv \frac{2D_t}{(S_t - \bar{S})^2} = G^* \omega_t. \quad (7)$$

By Bayes’s theorem, the posterior on G^* , $\Pr(G^* | \text{data}_t)$, is proportional to the product of the likelihood function, $\Pr(\text{data}_t | G^*)$, and the prior, $\Pr(G^*)$. The numerical value of the data at t depends on G^* and ω_t , but not on S_t (A change in S_t causes an offsetting change in D_t , leaving unchanged the middle expression in Eq. (7)). Therefore, $\Pr(\text{data}_t | G^*)$ is independent of S_t ; consequently, the posterior $\Pr(G^* | \text{data}_t)$ is independent of S_t . In other words, changing S_t does not change the information (about G^*) that the regulator obtains from observing $G^* \omega_t$.² The regulator is not able to affect the amount of learning by manipulating the pollutant stock.

The value of G^* might be much higher than the regulator’s expectation of this variable, so damages could be much higher than currently believed. To this extent, the model captures the uncertainty about global warming. However, the model makes a number of assumptions that may not hold for global warming, as we discuss in Section 6.

At time t the regulator’s subjective expectation of the value of G^* is $G_t = E_t G^*$; the operator E_t denotes the expectation conditional on information available at time t . In any period, the expectation of the single-period payoff is linear in G_t . This linearity implies that if G_t were a constant \bar{G} (i.e. there is parameter uncertainty but

²Consider the alternative damage function $\frac{G^*}{2}(S_t - \bar{S})^2 + \omega_t$, where ω appears additively rather than multiplicatively. In that case, the data at time t is (D_t, S_t) ; a larger value of $(S_t - \bar{S})^2$ causes G^* to explain a greater proportion of the variation in damages. For this additive model, there exists the possibility of active learning.

no learning), we could solve the control problem by replacing G^* with \bar{G} and simply ignore the uncertainty regarding G^* . \bar{G} is the certainty equivalent value of G^* in the model where G_t is constant. However, parameter uncertainty together with anticipated learning leads to a non-trivial change in the optimization problem.³

In order to be able to use standard dynamic programming methods, we need to be able to describe the subjective distribution of G^* using a finite number of parameters. Those parameters are elements of the state vector. In our model the subjective distribution of G^* at time t is defined by two moments, the mean and variance, $\chi_t \equiv (G_t, \sigma_{G,t}^2)$. However, our proofs do not depend on whether $\sigma_{G,t}^2$ is a vector of higher moments or a scalar (the variance). The regulator cannot predict his future subjective expectation, so his current subjective expectation is an unbiased estimator of its future value, i.e. $E_t G_{t+\tau} = G_t$ for $\tau \geq 0$.

For the purpose of nesting special cases in a more general model, we use the non-negative integer n to denote the number of periods during which the regulator expects to obtain information about the damage parameter. In order to avoid uninteresting special cases, we assume that if $n \geq 1$, learning begins in the current period and continues for n consecutive periods.

The variance of the subjective distribution (or more generally, the higher moments) changes stochastically. On average, we expect that learning decreases the subjective variance. However, if the regulator receives a surprising piece of information, he may decide that he is less certain about the unknown parameter than he previously thought. In that case, σ_G^2 increases. We assume that if the variance ever falls to 0, i.e., if the regulator ever becomes certain of the value of G^* , it does not subsequently increase. In addition, we assume that the variance approaches 0 only asymptotically, if at all. The last assumption (adopted only to simplify the notation) means that the control problem for very large but finite n and for $n = \infty$ are not exactly the same; they can, of course, be very similar.

If $n = 0$ the regulator expects never to obtain information about G^* ; in that case the regulator solves a standard control problem without anticipated learning. For given G , the two control problems with $n = 0$ and arbitrary σ_G^2 , or with $\sigma_G^2 = 0$ and arbitrary n , are equivalent. In these two cases the regulator never changes his subjective mean of the damage parameter, either because he never acquires new information ($n = 0$) or because he is convinced that he already knows the truth ($\sigma_G^2 = 0$).

This model of learning is quite general. It allows for the possibility that the regulator learns about G^* quickly or slowly. The regulator may discover that it is likely that damages are extremely high, and very sensitive to changes in the stock, and his subjective variance might either increase or decrease.

2.3. The optimization problem

The model has three types of state variables, the stock S , the moments of the subjective distribution $\chi_t \equiv (G, \sigma_G^2)$, and the number of periods of future learning, n . The state χ (and under taxes, the state S) changes stochastically and the state n is deterministic.

The expected payoff in a period is equal to the expected benefits of emissions minus the expected damages. The expectation is taken with respect to the cost shock, θ , the damage shock ω and the unknown parameter G^* . Under taxes, where emissions are given by Eq. (2), the expected single period payoff is

$$f_t + az_t - \frac{bz_t^2}{2} + \frac{\sigma_\theta^2}{2b} - c_t S_t - \frac{G_t}{2} S_t^2 \tag{8}$$

with $f_t \equiv \tilde{f} - \frac{G_t}{2} \bar{S}^2$, $c_t \equiv -G_t \bar{S}$.

We discuss the control problem in which the regulator uses taxes. We can obtain the solution under quotas immediately from the solution under taxes, simply by replacing z with x and setting $\sigma_\theta^2 = 0$. (Compare the expressions (3) and (4).)

The parameter f_t affects the value of the payoff but it does not interact with either the stock or the control, so it has no effect on the optimal policy or on any of our results. Therefore, to simplify notation we replace f_t

³Although the single period payoff is linear in the subjective expectation of G^* , the value function is non-linear in this parameter. This fact means that anticipated learning about G^* affects the optimal program; in contrast, parameter uncertainty in the absence of learning does not change the optimization problem. Even with a more general specification of the payoff, parameter uncertainty in the absence of learning is not of any particular interest. For example, if the single period payoff is $h(G^*, S_t, z_t, \theta_t)$ and there is no learning, we can solve the control problem as if there were no parameter uncertainty, replacing the single period payoff with $H(S_t, z_t, \theta_t) \equiv E_{G^*} h(G^*, S_t, z_t, \theta_t)$.

with a constant, f . Our formulation of damages, Eq. (6), implies that the intercept of marginal damages is the unknown constant $-G^*\bar{S}$, so the subjective expectation of the intercept is $c_t = -G_t\bar{S}$. The critical feature in our model is that the slope of marginal damages is uncertain. The uncertainty about the intercept is an incidental feature. In order to establish this point, we also consider an alternative model of damages, in which the intercept is a *known* constant, c . With this alternative, we need a restriction on the magnitude of the constant; that restriction is automatically satisfied when the slope is $c_t = -G_t\bar{S}$. We state our results for both the cases where the intercept of marginal damages is a constant and where it equals $-G_t\bar{S}$. To simplify notation we drop the time subscript on c_t unless it is needed for emphasis.

With a discount factor β , the tax-setting regulator’s maximized expected payoff at time t is

$$J(S_t, G_t, \sigma_{G,t}^2, n_t) = \max E_t \sum_{\tau=0}^{\infty} \beta^\tau \left\{ f + az_{t+\tau} - \frac{bz_{t+\tau}^2}{2} + \frac{\sigma_\theta^2}{2b} - cS_{t+\tau} - \frac{G^*}{2} S_{t+\tau}^2 \right\}. \tag{9}$$

The dynamic programming equation is

$$J(S_t, \chi_t, n_t) = \max_z f + az - \frac{bz^2}{2} + \frac{\sigma_\theta^2}{2b} - cS_t - \frac{G_t}{2} S_t^2 + \beta E_{\chi_{t+1}}(E_{\theta_t} J(S_{t+1}, \chi_{t+1}, n_{t+1}))$$

subject to $S_{t+1} = \Delta S_t + z + \frac{\theta_t}{b}$, $n_{t+1} = \max(n_t - 1, 0)$. (10)

To evaluate the continuation payoff we take expectations with respect to S_{t+1} and χ_{t+1} , the stochastic states. The maximization in problem (10) is subject to the equations of motion for χ , the moments of the subjective distribution. Since the analytic results do not depend on these equations, we do not specify them at this time.

There are at least three ways that we can think about increasing learning in this model: (i) An increase in n , the number of times that new information will arrive; (ii) An increase in σ_G^2 , the measure of uncertainty about the unknown parameter—if σ_G^2 is close to zero, there is little scope for learning, and if σ_G^2 is large, potential learning is also large; and (iii) An increase in the precision of future information (Section 4 formalizes the meaning of this third possibility.). The first two changes alter an argument in the value function, and the third change alters the equation of motion for σ_G^2 , thereby altering the value function itself.

3. The effect of learning on emissions and policy ranking

We begin by examining the case where $n = 1$ in order to show the relation between our model and previous models in which learning occurs only once. The next subsection states our two major results: anticipated learning increases emissions in the current period, and it favors the use of taxes rather than quotas. Most proofs are in the Appendix.

3.1. One-time learning ($n = 1$)

For any n , the first order condition to problem (10),

$$a - bz = -\beta E_{\chi_{t+1}}(E_{\theta_t} J_S(S_{t+1}, \chi_{t+1}, n_{t+1})), \tag{11}$$

states that the expected marginal benefit of emissions in the current period should equal the discounted expectation of pollution’s shadow cost (defined as the negative shadow value, i.e. $-J_S$).

Denote the value of G_t when there is no anticipated learning ($n = 0$) as G^0 . In the absence of learning, this value does not change, so we do not use a time subscript. Because G^* enters each period’s payoff linearly, the value function depends on G^0 but not on the higher moment(s) $\sigma_{G,t}^2$ when $n = 0$. Consequently, when $n = 0$ the value function can be written as $J(S_t, G_t, \sigma_{G,t}^2, 0) \equiv \tilde{J}(S_t; G^0)$.

We can apply the logic used in previous models where learning occurs only once, e.g. Ulph and Ulph [25], to compare emissions when $n = 1$ and $n = 0$. In both of these cases, $n = 0$ in the *next* period. Denote the value of G_t in the current period as G , so for $n = 0$, $G^0 = G$; for $n = 1$, $E_t G_{t+1} \equiv E_t G^0 = G$.

In the two cases where $n = 0$ or $n = 1$ Eq. (11) specializes to

$$\begin{aligned} n = 0 : a - bz &= -\beta(E_{\theta_t} \tilde{J}_S(S_{t+1}; G^0)) \\ n = 1 : a - bz &= -\beta E_{G^0|G}(E_{\theta_t} \tilde{J}_S(S_{t+1}; G^0)). \end{aligned}$$

The optimal z is larger under learning ($n = 1$) relative to no-learning ($n = 0$) if and only if the function $E_{\theta_t} \tilde{J}_S(S_{t+1}, G^0)$ is convex in G^0 .⁴ If this function is convex, then

$$E_{G^0}(E_{\theta_t} \tilde{J}_S(S_{t+1}, G^0)) > (E_{\theta_t} \tilde{J}_S(S_{t+1}, E_t G^0))$$

by Jensen’s inequality. When moving from no-learning to learning, z must increase in order to maintain the equality in the first order (11).

For the linear-quadratic specification we have an explicit expression for $\tilde{J}_S(S_t; G^0)$. It is easy to confirm that this function is convex in G^0 , so the optimal level of emissions is higher when $n = 1$ compared to $n = 0$. This fact provides the starting point for an inductive proof that establishes that an increase in n increases the level of emissions.

3.2. Statement of results

We begin with the following simple but useful result.

Lemma 1. *For both taxes and quotas, and for any integer $n \geq 0$, the value function is quadratic in S ; that is, the value function has the form $J(S, G, \sigma_G^2, n) = \lambda_n + \mu_n S + \frac{\rho_n}{2} S^2$, where λ_n, μ_n and ρ_n are functions of (G, σ_G^2, n) .*

This lemma is important for our analytic results, and it is also useful for numerical work. It enables us to express the solution to the optimization problem as an explicit functional of λ_n, μ_n and ρ_n . We can obtain those three functions by solving a recursive system of functional equations. This system *does not involve optimization*, a fact that greatly simplifies the numerical solution to the system. The proof of Lemma 1 presents this system.

As a consequence of this lemma we have

Proposition 1. *The Principal of Certainty Equivalence with respect to the cost shock θ holds for any integer $n \geq 0$. Consequently, the expected level of emissions under the optimal tax equals the optimal quota.*

Proof. Both statements follow from inspection of the control rules, given in the proof of Lemma 1. These control rules are independent of the variance of the cost shock, and they are identical for taxes and quotas. □

This fact has been previously noted in a model that does not involve learning about damages [11].

The following lemma identifies a restriction on parameter values needed to insure that when the stock is sufficiently close to 0, the optimal level of emissions is positive. This restriction makes the problem economically meaningful. The benefit of emitting in the current period must be great enough to induce the regulator to allow positive emissions, at least when the stock is low. Here we explicitly consider both the case where the intercept of marginal damages is a constant c , and where the intercept equals $c_t = -G_t \bar{S}$.

Lemma 2. *When c is a known constant, in the absence of learning the optimal level of expected emissions is positive for $S = 0$ iff*

$$c < \frac{a(1 - \beta\Delta)}{\beta}. \tag{12}$$

When $c_t = -G_t \bar{S}$ with $G^* \geq 0$ (so that $G_t \geq 0$) and $\bar{S} \geq 0$ the optimal level of expected emissions is positive for $S = 0$.

⁴The validity of this assertion depends on the fact that the single period payoff is linear in G^* . If the single period payoff were non-linear in G^* the comparison between learning and no learning would depend on the convexity $E_{\theta_t} J_S$ with respect to the distribution of G^* , as in Epstein [4].

The following lemma provides the basis for understanding the effect of learning on the optimal level of emissions, and on the comparison of taxes and quotas. The lemma uses the functions $\rho_n(\chi)$ and $\mu_n(\chi)$ introduced in Lemma 1.

Lemma 3. *When the initial variance is $\sigma_G^2 > 0$: (a) The function $\rho_n(\chi)$ is increasing in n ; (b) The function $\mu_n(\chi)$ is increasing in n if $c_t = -G_t \bar{S} < 0$ or if c is a constant and inequality (12) holds.*

The geometric intuition for this lemma is straightforward. Additional opportunities to learn must increase the payoff (provided that $\sigma_G^2 > 0$). Thus, for $n \geq 1$

$$J(S, \chi, n) - J(S, \chi, n-1) = (\lambda_n - \lambda_{n-1}) + (\mu_n - \mu_{n-1})S + \frac{1}{2}(\rho_n - \rho_{n-1})S^2 > 0. \quad (13)$$

This inequality must hold for all S , so it must be the case that $\rho_n - \rho_{n-1} > 0$ (It must also be true that $\lambda_n - \lambda_{n-1} > 0$, but this inequality does not influence our two major results, which depend only on the shadow cost of pollution, equal to $-\mu_n - \rho_n S$). There is no reason to suppose—in a general linear-quadratic control problem—that $\mu_n - \mu_{n-1} > 0$. However, that inequality does hold for parameter values that lead to a positive level of emissions at small stock levels.

Our primary result is

Proposition 2. *Suppose that the initial $\sigma_G^2 > 0$. (i) An increase in the opportunities for learning (an increase in n) always increases current emissions provided that S is sufficiently large. (ii) An increase in the opportunities for learning increases current emissions for all $S \geq 0$ if c is a constant and inequality (12) holds or if $c_t = -G_t \bar{S} < 0$.*

Proof. By Lemma 1, we can write the first order condition given by Eq. (11) as

$$a - bz = -\beta E_\chi[\mu_{n-1}(\chi) + \rho_{n-1}(\chi)S]. \quad (14)$$

By Lemma 3(a), the right side is a decreasing function of n for large S . By Lemma 3(a) and (b) it is a decreasing function of n for all $S \geq 0$ if c is a constant and inequality (12) holds, or if $c_t = -G_t \bar{S} < 0$. Under these conditions, an increase in n requires an increase in current emissions in order to retain equality between the marginal utility of current emissions and the shadow cost of the stock of pollution. \square

Proposition 2 shows (under the stated conditions) that anticipated learning increases emissions in our linear-quadratic setting. We know from earlier work (especially [25,6]) that in some settings increased learning has an ambiguous effect on emissions; therefore, anticipated learning might have an ambiguous effect in a more general dynamic model.

Despite this lack of generality, our result describes a plausible effect of learning. As is evident from the first order condition, anticipated learning increases emissions if and only if it decreases the expectation of discounted shadow costs. A sufficient condition for that decrease is for learning (higher n) to decrease the shadow cost of pollution at any state (i.e., any (χ, S)). The shadow cost $(-\mu_n - \rho_n S)$ equals the amount that the regulator would pay for a marginal decrease in the stock of pollution. It is “reasonable” for anticipated learning to reduce this shadow cost, because policies can be adjusted to accommodate new information. Here, the anticipation of learning reduces not only the cost of the stock (i.e., it increases the value of the program, as inequality (13) states), but it also reduces the marginal cost of the stock.

Proposition 2 describes the effect of anticipated learning at a given information set, i.e. for initial beliefs. Of course, an important effect of learning is that it changes those beliefs. If learning eventually eliminates uncertainty, i.e. if $\sigma_{G,t}^2 \rightarrow 0$ and $G_t \rightarrow G^*$ as $t \rightarrow \infty$, the subjective distribution collapses to the true parameter value. In this case, the regulatory program approaches the abatement rule under full information (with respect to G^*). The control rule under full information and under no-learning (Eq. (24) in the Appendix) implies that an increase in G^* or G^0 decreases emissions. This observation implies the following:

Remark 1. If $\sigma_{G,t}^2 \rightarrow 0$ and $G_t \rightarrow G^*$ as $t \rightarrow \infty$, anticipated learning eventually increases abatement and reduces the stock trajectory (relative to no-learning) if and only if $G_1 < G^*$.

If the regulator initially underestimates damages ($G_1 < G^*$) but is able to learn the true relation between stocks and damages, learning eventually increases abatement.

Using superscripts T and Q to denote the value functions under taxes and quotas, we state our second major result (This result holds regardless of whether the intercept of the marginal damage function equals $-G^*\bar{S}$ or a known constant; in the latter case, it does not matter whether inequality (12) is satisfied.).

Proposition 3. For $\sigma_G^2 > 0$, $J^T(S, \chi, n) - J^Q(S, \chi, n)$ is an increasing function of n : Increased opportunities for learning favor the use of taxes rather than quotas.

In the static linear-quadratic problem (where damages are associated with the flow rather than the stock of pollution), taxes are preferred to quotas if and only if the slope of marginal abatement cost exceeds the expected slope of marginal damages [26]. In the dynamic version of this problem (i.e., where damages are caused by the stock) without learning, taxes are preferred to quotas if and only if the slope of marginal abatement costs exceeds the discounted slope of the shadow cost of the stock, $-\beta\rho_0$ [11]. The function $-\beta\rho_0$ is increasing in the expected value of G^* . The intuition for policy ranking in the static and dynamic problems (without learning about the damage parameter) is essentially the same.⁵ Since ρ_0 is convex in $G = EG^*$ (as the proof of Lemma 3 shows), one-time learning ($n = 1$) decreases $E(-\beta\rho_0)$. Learning thus has an effect on the policy ranking that is comparable to a decrease in G , so learning favors the use of taxes. Increased opportunities to learn (a higher value of ρ_n corresponding to a larger value of n) reinforce this effect, further favoring the use of taxes.

4. The log-normal learning model

In order to calibrate a model for greenhouse gasses, we need an explicit learning rule. We assume that the distribution of the damage shock in Eq. (6) is lognormal:

$$\omega_t \sim i.i.d. \text{ lognormal} \left(-\frac{\sigma_\omega^2}{2}, \sigma_\omega^2 \right). \tag{15}$$

We express the subjective moments in terms of $g \equiv \ln G$. The regulator begins in period t with normal priors on $g^* = \ln G^*$, with mean g_t and variance $\sigma_{g,t}^2$:

$$g^* \sim N(g_t, \sigma_{g,t}^2). \tag{16}$$

Given distribution (16), the subjective distribution of G^* is log-normal with

$$E_t G^* \equiv G_t = \exp \left(g_t + \frac{1}{2} \sigma_{g,t}^2 \right), \sigma_{G,t}^2 \equiv \text{var}_t(G^*) = \exp(2g_t + \sigma_{g,t}^2) (\exp(\sigma_{g,t}^2) - 1). \tag{17}$$

Since damages are a product of independent log-normally distributed variables, the regulator has log-normal priors on damages. After observing damages and the current stock, the Bayesian regulator updates his belief about g^* . The moment estimator of g^* , denoted \hat{g}_t , is

$$\hat{g}_t = \ln \frac{2D_t}{(S_t - \bar{S})^2} + \frac{\sigma_\omega^2}{2} \tag{18}$$

with variance $\sigma_{\hat{g}}^2 = \sigma_\omega^2$. The posterior for g^* is normally distributed with the posterior mean g_{t+1} and posterior variance $\sigma_{g,t+1}^2$:

$$g_{t+1} = \frac{\sigma_\omega^2}{\sigma_\omega^2 + \sigma_{g,t}^2} g_t + \frac{\sigma_{g,t}^2}{\sigma_\omega^2 + \sigma_{g,t}^2} \hat{g}_t, \tag{19}$$

$$\sigma_{g,t+1}^2 = \frac{\sigma_{g,t}^2 \sigma_\omega^2}{\sigma_\omega^2 + \sigma_{g,t}^2} \Rightarrow \sigma_{g,t}^2 = \frac{\sigma_{g,0}^2 \sigma_\omega^2}{\sigma_\omega^2 + t \sigma_{g,0}^2}, \tag{20}$$

where $\sigma_{g,0}^2$ is the prior at the beginning of the initial period, $t = 0$ [7, pp. 407–410].

⁵With quotas the regulator chooses emissions exactly, and with taxes the regulator chooses the expected value of emissions. Since the damage function is convex, expected damages are higher when emissions (in the static problem) or the stock (in the dynamic problem) are random variables—as they are under taxes. A larger value of G increases the convexity of damages and therefore favors quotas.

A smaller value of σ_ω^2 is equivalent to greater precision of future information. Using (19), greater precision of information implies that this period's posterior mean, g_{t+1} , is more responsive to information obtained in the current period. Using (20), greater precision of information means that the posterior variance decreases over time more rapidly. Thus, greater precision of information increases the amount of learning, as stated in Section 2.3.

The subjective distribution for the unknown damage parameter G^* collapses to the true value of this parameter as the number of observations approaches infinity. Appendix B1, available through *JEEM*'s online archive for supplementary material at <http://www.aere.org/journal/index.html>, proves this result.

If the regulator begins with too optimistic a prior ($g_0 < g^*$) g_t increases over time, on average. This increase can be enough to offset the decrease in $\sigma_{g,t}^2$, leading to an increase in $\text{var}_t(G_t)$ (using (17)). In this case, during a phase of the learning process the regulator becomes less certain about the value of G^* , although he eventually learns the correct value with probability 1. It is also straightforward to show that the regulator's current expectation of G^* is an unbiased estimate of the future expectation: $E_t G_{t+\tau} = G_t, \forall \tau \geq 0$.

In the absence of anticipated learning, the regulator solves the control problem treating G_t as a constant. In this case the constant $\bar{G} \equiv G_t = \exp(g_t + \frac{\sigma_{g,t}^2}{2})$ is the certainty equivalent value of G^* .

5. Quantitative results

We calibrate the model to describe the problem of controlling CO_2 emissions in order to limit the possible damages caused by global warming. Most global warming models contain a more complex relation between greenhouse gas stocks and environmental damages. In some respects these models reflect more accurately the current state of art of the physical sciences.

This model is much simpler. It is easy to discover how assumptions about the likely consequences of increased carbon stocks and about abatement costs determine the optimal level and method of abatement, and to explore the role of learning. Our model is consistent with the more complex models, because our calibration uses much of the same data and opinions. The numerical results provide an indication of the quantitative effect of learning on both the optimal level of abatement and on the choice between taxes and quotas.

5.1. Calibration of a global warming model

Most readers would find it difficult to decide whether a particular value of g (or G) should be considered large or small. Therefore, we describe our calibration in terms of the parameter ϕ , defined as the expected percentage reduction of Gross World Product (GWP) due to a doubling of stocks from their pre-industrial level. The parameters ϕ and G are linearly related, as described in the online Appendix B2. The values $\phi = 0.3$, $\phi = 1.33$ and $\phi = 3.6$ represent low, moderate, and high estimates of damages. Table 1 contains the

Table 1
Base-line parameters

Parameter	Note	Value
σ_ω^2	Variance of ln(damage shock)	0.6349
\bar{S}	Zero damage stock, billion tons of carbon	590
a	Intercept of the marginal benefit, \$(/ton of carbon)	224.26
\bar{x}	BAU decade emissions, billion tons of carbon	116.73
b	Slope of the marginal benefit, billion \$(/billion tons of carbon) ²	1.9212
σ_θ	Standard deviation of cost shock, \$(/ton of carbon)	5.5945
λ	An annual decay rate of 0.0083	0.9204
β	A continuous yearly discount rate of 3%	0.7408

baseline parameter values. The online Appendix B2 explains how we obtain these values, and the relation between our calibration and previous models.

5.2. Numerical results

We present four sets of simulations in order to assess the magnitude of the effect of learning on abatement and on the comparison between taxes and quotas. We assume that learning continues indefinitely ($n = \infty$). We obtain the control rule in this case by numerically solving the recursive fixed point equations (27) and (28) in the Appendix.

Fig. 1 shows the optimal abatement in the first period, when $S = 781$ billions tons of carbon, equivalent to the current atmospheric CO_2 concentration. This abatement is expressed as a percentage of the BAU level of emissions, using three different values of ϕ (defined as the regulator's initial point estimate of the annual percentage loss in GWP due to a doubling of carbon stocks). As noted above, the values $\phi = 0.3$, $\phi = 1.33$ and $\phi = 3.6$ represent low, moderate, and high estimates of damages. In performing this simulation we change $\sigma_{g,1}^2$ and make offsetting changes in g_1 so that $G_1 = E_1 G^* = \exp(g_1 + \frac{1}{2}\sigma_{g,1}^2)$ remains constant. As we hold G_1 fixed and increase $\sigma_{g,1}^2$, the initial expectation of damages remains constant but the amount of uncertainty increases. Consequently, the potential for learning increases. We show the results as $\sigma_{g,1}^2$ varies from the minimum level, 0, to the level in our calibration, 0.63. As $\sigma_{g,1}^2$ varies over this range, the coefficient of variation of damages varies from 0.94 to 1.6.

As we previously noted, in the absence of anticipated learning, the optimal decision depends on the certainty equivalent parameter G_1 , but it does not depend on the amount of uncertainty about the parameter G^* . Therefore, the dotted lines labeled “without learning” are constant with respect to $\sigma_{g,1}^2$. In the absence of learning, optimal abatement is sensitive to the estimate of damages; a 170% increase in the estimate of damages, from $\phi = 1.33$ to $\phi = 3.6$, results in a 138% increase in abatement.

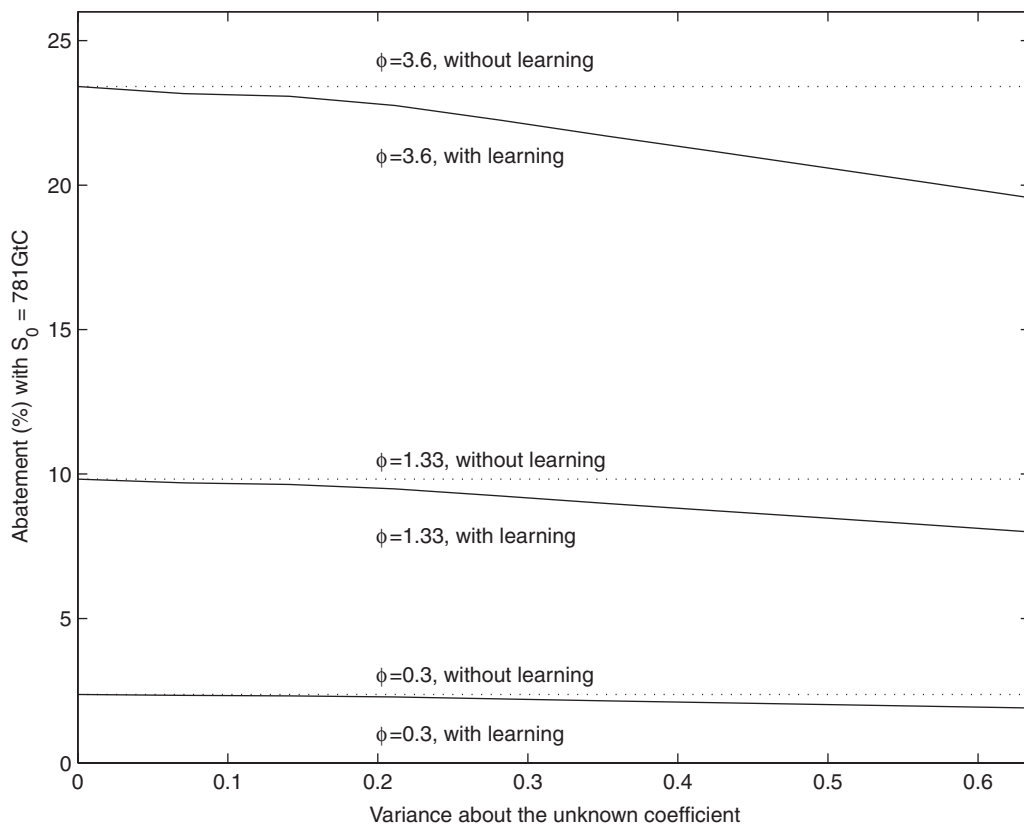


Fig. 1. Abatement with learning as a function of initial uncertainty.

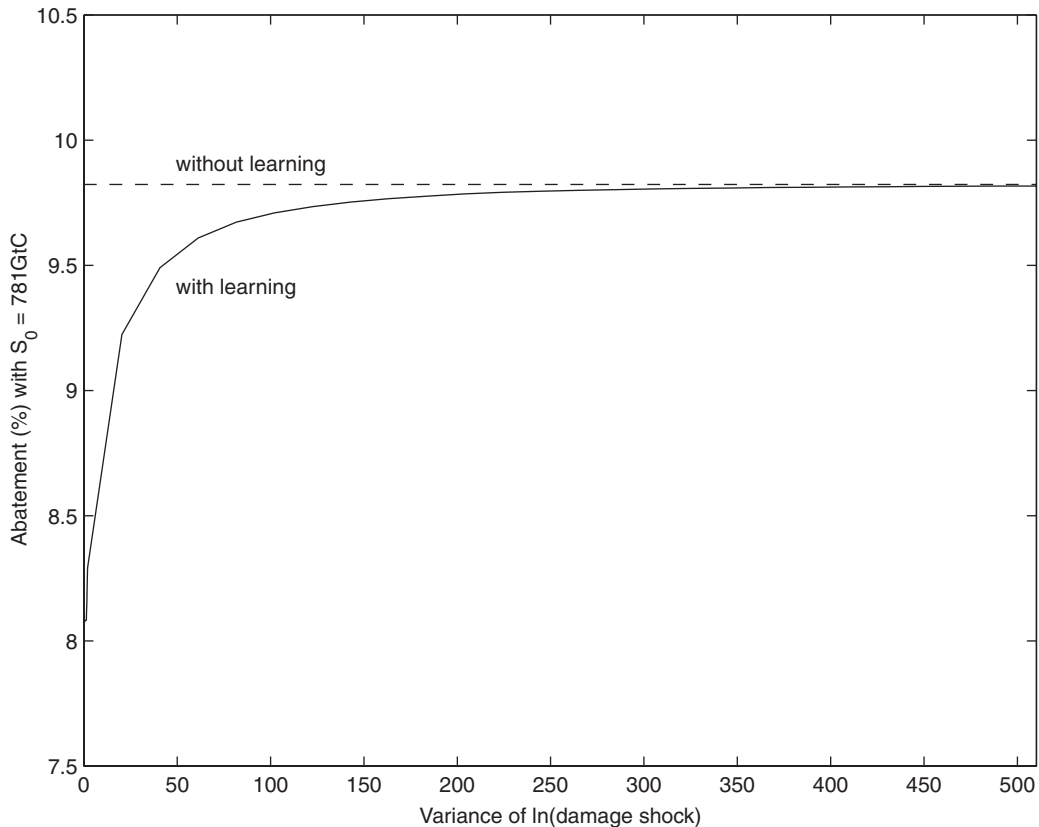


Fig. 2. Abatement with learning as a function of the variance of the signal ($\phi = 1.33$, $\sigma_g^2 = 0.63$).

The potential for learning increases with the amount of parameter uncertainty. Not surprisingly, the difference between the optimal level of abatement with and without learning also increases with this uncertainty. When ϕ takes the values 0.3, 1.33 and 3.6, (fixing $\sigma_g^2 = 0.63$) the potential for learning decreases the level of abatement by 16%, 19% and 20%, respectively.

The second experiment studies the effect of learning as a function of the variance of the damage shock. As this variance increases, the signal becomes less informative. Consequently, learning occurs more slowly, so the anticipation of learning has a smaller effect on the optimal decision. Our base-line calibration assumes that the damage shock and the parameter uncertainty contribute equally to the overall level of uncertainty about damages: $\sigma_{g,0}^2 = \sigma_{\omega}^2 = 0.63$.

Fig. 2 shows how the variance of the damage shock affects the optimal level of abatement. For very large variances (e.g. $\sigma_{\omega}^2 > 400$), learning occurs so slowly that it is virtually worthless, and there is a negligible difference between the optimal first period policy with and without learning. However, even if the variance of the damage shock is substantially larger than in our calibration, the effect of anticipated learning remains significant.

The third experiment investigates the importance of anticipated learning in the ranking of taxes and quotas. For our calibration, the expected payoff under taxes is approximately 30 billion dollars larger than the expected payoff under quotas, for $1.3 \leq \phi \leq 3.6$ (Fig. 3). The difference decreases with ϕ , in line with previous analytic results [11].⁶ For $\phi = 3.6$ and $\sigma_{g,0}^2 = 0.63$, anticipated learning increases the difference in payoffs

⁶The difference in payoffs under taxes and quotas is proportional to the variance of the cost shock, a parameter about which we have little information. Nevertheless, \$30 billion is a fairly small amount, since it is the difference in value functions. With our decade discount factor of $\beta = 0.74$, \$30 billion is equivalent to a flow of approximately \$7.8 billion per decade.

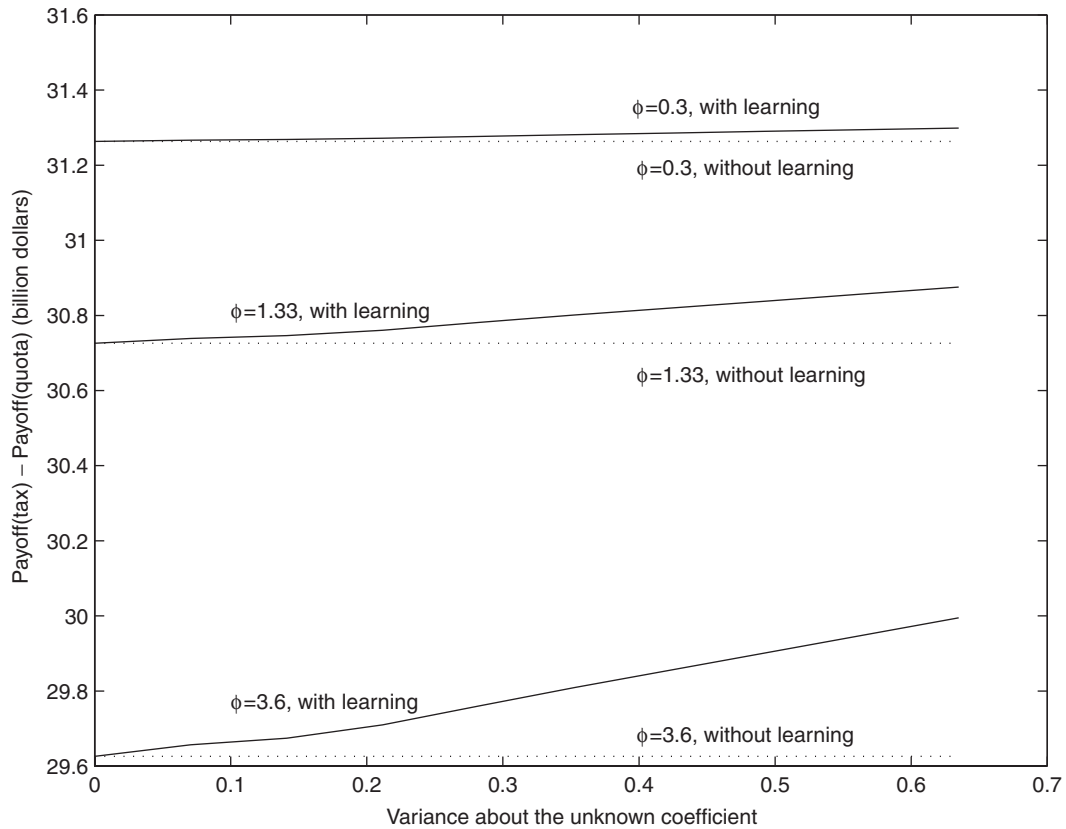


Fig. 3. The effect of learning on the comparison of taxes and quotas.

under taxes and quotas from \$29.6 to \$30 billion, an increase of about 1.3%. For this calibration, anticipated learning has a very small effect on the policy ranking.

The final experiment illustrates the effect of learning about the damage parameter on the expected stock trajectory. Fig. 4 shows the expected stock trajectories under four scenarios: Business as Usual; the case where the regulator believes that $\phi = 1.33$ and does not learn; the case where he knows that $\phi = 3.6$; and the case where the true value is $\phi = 3.6$, the regulator begins with the belief that $\phi = 1.33$ and he anticipates learning. The other parameters equal the baseline values in Table 1, and we use $\sigma_{q,1}^2 = 0.63$. The right panel shows the trajectories over a horizon of 50 periods (500 years) and the left panel shows the trajectories during the first three periods.

All of the stock trajectories begin at the same initial level. Abatement is positive with or without learning, so the stock trajectories under regulation always lie below the BAU trajectory. Initial abatement with learning is lower (emissions are higher) under anticipated learning, compared to no-learning, for given initial beliefs that $\phi = 1.33$. Therefore, the expected trajectory under anticipated learning lies above the expected trajectory without learning, for early periods (the left panel). Both of those trajectories lie above the expected trajectory when the regulator is certain that $\phi = 3.6$. With learning, the regulator increases his subjective expectation of G^* . The level of abatement in the scenario with learning is eventually greater than under no-learning, and the expected stock is lower in the former case. Within 5 periods (fifty years) the expected stock is very close to the level under perfect information about G^* .

For our parameterization, learning occurs quickly enough that the stock remains close to its optimal level, even though the regulator's initial belief about damages is much too optimistic. The stock decays slowly and emissions during a decade are a small fraction of the stock; the stock changes slowly relative to the speed of learning. This example illustrates Remark 1.

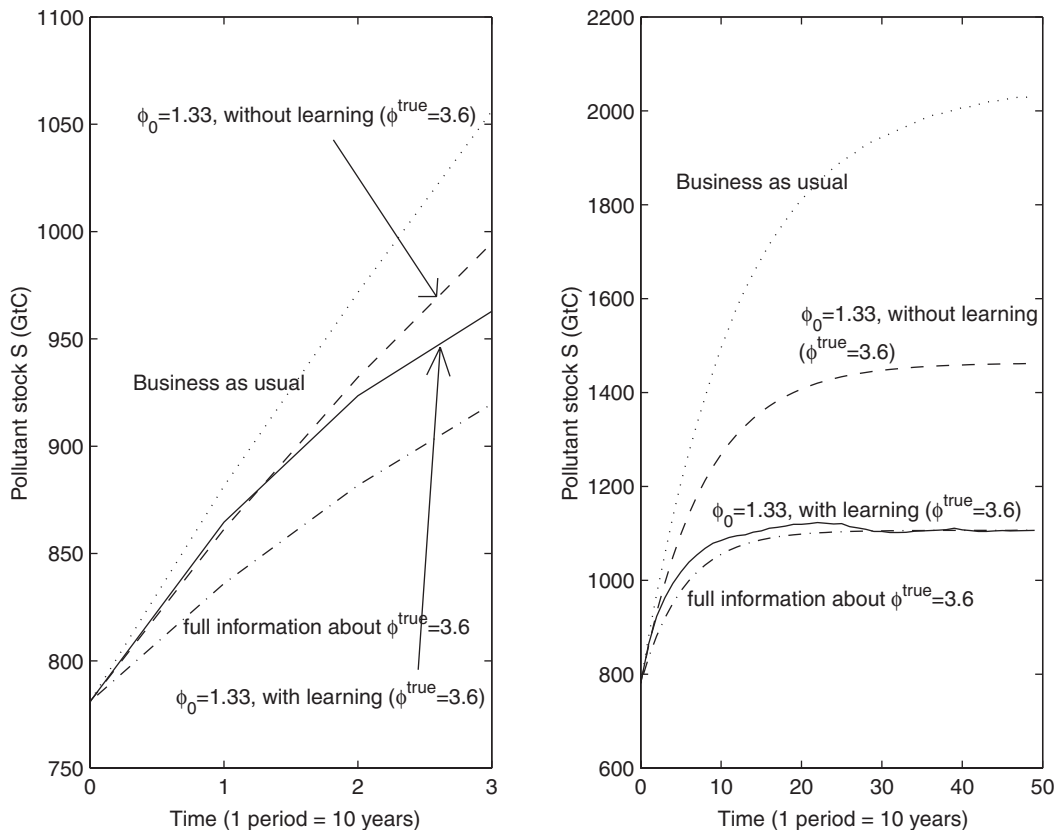


Fig. 4. The effect of learning on the stock trajectory.

6. Model sensitivity

The functional forms for abatement costs and environmental damages (but not for learning in the theoretical model) are restrictive, and the model ignores a number of features of the global warming problem. In assessing the applicability of the model, it is worth distinguishing between these two types of limitations. This section discusses the effect of allowing inequality constraints, catastrophic changes, or different types of abatement activities.

Our results hold even if $\Delta = 1$, i.e. if the stock does not decay. However, the model does not include inequality constraints, such as $x_t \geq 0$. Thus, even when $\Delta = 1$ the stock is reversible. The possibility (in a more general model) that an irreversibility constraint might bind is one reason that anticipated learning could increase abatement [2]. This possibility does not arise in our model.

Kolstad [16] finds that a non-negativity constraint on emissions does not bind for reasonable parameterizations of the DICE model, provided that the stock of abatement capital is reversible. In this situation, in his simulations anticipated learning has negligible effect on abatement. Ulph and Ulph [25] find that a non-negativity constraint on emissions binds only for extreme parameter values, using Maddison's [19] model. When the constraint does bind, the effect of anticipated learning is ambiguous. When it does not bind, learning decreases abatement, typically by a small amount.

We conducted numerical experiments (reported in the online Appendix B3) which show that the probability that it is optimal to set emissions less than 0 is not measurably different from 0. In our model and for our calibration, the constraint $x \geq 0$ is (essentially) never binding. Thus, imposing the constraint $x \geq 0$ would not alter our qualitative results.

The model excludes the possibility of catastrophic changes. There are a variety of ways to model such a change, but the two obvious alternatives are to assume that the probability that the catastrophe occurs is a function of the stock [3] or that the catastrophe occurs when the state crosses an unknown threshold [23]. Tsur and Zemel [24] show that under plausible circumstances, either type of risk reduces the steady state stock (conditional on the catastrophe not having yet occurred). In this case, the risk increases abatement, at least asymptotically.

We are not aware of any analysis of the effect of anticipated learning about such a risk. This anticipation increases the value of being in a pre-catastrophe state in the next period. If, as seems likely, anticipated learning also increases the shadow cost of the stock in a pre-catastrophe state, it encourages abatement in the current period. Again, we would have to use numerical methods to test whether this conjectured effect exists, and if so, whether it would outweigh the effect described in Proposition 2.

In our model, abatement in the current period and in future periods both decrease future stocks, relative to BAU levels. In that sense, the current and future actions are substitutes. In some situations, actions in different periods might be complements. For example, in the current period it may be possible to undertake research (or some other type of investment) that can only be used in subsequent periods.

Karp and Zhang [12] study the case where investment increases a stock of abatement capital that reduces future marginal abatement costs. Current investment and future abatement are complements, but the relation between the two is independent of the information about environmental damages. In this setting, anticipated learning about environmental damages is likely to have the same effect as in the model without capital. Anticipated learning about damages decreases the shadow cost of the stock of pollution, and therefore decreases the level of investment and the level of abatement at a given information state.

However, the fruits of current research (or investment) might be more useful the more we know about global warming. For example, research might enable us to respond more flexibly to future information about global warming. The value of this flexibility might depend on the quality of our information. In this case, anticipated learning increases the shadow value of current research, increasing current R&D. Examples that go in the opposite direction are also easy to construct. For example, anticipated learning might increase the benefit of waiting to invest, until we learn what type of technology is appropriate.

7. Conclusion

There is tremendous scientific uncertainty regarding the relation between greenhouse gasses and global warming; the science is likely to improve. There are many reasons why people disagree about the appropriate response to the danger of greenhouse gasses. One reason is that they hold different views about how the anticipation of learning should affect the regulatory decision. There is good reason for these differing views: even in two-period models the effect of learning is ambiguous. In a more realistic multi-period problem the comparison will also be ambiguous.

Despite the impossibility of a general answer to the question “How does anticipated learning affect optimal regulation of greenhouse gasses?”, economic models can shed light on the issue. We adapted a linear-quadratic model to include anticipated learning about a damage parameter. In this model, anticipated learning always reduces abatement, for a given set of beliefs (Learning eventually increases or decreases abatement, relative to no-learning, depending on how the beliefs change.). The intuition for this result is simple: the ability to respond to new information reduces the threat of future damages, and therefore has an effect that is similar to a more optimistic view of future damages. The simplicity of this intuition is important because it suggests that the result is robust to functional forms. We also showed that anticipated learning favors the use of taxes rather than quotas.

We confirmed numerically that the absence of an explicit irreversibility constraint on the level of emissions is not important in our model. However, the possibility of irreversible catastrophic changes resulting from the accumulation of greenhouse gasses, would be likely to weaken, and might overturn the conclusion that anticipated learning reduces abatement efforts. The assessment of that possibility requires a more complicated model, which could probably be analyzed only by using numerical methods.

An important advantage of the linear-quadratic formulation is that it permits a simple calibration. We know little about the relation between greenhouse gas stocks and global warming, and little about the relation

between global warming and economic costs. Rather than attempting to model both of these relations, we posit a direct relation between stocks and damages which we calibrate using estimates that have appeared in the literature.

We find that the effect of anticipated learning causes a 15–20% reduction in the optimal level of abatement. Even if learning occurs much more slowly than our baseline assumes, it still causes a significant reduction in abatement. Learning has a small effect on the ranking of taxes and quotas. Even if the regulator begins with priors that are much too optimistic, the expected stock trajectory remains close to the full information optimal level. The numerical results suggest that a substantial level of abatement is optimal even with anticipated learning. The results therefore do not support a policy of ignoring the dangers of global warming while learning takes place.

Appendix A. Proofs

The Appendix proves the results stated in Section 3.2. In the proofs, the operator E takes expectations of the moments in the next period, $\chi' = (G', \sigma_G^2)$. Note that the proofs do not require that σ_G^2 be a scalar. In the text we refer to σ_G^2 as the variance, but we pointed out that it could also be viewed as a vector of higher moments.

Proof of Lemma 1. We use a proof by induction. We begin the induction with $n = 0$, where we have the well-known linear-quadratic model without anticipated learning. The value function when $n = 0$ is quadratic in S , i.e. $J(S, \chi, 0) = \lambda_0 + \mu_0 S + \frac{1}{2} \rho_0 S^2$. Here we merely present the formulae for the coefficients of the value function and the control rule when the regulator uses taxes. By comparing the expected payoff functions under taxes and quotas, we can obtain the coefficients of the value function and of the control under quotas simply by replacing z with x and by setting $\sigma_\theta^2 = 0$.

We use the definition:

$$\Psi \equiv \sqrt{(G^2 \beta^2 + 2\beta Gb + 2G\beta^2 b \Delta^2 + b^2 - 2b^2 \beta \Delta^2 + b^2 \beta^2 \Delta^4)} > 0.$$

The formulae for the coefficients of the value function are

$$\rho_0 = \frac{1}{2\beta} (-G\beta + b - b\beta \Delta^2 - \Psi) < 0, \tag{21}$$

$$\mu_0 = \frac{-c\beta\rho_0 + bc - a\beta\rho_0\Delta}{-b + \beta\rho_0 + b\beta\Delta}, \tag{22}$$

$$\lambda_0 = \frac{1}{2} \frac{\beta\rho_0 + b}{(1 - \beta)b^2} \sigma_\theta^2 + \frac{-a^2 - \beta^2 \mu_0^2 - 2a\beta\mu_0 - 2fb + 2f\beta\rho_0}{2(-b + \beta\rho_0)(1 - \beta)}. \tag{23}$$

The optimal control is

$$z_0 = \frac{a + \beta\mu_0 + \beta\rho_0\Delta S}{b - \beta\rho_0}. \tag{24}$$

For $\infty > n \geq 1$ we use an inductive argument to show that the value function is quadratic in S and to obtain the formulae for the coefficients of the value function and the control rule. Suppose that the value function is quadratic at $n - 1$ (as we know is true when $n = 1$): $J(S, G, \sigma_G^2, n - 1) = \lambda_{n-1} + \mu_{n-1} S + \frac{\rho_{n-1}}{2} S^2$. The dynamic programming equation when there are n future learning periods is

$$J(S, G, \sigma_G^2, n) = \max_z f + az - \frac{bz^2}{2} + \frac{\sigma_\theta^2}{2b} - cS - \frac{G}{2} S^2 + \beta E \left(\lambda_{n-1} + \mu_{n-1} (\Delta S + z) + \frac{\rho_{n-1}}{2} (\Delta S + z)^2 + \frac{\rho_{n-1}}{2} \frac{\sigma_\theta^2}{b^2} \right). \tag{25}$$

The functions λ_{n-1}, μ_{n-1} , and ρ_{n-1} depend on the index n and also on the values of the state in the next period, G', σ_G^2 ; we suppress those arguments.

The operator E takes the expectation of the next period value (G', σ_G^2) , conditional on the subjective moments in the current period, (G, σ_G^2) . In writing the DPE (25) we took expectations with respect to the current cost shock, θ . This operation accounts for the presence of the terms involving the variance of the cost shock, σ_θ^2 . We also took expectations with respect to the damage shock, ω , using $E\omega = 1$.

The optimal control rule, obtained by performing the maximization, is

$$z_n = \frac{a + \beta E\mu_{n-1} + \beta E\rho_{n-1}\Delta S}{b - \beta E\rho_{n-1}}. \tag{26}$$

Substituting this control rule into the DPE, using the trial solution $J(S, G, \sigma_G^2, n) = \lambda_n + \mu_n S + \frac{\rho_n}{2} S^2$ and then equating coefficients in orders of S , we obtain the coefficients in the current period

$$\rho_n = \frac{(G\beta + b\beta\Delta^2)E\rho_{n-1} - bG}{b - \beta E\rho_{n-1}} < 0, \tag{27}$$

$$\mu_n = \frac{-bc + \beta(a\Delta + c)E\rho_{n-1} + b\beta\Delta E\mu_{n-1}}{b - \beta E\rho_{n-1}}, \tag{28}$$

$$\lambda_n = \frac{1}{2} \frac{b + \beta E\rho_{n-1}}{b^2} \sigma_\theta^2 + \frac{(-2f\beta - 2\beta^2 E\lambda_{n-1})E\rho_{n-1} + 2a\beta E\mu_{n-1} + 2b\beta E\lambda_{n-1} + 2bf + \beta^2 (E\mu_{n-1})^2 + a^2}{2(b - \beta E\rho_{n-1})}. \tag{29}$$

In order to establish the inequality in Eq. (27) we use the definitions of ρ_n and induction. We start the inductive chain using the inequality in Eq. (21).

If we set $c = \bar{S}G^*$, an unknown parameter, we repeat the steps above. The parameter c is replaced by $\bar{S}G$, its current estimate, in the above equations. \square

Proof of Lemma 2. Using the formula for μ_0 and z_0 , Eqs. (22) and (24), the value of z_0 when $S = 0$ is

$$\frac{a + \beta\mu_0}{b - \beta\rho_0} = \frac{a(\beta\Delta - 1) + \beta c}{\beta\rho_0 + b(\beta\Delta - 1)}.$$

Since $\rho_0 < 0$, this expression is positive if and only if Eq. (12) holds. \square

Proof of Lemma 3. Part (a). We use an inductive proof. In Step 1 we start the induction by showing that $\rho_0(G)$ is a convex function of G , which implies that $\rho_1(G, \sigma_G^2) > \rho_0(G)$. (This inequality is part of the condition that insures that a single learning period reduces emissions, as described in Section 3.1.) In Step 2 we complete the induction.

Step 1: Recall our comment in the text that $J(S, G, \sigma_G^2, 0)$ is independent of σ_G^2 . Consequently, ρ_0 is independent of σ_G^2 . If the regulator does not expect to learn in the future, the optimal decision depends on the expectation of G^* , but not the higher moments of the subjective distribution. The formula for ρ_0 , Eq. (21), implies

$$\frac{d^2\rho_0}{dG^2} = \frac{2\beta^2 b^2 \Delta^2}{(\Psi)^3} > 0,$$

so ρ_0 is a convex function of G . Jensen’s inequality implies,

$$E[\rho_0(G') \mid (G, \sigma_G^2)] > \rho_0(E[G' \mid (G, \sigma_G^2)]) = \rho_0(G) \tag{30}$$

whenever $\sigma_G^2 > 0$. The equality in (30) is a consequence of the fact that $J(S, G, \sigma_G^2, 0)$ is independent of σ_G^2 .

In order to ease the notation, define the right side of Eq. (27) as the function

$$h(r; G) \equiv \frac{G\beta r + b\beta\Delta^2 r - bG}{b - \beta r}.$$

In this function, r is the proxy for $E\rho_{n-1}$; $h(\cdot)$ is strictly increasing in r : $h_r > 0$. Using this definition we rewrite ρ_0 and ρ_1 as

$$\begin{aligned} \rho_0(G) &= h(\rho_0; G), \\ \rho_1(G, \sigma_G^2) &= h(E[\rho_0(G') \mid (G, \sigma_G^2)]; G). \end{aligned}$$

In view of the fact that h is increasing in its first argument, and using inequality (30), we have

$$\rho_1(G, \sigma_G^2) > h(\rho_0(E[G' \mid (G, \sigma_G^2)]); G) = h(\rho_0(G); G) = \rho_0(G)$$

for all G and for all $\sigma_G^2 > 0$.

Step 2: We now compare the functions $\rho_n(G, \sigma_G^2)$ and $\rho_{n-1}(G, \sigma_G^2)$ for $n \geq 2$. Note that we are comparing these two functions evaluated at the same argument, (G, σ_G^2) . We want to show how the number of opportunities for learning, n , affects the functions for given beliefs (i.e., given subjective moments).

Suppose that for some $n \geq 2$ the following relation holds: $\rho_{n-1}(G, \sigma_G^2) > \rho_{n-2}(G, \sigma_G^2)$ for all G and for all $\sigma_G^2 > 0$. (We know from Step 1 that this relation is true for $n = 2$.) This inequality implies that

$$r_{n-1} \equiv E[\rho_{n-1}(G', \sigma_G^2) \mid G, \sigma_G^2] > E[\rho_{n-2}(G', \sigma_G^2) \mid G, \sigma_G^2] \equiv r_{n-2}$$

for all G and for all $\sigma_G^2 > 0$. Consequently,

$$\rho_n = h(r_{n-1}; G) > h(r_{n-2}; G) = \rho_{n-1}.$$

Part (b). We concentrate on the case where c is a known constant and then briefly consider the case where $c_t = -\bar{S}G_t$. We first obtain an intermediate result, we then start the inductive chain by considering the case where $n = 0$ and we then complete the inductive argument.

Step 1: We first note some characteristics of the mapping in Eq. (28), which we repeat for convenience.

$$\mu_n = \frac{-bc + a\beta\Delta E\rho_{n-1} + c\beta E\rho_{n-1} + b\beta\Delta E\mu_{n-1}}{b - \beta E\rho_{n-1}}.$$

Using the right side of this equation, we define the function

$$s \equiv \frac{-bc + a\beta\Delta r + c\beta r + b\beta\Delta q}{b - \beta r}.$$

In this function, r is the proxy for $E\rho_{n-1}$ and q is the proxy for $E\mu_{n-1}$. We note that

$$\frac{ds}{dq} = \frac{b\beta\Delta}{b - \beta r} > 0 \quad \text{for } b - \beta r > 0, \tag{31}$$

$$\frac{ds}{dr} = \beta\Delta b \frac{a + \beta q}{(b - \beta r)^2} > 0 \quad \text{for } a + \beta q > 0, \tag{32}$$

$$\frac{d^2s}{dqdr} = \frac{b\beta^2\Delta}{(b - \beta r)^2} > 0. \tag{33}$$

An increase in q and r when $b - \beta r > 0$ and $a + \beta q > 0$ increases the value of the function s .

The inequalities (31)–(33) imply that the following is a set of sufficient conditions to conclude that $\mu_n > \mu_{n-1}$.

Condition 1. (a) $b - \beta E\rho_{n-1} > 0$. (b) $E\rho_{n-1} > E\rho_{n-2}$. (c) $E\mu_{n-1} > E\mu_{n-2}$. (d) $a + \beta E\mu_{n-1} > 0$.

The inequality in Eq. (27) establishes that Condition 1(a) holds for all $n \geq 1$ and the proof of Lemma 3(a) shows that condition 1(b) holds for all $n \geq 2$.

Step 2: We now establish that $\mu_1 > \mu_0$. To verify this inequality we begin by showing that μ_0 is convex in G . We use the chain rule to obtain

$$\frac{d\mu_0}{dG} = \frac{\partial\mu_0}{\partial\rho_0} \frac{\partial\rho_0}{\partial G}$$

and

$$\frac{d^2\mu_0}{dG^2} = \frac{\partial^2\mu_0}{\partial\rho_0^2} \left(\frac{\partial\rho_0}{\partial G}\right)^2 + \frac{\partial\mu_0}{\partial\rho_0} \left(\frac{\partial^2\rho_0}{\partial G^2}\right).$$

Eq. (22), the formula for μ_0 , implies

$$\begin{aligned} \frac{\partial\mu_0}{\partial\rho_0} &= -\beta b\Delta \frac{c\beta - a + a\beta\Delta}{(-b + \beta\rho_0 + b\beta\Delta)^2}, \\ \frac{\partial^2\mu_0}{\partial\rho_0^2} &= 2\beta^2 b\Delta \frac{c\beta - a + a\beta\Delta}{(-b + \beta\rho_0 + b\beta\Delta)^3}. \end{aligned}$$

The last three equalities imply

$$\begin{aligned} \frac{d^2\mu_0}{dG^2} &= 2\beta^2 b\Delta \frac{c\beta - a + a\beta\Delta}{(-b + \beta\rho + b\beta\Delta)^3} \left(\frac{d\rho_0}{dG}\right)^2 - \beta b\Delta \frac{c\beta - a + a\beta\Delta}{(-b + \beta\rho + b\beta\Delta)^2} \left(\frac{d^2\rho_0}{dG^2}\right) \\ &= \frac{(c\beta - a + a\beta\Delta)}{(-b + \beta\rho_0 + b\beta\Delta)^2} \left[\frac{2\beta^2 b\Delta}{(-b + \beta\rho_0 + b\beta\Delta)} \left(\frac{d\rho_0}{dG}\right)^2 - \beta b\Delta \left(\frac{d^2\rho_0}{dG^2}\right) \right]. \end{aligned} \tag{34}$$

The term in square brackets is negative, so μ_0 is a convex function of G iff $c\beta - a + a\beta\Delta < 0$, i.e. iff Eq. (12) holds. For all G, σ_G^2 Jensen’s inequality implies,

$$E\mu_0(G, \sigma_G^2) > \mu_0(EG, \sigma_G^2) \tag{35}$$

(iff $c < \frac{a(1-\beta\Delta)}{\beta}$) (Recall that μ_0 , like ρ_0 , depends on G but not on σ_G^2 . The expectation of μ_0 with respect to G obviously does depend on σ_G^2).

Eqs. (31)–(33), (35), and the facts that $E\rho_0 > \rho_0$ and $b - \beta E\rho_0 > 0$ establish that $\mu_1 > \mu_0$. Consequently, $E\mu_1 > E\mu_0$ and $a + \beta E\mu_1 > a + \beta E\mu_0 > 0$ (by Lemma 2, given that Eq. (12) holds). Therefore, Conditions 1(a)–(d) hold for $n = 2$.

Step 3: We now consider the case for $n \geq 2$. Suppose that Condition 1 holds for some $n \geq 2$ (The previous paragraph confirms this hypothesis for $n = 2$). For this value of n we have $\mu_n > \mu_{n-1}$ (by part b, Step 1) so $E\mu_n > E\mu_{n-1}$; thus, Condition 1(c) holds for $n + 1$. In addition, $a + E\mu_n > a + E\mu_{n-1} > 0$, so Condition 1(d) holds for $n + 1$. Condition 1(b) holds by virtue of part a of this proof. Condition 1(a) holds in view of Eq. (27). Thus Condition 1 holds for $n + 1$. This completes the proof when c is a known constant.

Step 4: In the case where $c_t = -\sqrt{S}G_t$, the outline of the argument is unchanged, but the formula for $\frac{d^2\mu_0}{dG^2}$ is more complicated because μ_0 now depends directly on $G = EG^*$. There is still the indirect effect of G on μ_0 via the parameter ρ_0 . Taking into account this direct effect, Eq. (34) is replaced by

$$\begin{aligned} \frac{d^2\mu_0}{dG^2} &= \frac{(c\beta - a + a\beta\Delta)}{(-b + \beta\rho_0 + b\beta\Delta)^2} \\ &\times \left[\frac{2\beta^2 b\Delta}{(-b + \beta\rho_0 + b\beta\Delta)} \left(\frac{\partial\rho_0}{\partial G}\right)^2 - \beta b\Delta \frac{\partial^2\rho_0}{\partial G^2} - \frac{\beta b\Delta b\beta^2\Delta}{(-b + \beta\rho_0 + b\beta\Delta)^2} \right]. \end{aligned}$$

The only difference is that an additional negative term appears in the square brackets. The rest of the argument remains the same as in the case where c is a known constant. \square

Remark 2. The proof of Lemma 3 shows that the functions μ_n and ρ_n are increasing in n , and the proof of Lemma 1 show that ρ_n is bounded above by 0. The value function $\lambda_n + \mu_n S + \frac{1}{2}\rho_n S^2$ is bounded above by $\frac{1}{1-\beta}(f + \frac{a^2 + \sigma_a^2}{2b})$, so both λ_n and μ_n are bounded above. Consequently, the two sequences of functions $\mu_n(G, \sigma_G^2)$ and $\rho_n(G, \sigma_G^2)$ converge to functions $\mu_\infty(G, \sigma_G^2), \rho_\infty(G, \sigma_G^2)$ as $n \rightarrow \infty$. These limits are the solution to the fixed point mapping obtained by removing the subscript n in Eqs. (27) and (28). We solve this fixed point mapping to obtain the control rule for $n = \infty$. \square

Proof of Proposition 3. We rearrange Eq. (29) and use a superscript T to denote taxes. Under taxes, the constant (with respect to S) in the value function obeys the difference equation

$$\lambda_n^T = \frac{1}{2} \frac{b + \beta E\rho_{n-1}}{b^2} \sigma_\theta^2 + \beta E\lambda_{n-1}^T + \frac{1}{2} \frac{-2f\beta E\rho_{n-1} + 2bf + 2a\beta E\mu_{n-1} + a^2 + \beta^2(E\mu_{n-1})^2}{b - \beta E\rho_{n-1}}.$$

The expression for the constant (with respect to S) in the value function under quotas, denoted λ_n^Q , obeys the same difference equation, except that the term that multiplies σ_θ^2 is absent (As we noted in the proof of Proposition 1, ρ_n and μ_n are the same under taxes and quotas.). Defining $D_n \equiv \lambda_n^T - \lambda_n^Q$, we have

$$D_n = \frac{1}{2} \frac{b + \beta E\rho_{n-1}}{b^2} \sigma_\theta^2 + \beta ED_{n-1}. \quad (36)$$

We want to show that $D_n > D_{n-1}$. We use an inductive proof, and first show that this inequality holds for $n = 1$. Using Eq. (23) to compute D_0 gives

$$D_0 = \frac{\sigma_\theta^2}{2(1-\beta)b^2} (b + \beta\rho_0).$$

Taking expectations at $n = 1$ and substituting this function into Eq. (36) gives

$$D_1 = \frac{\sigma_\theta^2}{2(1-\beta)b^2} (b + E\beta\rho_0).$$

Using the convexity of ρ_0 in G we confirm that $D_1 > D_0$.

Now suppose that $D_{n-1} > D_{n-2}$ for some $n \geq 2$ (We have already confirmed that this hypothesis is true, using $n = 2$). This hypothesis implies $ED_{n-1} - ED_{n-2} > 0$. We have

$$D_n - D_{n-1} = \frac{\beta(E\rho_{n-1} - E\rho_{n-2})}{2b^2} \sigma_\theta^2 + \beta(ED_{n-1} - ED_{n-2}).$$

The first term on the right side is positive by Lemma 3a and the second is positive by the hypothesis, thus confirming $D_n - D_{n-1} > 0$.

This proof does not involve the parameter c , so it does not matter whether we view it as a known constant or as $\bar{S}G$.

References

- [1] K. Arrow, A. Fisher, Environmental preservation, uncertainty, and irreversibility, *Quart. J. Econ.* 88 (1974) 312–319.
- [2] G. Chichilnisky, G. Heal, Global environmental risks, *Soc. Choice Welfare* 7 (4) (1993) 65–86.
- [3] H. Clarke, W. Reed, Consumption/pollution tradeoffs in an environment vulnerable to pollution-related catastrophic collapse, *J. Econ. Dynam. Control* 18 (1994) 991–1010.
- [4] L. Epstein, Decision-making and the temporal resolution of uncertainty, *Int. Econ. Rev.* 21 (1980) 269–284.
- [5] A. Fisher, U. Narain, Global warming, endogenous risk and irreversibility, *Environ. Resource Econ.* 25 (2003) 395–416.
- [6] C. Gollier, B. Jullien, N. Treich, Scientific progress and irreversibility: an economic interpretation of the ‘precautionary principle’, *J. Public Econ.* (107) (2000) 463–473.
- [7] W. Greene, *Econometric Analysis*, Prentice-Hall, New Jersey, 2000.
- [8] G. Heal, B. Kristrom, Uncertainty and climate change, *Environ. Resource Econ.* 22 (1–2) (2002) 3–39.
- [9] C. Henry, Investment decisions under uncertainty: the irreversibility effect, *Amer. Econ. Rev.* 64 (1974) 1006–1012.
- [10] M. Hoel, L. Karp, Taxes and quotas for a stock pollutant with multiplicative uncertainty, *J. Public Econ.* 82 (2001) 91–114.
- [11] M. Hoel, L. Karp, Taxes versus quotas for a stock pollutant, *Resource Energy Econ.* 24 (2002) 367–384.
- [12] L. Karp, J. Zhang, Controlling a stock pollutant with endogenous abatement capital and asymmetric information, *Giannini Foundation of Agricultural Economics, Working Paper no. 928, University of California, Berkeley, 2002, <http://are.Berkeley.EDU/karp/>.*
- [13] L. Karp, J. Zhang, Regulation of stock externalities with correlated abatement costs, *Environ. Resource Econ.* 32 (2005) 273–300.
- [14] D.L. Kelly, C.D. Kolstad, Bayesian learning growth and pollution, *J. Econ. Dynam. Control* 23 (4) (1999) 491–518.
- [15] P. Kennedy, Learning about environmental damage: implications for emissions trading, *Can. J. Econ.* 32 (5) (1999) 1313–1327.
- [16] C.D. Kolstad, Fundamental irreversibilities in stock externalities, *J. Public Econ.* 60 (2) (1996) 221–233.
- [17] C.D. Kolstad, Learning and stock effects in environmental regulation: the case of greenhouse gas emissions, *J. Environ. Econ. Manage.* 31 (1) (1996) 1–18.

- [18] A. Leach, The climate change learning curve, University of Montreal, Working Paper, 2004.
- [19] D. Maddison, A cost-benefit analysis of slowing climate change, *Energy Policy* 23 (1995) 337–346.
- [20] R. Newell, W. Pizer, Regulating stock externalities under uncertainty, *J. Environ. Econ. Manage.* 45 (2003) 416–432.
- [21] W.D. Nordhaus, D. Popp, What is the value of scientific knowledge? An application to global warming using the PRICE model, *Energy J.* 18 (1) (1997) 1–45.
- [22] S.C. Peck, T.J. Teisberg, Global warming uncertainties and the value of information: an analysis using CETA, *Resources Environ. Econ.* 15 (1) (1993) 71–97.
- [23] Y. Tsur, A. Zemel, Accounting for global warming risks: resource management under event uncertainty, *J. Econ. Dynam. Control* 20 (1996) 1289–1305.
- [24] Y. Tsur, A. Zemel, Pollution control in an uncertain environment, *J. Econ. Dynam. Control* 22 (1998) 967–975.
- [25] A. Ulph, D. Ulph, Global warming, irreversibility and learning, *Econ. J.* 107 (1997) 636–650.
- [26] M. Weitzman, Prices versus quantities, *Rev. Econ. Stud.* 41 (4) (1974) 477–491.