



The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

Papers downloaded from AgEcon Search may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

926

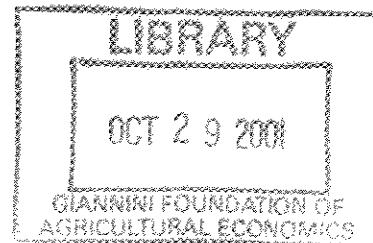
DEPARTMENT OF AGRICULTURAL AND RESOURCE ECONOMICS AND POLICY
DIVISION OF AGRICULTURE AND NATURAL RESOURCES
UNIVERSITY OF CALIFORNIA AT BERKELEY

WORKING PAPER NO. 926

BAYESIAN LEARNING AND THE REGULATION OF
GREENHOUSE GAS EMISSIONS

by

Larry Karp and Jiangfeng Zhang



Copyright © 2001 by Larry Karp and Jiangfeng Zhang. All rights reserved. Readers may make verbatim copies of this document for noncommercial purposes by any means, provided that this copyright notice appears on all such copies.

California Agricultural Experiment Station
Giannini Foundation of Agricultural Economics
August 2001

Bayesian Learning and the Regulation of Greenhouse Gas Emissions

Larry Karp and Jiangfeng Zhang*

August 15, 2001

Abstract

We study the importance of anticipated learning – about both environmental damages and abatement costs – in determining the level and the method of controlling greenhouse gas emissions. We also compare active learning, passive learning, and parameter uncertainty without learning. Current beliefs about damages and abatement costs have an important effect on the optimal level of emissions. However, the optimal level of emissions is not sensitive either to the possibility of learning about damages, or to the type of learning (active or passive). Taxes dominate quotas, but by a small margin.

JEL Classification numbers: C11; C61; D8; H21; Q28.

Key Words: Climate change; Uncertainty; Bayesian learning; Asymmetric information; Choice of instruments; Dynamic optimization.

*Department of Agricultural and Resource Economics, University of California, Berkeley, CA 94720. Email: karp@are.berkeley.edu; jzhang@are.berkeley.edu.

We thank Alan Auerbach, Trudy Cameron, Jean-Daniel Saphores, and Martin Weitzman for their helpful comments on an earlier version of this paper. We also benefited from the comments of seminar participants at the 2001 AERE Workshop on Assessing and Managing Environmental and Public Health Risks in Bar Harbor, and the 7th International Conference of the Society for Computational Economics at Yale University.

1 Introduction

Many pollutant stock externality problems – notably the problem of controlling greenhouse gasses – are complicated by the uncertainty of abatement costs and environmental damages. In setting environmental policies, a regulator who has imperfect information about these economic and environmental costs and damages should recognize that information may improve over time. We construct a model of a stock externality that includes endogenous learning about both abatement costs and stock-related damages. We calibrate the model to describe the problem of global warming and solve it numerically. The results show how uncertainty and learning affect both the optimal level of control and the comparison of taxes and quantity restrictions.

Several papers ([7], [8], [13], and [26]) 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 [37] static model continues to hold: a steeper marginal environmental damage curve, or a flatter marginal abatement cost curve favors the use of quotas. These models assume that the regulator knows the parameters of the damage function.

We extend these models to describe the situation where the regulator does not know – but learns about – the true relation between pollutant stocks and environmental damages. We identify the effect of parameter uncertainty in the absence of learning by solving a certainty equivalent version of this model. The possibility of learning about these uncertain parameters causes a qualitative change in the optimization problem. This difference enables us to identify the effect of learning, as distinct from the intrinsic effect of parameter uncertainty.

The uncertainty about both abatement costs and stock-related damages, coupled with the belief that we will obtain better information over time, is central to the current debate over efforts to reduce carbon emissions. If we were convinced that this uncertainty would persist indefinitely, we could model it like any other form of randomness. The possibility that we will learn more about the relation between greenhouse gasses and global warming 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.

Chichilnisky and Heal [1] explain why anticipated learning may lead to greater initial abatement when irreversibilities are important. Ulph and Ulph [36] explain why the relation between learning and the amount of abatement is ambiguous. Their numerical results suggest that under plausible circumstances, anticipated learning decreases initial abatement, and that usually the

magnitude of the effect of learning is small. Our results are consistent with [36], although our model is very different.

The uncertainty about abatement costs is also an important component of the debate. Opponents of the Kyoto Protocol frequently claim that the economic cost of reducing emissions is large; proponents point to reasons (e.g. positive externalities in innovation) why abatement costs will be small. We do not know the actual costs.

Much of the existing literature concerning climate change uncertainty assumes that it will eventually be resolved¹ (e.g., Kennedy [18], Kolstad [19] [20], Manne and Richels [21], Nordhaus and Popp [30], Peck and Teisberg [31]). Nordhaus and Popp [30] and Peck and Teisberg [31] consider the difference between “act and learn” and “learn and act”. All these papers focus on the effect of passive learning; the exogenous arrival of information decreases uncertainty. Passive learning may occur all at once as in Kennedy [18], Kolstad [19], or more gradually as a function of time as in Kolstad [20].²

The assumption of passive learning ignores the possible impact of the regulator’s decisions on the learning process. Policy decisions will affect future levels of stock, and the magnitude of these levels may affect the amount of information that the regulator acquires. The regulator is unlikely to manipulate global carbon stocks in order to learn the true relation between stocks and damages.³ However, the regulator should recognize that there is a relation between control decisions and learning.

Kelly and Kolstad [17] consider active learning about the relation between greenhouse gas levels and global mean temperature changes. Their simulations show that abatement is sensitive to the state of knowledge, and they find that learning occurs slowly.

We model active learning about stock-related damages. There are two main differences in focus between our paper and [17], in addition to many technical differences. First, we allow for uncertainty and learning about abatement costs as well as stock-related damages. Arguably, uncertainty about abatement costs is as important an issue as uncertainty about environmental damages. Second, we compare the use of taxes and quotas in this setting. Our paper contributes to two related but distinct literatures, the control of emissions under learning about damages, and the control of emissions under asymmetric information and learning about abatement costs.

¹Pizer [32] considers persistent uncertainty. He studies an open-loop equilibrium, where no learning occurs.

²Another literature emphasizes the role of learning and irreversibilities in abatement capital and pollution stocks; see Fisher and Narain [5].

³For some problems, particularly where the stock decays rapidly, that kind of active learning might make sense.

Our model has several key features. (i) It allows for the possibility that additional stocks cause very large increases in damages, even if current expected damages are moderate. (ii) There is an objective stochastic relation between stocks and damages. The regulator does not know one of the parameters of this relation. By observing stocks and stock-related damages the regulator learns about the unknown parameter, but does not learn its exact value in finite time. (iii) Our model makes it easy to distinguish between the intrinsic effect of parameter uncertainty, and the effect of anticipated learning about the uncertain parameter(s). Thus, we can isolate the effect of learning. We can also distinguish between the effect of active and passive learning. (iv) The firms have better information than the regulator concerning abatement costs. These costs change over time, and may be serially correlated. Under some circumstances, the regulator is able to learn about abatement costs, but firms always remain better informed about them.

Our principal policy conclusion is that the optimal level of emissions is insensitive to the anticipation of (either active or passive) learning about environmental damages. This conclusion is important, because both opponents and proponents of imposing stricter abatement rules now have used this anticipated learning to support their position. Our results suggest that this hope for improved science may be a red herring in the discussion of current abatement policies. Those policies should be based on our current beliefs about the relation between carbon stocks and environmental damages – i.e. on the best current science – rather than on our beliefs that information will improve.⁴ Of course, this conclusion is based on a particular parameterization of a particular model.

The next section discusses some basic ideas that are important in understanding our model and results. We then describe how the regulator learns about stock-related damages and abatement costs. Using previously published data and conjectures about global warming, we calibrate the model, and then solve it to show how learning affects the optimal level of emissions and the choice between taxes and quotas.

⁴Our model has nothing to say about the argument that we should postpone abatement until technological improvements reduce abatement costs.

2 Background

Since we draw on ideas from several specialized fields, we begin with a basic discussion. Here we explain the intuition for the ranking of taxes and quotas when there is asymmetric information about abatement costs. Then we describe the difference between active and passive learning, and we discuss the role of the Principal of Certainty Equivalence.

In a period, the firm – but not the regulator – knows the current marginal abatement cost function. If faced with an emissions tax, the firm sets marginal abatement costs equal to the tax. Thus, under taxes the regulator chooses the current level of marginal cost and regards emissions as random. Under quotas the regulator chooses the current level of emissions and regards the level of marginal cost as random.⁵ This difference causes the expected payoffs to differ under the two policies.

The convexity of damages means that a mean-preserving spread in emissions (or stocks) increases expected damages. Steeper marginal damages – i.e. a more convex damage function – increase the importance of controlling emissions (or stocks) exactly, and therefore favor the use of quotas. When marginal abatement costs are steeper, it is more important for the firm to be able to adjust emissions in light of the current cost shock. Thus, steeper marginal abatement costs favor the use of taxes.

It is more important to control emissions (and stocks) exactly when the stock is more persistent. Thus, a more persistent stock favors the use of quotas. Since current emissions cause damages in the future, a higher discount factor increases the importance of controlling exactly the current level emissions. Thus, a higher discount factor increases the range of other parameter values for which a quota is preferred. However, the ranking of policy depends on the comparison of the present discounted value of abatement costs plus damages under the two policies. A larger discount factor increases the present discounted streams under both policies, and therefore can increase the magnitude by which either policy is preferred to the other. The effect on policy ranking of the degree of autocorrelation between cost shocks depends on the information structure. We return to this issue in Section 3.3.

In our model, future observations of stocks and damages enable the regulator to learn about the true relation between these variables. Since the “signal” (damages) depends on the stock,

⁵In common with most papers on this topic, we assume that the quota is binding with probability 1. Costello and Karp [3] examine a model with flow-dependent damages, in which the regulator can use a non-binding quota to learn about abatement costs.

and since the regulator can influence the evolution of the stock, he is able to influence the future information. If the regulator understands this relation, learning is “active”. If instead he takes the future signals as exogenous, learning is “passive”. The regulator is unlikely to manipulate global carbon stocks in order to provide better information about how these stocks affect environmental damages. That is, he is likely to behave almost as if future signals were exogenous. If he does so, active and passive learning would lead to almost the same control rule. It is easier to solve the model of passive rather than active learning. Of course, if future information really is unrelated to stock levels, the model of passive learning is more accurate. For these reasons, we also consider the model of passive learning.

We assume that abatement costs are quadratic in abatement, which is defined as the “business-as-usual” (BAU) level of emissions minus the actual level of emissions. BAU emissions are realizations of a random process, and in each period the firm but not the regulator knows the current value of this random variable. This formulation means that the cost shock affects the intercept but not the slope of marginal costs. We also assume that environmental damages are quadratic in carbon stocks.

These functional assumptions mean that *the model with passive learning* satisfies the Principal of Certainty Equivalence, and this in turn means that the expected stock and flow trajectories (but not their higher moments) are identical under optimal taxes and quotas. In this case, the ranking of taxes and quotas depends on higher moments of the trajectories. This Certainty Equivalence property does not hold in the model of active learning; however, as explained above, we expect that in the case of global warming, the two models would lead to similar results. In that case, the Certainty Equivalence property holds approximately under active learning.

The functional assumptions greatly facilitate the solution to the model of passive learning, but not the model of active learning – which requires different methods. Since the functional assumptions are unimportant for the latter model, and since some readers might regard them as unattractive, it is worth explaining why we use them. There are two reasons.

First, these assumptions lead to a relatively simple calibration of the model. Since our results are model- and parameter-specific, this transparency is important. The calibration depends on the assumed magnitude of abatement cost and of environmental damage. We can easily determine how robust the policy implications are to these magnitudes.

Second, the functional assumptions imply the (exact or approximate) equivalence of ex-

pected trajectories under taxes and quotas. If the research objective was to understand how different policies affect the optimal level of abatement, this Certainty Equivalence property would obviously be a disadvantage (since it renders the comparison trivial). However, if our goal is to compare the two policies given the same target level of stock trajectories, the Certainty Equivalence property is a great advantage. In a general model, such a comparison might be sensitive to the particular target trajectory. For example, quotas might achieve a higher payoff given the target that is optimal under quotas, and taxes might achieve a higher payoff given the target that is optimal under taxes. This kind of sensitivity cannot arise in our model, where the optimal targets are (exactly or approximately) the same.

3 The Model

We discuss environmental damages first. The regulator learns about an unknown parameter that determines the stochastic relation between stocks and damages. Our formulation has two important features: the variability of damages increases with the stock, and the marginal damage can be extremely large but never negative. We then discuss how the regulator learns about abatement cost shocks by using a tax, and we present the regulator's optimization problem. We explain how this model enables us to distinguish the intrinsic effect of parameter uncertainty from the effect of anticipated learning. We also explain how to modify the model to replace active with passive learning.

3.1 Uncertain Environmental Damages

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 pollutant stock S_t decays at a constant rate. With the fraction $\Delta > 0$ of the pollutant stock lasting into the next period, the growth equation for S_t is:

$$S_{t+1} = \Delta S_t + x_t.$$

In period t the stochastic stock-related environmental damage equals

$$D(S_t, \omega_t; g) = e^g (S_t - \bar{S})^2 \omega_t, \quad \omega_t \sim \text{i.i.d. lognormal} \left(-\frac{\sigma_\omega^2}{2}, \sigma_\omega^2 \right). \quad (1)$$

The slope of marginal damages is $2e^g$. For our purposes, the important variable is e^g rather than g , but the formulation in equation (1) provides a convenient way to introduce parameter uncertainty. The parameter g (and therefore, the slope of marginal damages) is unknown, and ω_t is a random damage shock. Damage is convex in the pollutant stock ($D_{SS} > 0$); \bar{S} is the stock level associated with zero environmental damage. For example, \bar{S} is the pre-industrial CO₂ level.⁶ The distributional assumption about ω imply⁷:

$$E(\omega) = 1, \quad Var(\omega) = e^{\sigma_\omega^2} - 1, \quad E(\ln \omega) = -\frac{\sigma_\omega^2}{2}, \quad Var(\ln \omega) = \sigma_\omega^2.$$

The relation between stocks and damages is both uncertain and stochastic. The random shock ω_t prevents the regulator from ever learning the true value of the unknown parameter g . The regulator does not know whether a high level of damage is caused by a large value of g or by a large realization of the random variable ω . Future observations on damages and stocks lead to better estimates of g .

The regulator begins with a prior belief on g , summarized in a normal distribution with mean g_t and variance $\sigma_{g,t}^2$:

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

The subscript t denotes information (a belief) at the beginning of period t . Given distribution (2), e^g has a lognormal distribution with expected value $G_t \equiv \exp(g_t + \frac{1}{2}\sigma_{g,t}^2)$ and variance $\exp(2g_t + \sigma_{g,t}^2)(\exp(\sigma_{g,t}^2) - 1)$. Since damages are a product of independent lognormally distributed variables, the subjective expectation of damages, given the current stock and current beliefs, is lognormally distributed with mean and variance:

$$E[D(S_t, \omega_t; g)|\Omega_t] = \exp(g_t + \frac{1}{2}\sigma_{g,t}^2) (S_t - \bar{S})^2 = G_t (S_t - \bar{S})^2, \quad (3)$$

$$Var[D(S_t, \omega_t; g)|\Omega_t] = \exp(2g_t + \sigma_{g,t}^2) [\exp(\sigma_{g,t}^2 + \sigma_\omega^2) - 1] (S_t - \bar{S})^4. \quad (4)$$

The information set is $\Omega_t \equiv [S_t, g_t, \sigma_{g,t}^2]$, the current pollutant stock level, and the subjective mean and variance of g . Both the expectation and the variance of environmental damages are increasing in each element of the information set (for $S_t > \bar{S}$).

⁶We obtain the special case in which environmental damages are caused by the flow rather than the stock of pollutants, by setting the stock persistence $\Delta = 0$ and defining \bar{S} as the flow level associated with zero damages.

⁷See Greene [6] page 69. Note that the convention here is that the parameters of the lognormal random variable ω denote the mean and variance of $\ln \omega$ – not the mean and variance of ω .

This model of damages has two appealing features. First, a higher stock increases the probability of extremely high damages, associated with extreme events. Some physical scientists believe that the greatest threat of greenhouse gasses is that they increase the probability of such events. Second, the slope of the marginal damage of stocks, $2e^g$, is always positive and might be arbitrarily large. Marginal expected damages are proportional to $(S_t - \bar{S})$ and expected damages are proportional to $(S_t - \bar{S})^2$. The difference in the magnitudes of these two expectations implies that this model allows for the possibility that a small increase in stock can cause a large increase in expected damages, even if expected damages at the current stock level are moderate.⁸

The coefficient of variation (CV) of damages is increasing in both $\sigma_{g,t}^2$ and σ_ω^2 :

$$CV [D(S_t, \omega_t; g) | \Omega_t] = [\exp(\sigma_{g,t}^2 + \sigma_\omega^2) - 1]^{\frac{1}{2}}. \quad (5)$$

3.2 Learning about Environmental Damages

Environmental damages in period t depend on the pollutant stock S_t , the true value of g and the damage shock ω_t . After observing damages and the current stock, the regulator updates his belief about g . The natural logarithm of damages, from equation (1), is normally distributed. The “moment estimator” of g , denoted \hat{g}_t , is

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

with variance $\sigma_{\hat{g}}^2 = \sigma_\omega^2$.

The Bayesian regulator updates his priors using the moment estimator. The posterior g is normally distributed with the posterior mean g_{t+1} and posterior variance $\sigma_{g,t+1}^2$ given by

$$g_{t+1} = \frac{\sigma_\omega^2}{\sigma_\omega^2 + \sigma_{g,t}^2} g_t + \frac{\sigma_{g,t}^2}{\sigma_\omega^2 + \sigma_{g,t}^2} \hat{g}_t, \quad (6)$$

$$\sigma_{g,t+1}^2 = \frac{\sigma_{g,t}^2 \sigma_\omega^2}{\sigma_\omega^2 + \sigma_{g,t}^2} \Rightarrow \sigma_{g,t}^2 = \frac{\sigma_{g,0}^2 \sigma_\omega^2}{\sigma_\omega^2 + t \sigma_{g,0}^2}, \quad (7)$$

⁸In an earlier working paper [14] we modeled damages as

$$D(S_t, \omega_t; g) = g (S_t - \bar{S})^2 + \omega_t, \quad \omega_t \sim i.i.d. \text{ normal}(0, \sigma_\omega^2), \quad \forall t \geq 0.$$

We abandoned this formulation because it has neither of the two features described in the text.

where $\sigma_{g,0}^2$ is the prior at the beginning of the initial period, $t = 0$. (Greene [6], pages 407-410). The posterior mean g_{t+1} depends on the prior g_t , the magnitude of uncertainty ($\sigma_{g,t}^2$ and σ_ω^2), and on the data S_t and D_t (via \hat{g}_t). The posterior variance $\sigma_{g,t+1}^2$ depends on t , σ_ω^2 , and $\sigma_{g,0}^2$, but not on the data. Note that the subjective variance of the slope of marginal damages, $2e^g$, *does* depend on the data. (See equation (4).)

The variance of g decreases monotonically with time. A regulator who begins with an imprecise estimate of g ($\sigma_{g,0}^2$ is large) initially puts a large weight on the moment estimator. As time progresses the regulator becomes more convinced about the true value of g , and puts less weight on the moment estimator. A smaller value of σ_ω^2 means that the new observations are more informative, so the regulator learns more quickly about the true value of g . Learning about damages is “active” because the amount of new information depends on stocks (via \hat{g}_t), a variable which the regulator is able to influence.

The unknown slope of marginal damages is twice the value of e^g ; the regulator’s expectation of this parameter is $G_t \equiv \exp(g_t + \frac{1}{2}\sigma_{g,t}^2)$. Since $\sigma_{g,t}^2$ is deterministic and decreases monotonically, it might seem that the current expectation of future values of G_t would decrease. However, we have (See Appendix 7.2 for details.):

Remark 1 *The regulator’s current belief about e^g is an unbiased estimate of the belief that he will hold in the future: $E_t G_{t+\tau} = G_t$, $\forall \tau \geq 0$.*

3.3 Learning about Abatement Cost

This section compares firms’ abatement cost and emission response under taxes or quotas. Suppose the representative firm’s business-as-usual (BAU) level of emissions in period t is $x_t^b = \bar{x} + \tilde{\theta}_t$ where $\tilde{\theta}_t$ is a random shock.⁹ With an actual emission level $x_t < x_t^b$, the firm’s abatement cost is a quadratic function of abatement $A(x_t) = \frac{b}{2}(x_t^b - x_t)^2$ with $b > 0$. The firm’s benefit (its cost saving) from higher emission equals the abatement costs that it avoids. Defining the cost shock $\theta_t \equiv b\tilde{\theta}_t$, we write the benefit as a linear-quadratic function, concave in

⁹In [15] we study the case where business-as-usual emissions change endogenously due to firm’s investment in abatement capital. There we ignore the uncertainty of environmental damages.

the emission with an additive cost shock¹⁰

$$B(x_t, \theta_t) = f + (a + \theta_t)x_t - \frac{b}{2}x_t^2.$$

The cost shock θ_t is the firm's private information, and it follows an $AR(1)$ process:

$$\theta_t = \rho\theta_{t-1} + \mu_t; \quad \mu_t \sim i.i.d. (0, \sigma_\mu^2), \quad \forall t \geq 1, \quad (8)$$

with $-1 < \rho < 1$. The regulator has a subjective prior with mean $\bar{\theta}_0$ and variance σ_0^2 on the initial cost shock θ_0 . The i.i.d. random process $\{\mu_t\}$ ($t \geq 1$) has mean 0 and common variance σ_μ^2 and is uncorrelated with θ_0 .

The firm observes θ_t before making its current emission decision. If the regulator sets a unit tax p_t on emissions, firms in each period maximize the abatement cost saving minus the tax payment:

$$\underset{x_t}{\text{Max}} \Pi_t = B(x_t, \theta_t) - p_t x_t = \left[f + (a + \theta_t)x_t - \frac{b}{2}x_t^2 \right] - p_t x_t.$$

The firms' first order condition implies that the optimal emission under a tax is

$$x_t^* = \frac{a - p_t}{b} + \frac{\theta_t}{b}. \quad (9)$$

The flow of emissions and thus the future pollutant stock S_{t+1} is stochastic since it depends on the cost shock.

Under taxes, the regulator infers the value of θ_{t-1} with a one period lag, using equation (9) and the observed x_{t-1}^* . The regulator's posterior beliefs about the cost shock satisfy $E^{Tax}(\theta_t | \theta_{t-1}) = \rho\theta_{t-1}$ and $Var^{Tax}(\theta_t | \theta_{t-1}) = \sigma_\mu^2$ ($t \geq 1$). Learning about the cost shock is "passive" rather than "active": the level of the tax does not affect the amount of cost-related information that the regulator acquires. Hereafter, when discussing taxes we take the regulator's control variable to be the expected level of emissions, z_t , rather than the tax p_t . Using equation (9) we have

$$z_t = E^{Tax}(x_t^* | \theta_{t-1}) = \frac{a - p_t}{b} + \frac{1}{b}E^{Tax}(\theta_t | \theta_{t-1}). \quad (10)$$

If firms have heterogenous cost shocks, and are allowed to trade quotas, the equilibrium quota price conveys the same information as does the equilibrium response to an emissions

¹⁰The parameters satisfy $f = -\frac{b}{2}\bar{x}^2$ and $a = b\bar{x}$. We ignore the effect of $\bar{\theta}$ on f since f has no effect on the regulator's control.

tax. If firms cannot trade quotas (or do not wish to do so because of transactions costs or because they are homogenous) then the regulator never learns the current value of the cost shock. Thus, there is an important informational difference between tradable and non-tradable quotas. Our previous paper [13] explores this difference in a model without uncertainty about damages. That paper shows that the informational advantage of taxes is a major reason that they dominate non-traded quotas. In the interests of brevity, here we consider only the case where quotas are not traded. Given the hostility (from some quarters) to international trade in carbon permits, we think that this is the relevant comparison. The assumption that quotas are not traded favors the use of taxes when $\rho \neq 0$.

Under a non-tradable quota, the regulator does not learn about cost shocks. The regulator's subsequent beliefs of the cost shock satisfy $E^{Quota}(\theta_t|\bar{\theta}_0) = \rho^t \bar{\theta}_0$. Since θ appears linearly in the regulator's objective function, the expected payoff does not depend on the variance of θ . If $\rho = 0$, neither policy provides an informational advantage.

Appendix 7.1 summarizes the equations of motion and the expectations of abatement costs under taxes and quotas.

3.4 Optimal Regulation

The regulator always uses taxes or always uses quotas.¹¹ He chooses the policy level in each period based on current information, in order to maximize the expectation of the discounted stream of future cost savings minus environmental damages:

$$E_t \sum_{j=0}^{\infty} \beta^j \{ B(x_{t+j}, \theta_{t+j}) - D(S_{t+j}, \omega_{t+j}; g) \}. \quad (11)$$

The discount factor is β . The regulator anticipates learning about g . A tax-setting regulator anticipates learning about the cost shock θ . E_t is the expectations operator conditional on the regulator's information. This information consists of $\Omega_t \equiv [S_t, g_t, \sigma_{g,t}^2]$ and $E^i(\theta_t|I_{i,t})$ ($i=\text{tax}$ or quota) with $I_{i,t}$ equal to θ_{t-1} or $\bar{\theta}_0$. The regulator takes expectations with respect to the cost shock θ , the damage shock ω , and the unknown damage parameter g .

Maximization of the expression in (11) results in a value function $J^i(\cdot)$ ($i=\{\text{tax, quota}\}$)

¹¹Since we rely on numerical methods, we could consider more sophisticated policies, such as a two-part tax. In order to focus on the policy choice under learning, we restrict the regulator to the limited policy menu of taxes and quotas.

satisfying the following dynamic programming equation (DPE)¹²:

$$J_t^i [S_t, g_t, \sigma_{g,t}^2, E^i(\theta_t | I_{i,t})] = \underset{c^i}{\text{Max}} \{ E^i [B(x_t, \theta_t) | I_{i,t}] - E [D(S_t, \omega_t; g) | \Omega_t] \\ + \beta E_t J_{t+1}^i [S_{t+1}, g_{t+1}, \sigma_{g,t+1}^2, E^i(\theta_{t+1} | I_{i,t+1})] \} \quad (12)$$

The control variable c^i is the quota level x_t under quotas, and the expected emission z_t under taxes. The superscript $i = (\text{taxes, quotas})$ on the expectations operator emphasizes that information may be different under taxes and quotas.

The maximization problem with quotas is a special case of the problem with taxes, obtained by setting $\mu_t \equiv 0 \equiv \sigma_\mu^2$. Thus, we focus on the problem with taxes. The regulator chooses the optimal control z_t in each period, and then observes firms' emission responses and environmental damages. These observations enable the regulator to update the priors on the cost shock θ_t and on g . In setting the optimal control z_t , the regulator considers its effect on current expected abatement costs and on future state variables. The control variable z_t affects the current abatement costs and the pollutant stock S_{t+1} directly. Although z_t has no direct effect on the posterior g_{t+1} , it influences future beliefs about g due to the dependence of g_{t+2} on S_{t+1} . The independence of the posterior mean g_{t+1} and variance $\sigma_{g,t+1}^2$ on z_t leads to the necessary condition

$$z_t = \frac{1}{b} \left[a + E^i(\theta_t | I_{i,t}) + \beta E_t \frac{\partial J_{t+1}^i [S_{t+1}, g_{t+1}, \sigma_{g,t+1}^2, E^i(\theta_{t+1} | I_{i,t+1})]}{\partial S_{t+1}} \right]. \quad (13)$$

Using the definition of z_t , equation (10), the optimal tax level equals the negative discounted shadow value of future pollutant stocks:

$$p_t^* = -\beta E_t \frac{\partial J_{t+1}^{tax} [S_{t+1}, g_{t+1}, \sigma_{g,t+1}^2, E^{tax}(\theta_{t+1} | \theta_t)]}{\partial S_{t+1}}.$$

3.5 Isolating the Effect of Anticipated Learning

Parameter uncertainty changes the optimal level of abatement, even in the absence of anticipated learning about damages. We refer to this change as the “intrinsic effect” of parameter uncertainty. The anticipation of learning may cause an additional change (“the learning effect”) in the level of abatement. The learning effect can differ depending on whether learning is active or passive.

¹²The time subscript for J denotes changes in the variance $Var^{Tax}(\theta_t | \theta_{t-1})$ between $t = 0$ and $t \geq 1$. The value function is independent of the variance under quotas.

This section considers two simpler versions of our model. The first version models parameter uncertainty without learning of any kind, in order to distinguish between the intrinsic effect of parameter uncertainty and the learning effect. The second version models uncertainty with passive learning. This model clarifies the distinction between active and passive learning, and it helps in understanding the relation between taxes and quotas under active learning. It is also useful for interpretation of numerical results described below. This model may also be of independent interest, since it shows how to modify the standard linear-quadratic optimal control problem to include exogenous learning about parameters.

3.5.1 Parameter Uncertainty without Learning

If the regulator never expects to acquire information about the uncertain parameter g , that parameter is like any other random variable. In this case, we can solve the optimization problem by replacing the damage function in equation (1) with

$$D(S_t) = G_1 (S_t - \bar{S})^2; \quad G_1 \equiv \exp(g_1 + \frac{1}{2}\sigma_{g,1}^2),$$

using equation (3) and $E(\omega) = 1$.

This certainty equivalent¹³ version of the problem with unknown g is identical to the linear-quadratic model studied in a number of previous papers. There is a closed form solution to this problem, given in terms of the solution to a Riccati matrix equation (Karp and Zhang [13]). The assumption that G_1 is convex in the unknown parameter means that greater uncertainty about g (a larger value of $\sigma_{g,1}^2$) increases the expected slope of marginal damages ($= 2G_1$), leading to higher abatement. This change is the intrinsic effect of parameter uncertainty. This control problem satisfies the Principle of Certainty Equivalence (see Section 2), which implies that the expected (stock and flow) trajectories under taxes and quotas are identical.

3.5.2 Passive Learning

The simplest way to model passive learning, and the closest to our model of active learning, is to assume that in each period the regulator receives a signal $\eta_t \sim N(g, \sigma_\omega^2)$. With this assumption,

¹³The certainty equivalent version of a problem with an unknown parameter (here g) replaces that parameter with a known parameter (here, $g_1 + \frac{\sigma_{g,1}^2}{2}$). By solving the certainty equivalent version, we obtain the solution to the original problem. In order to distinguish the certainty equivalent approach of solving the problem with parameter uncertainty and no learning, from the Principle of Certainty Equivalence in the linear-quadratic control problem, we capitalize the latter.

the model from Section 3.2 is unchanged, except that the exogenous random variable η_t replaces the endogenous random variable \hat{g}_t . Learning is “passive” here, because the signal does not depend on anything that the regulator can influence, such as the stock of pollution. Remark 1 continues to hold. (The current belief about e^g is an unbiased estimate of future beliefs.)

The control problems with both active and passive learning have four state variables, $(g_t, \sigma_{g,t}^2, S_t, E^i(\theta_{t+1}|I_{i,t+1}))$. Under passive learning, the state g_t changes exogenously. This fact leads to an important simplification of the solution, summarized in the following proposition.

Proposition 1 *Under passive learning: (i) The value function is quadratic in $w_t \equiv (S_t, E^i_t \theta_t)'$. That is, the value function is of the form $\Psi_0^i + \Psi_1 w + w' \Psi_2 w$, $i = (\text{taxes, quotas})$. The (scalar, vector and matrix) Ψ_j , $j = 0, 1, 2$ are functions of g_t and $\sigma_{g,t}^2$. (ii) The vector and matrix functions Ψ_1 and Ψ_2 are identical under taxes and quotas. This fact implies that the control rules for quotas and for expected emissions under taxes are identical linear functions of w_t under passive learning. (iii) The scalar function Ψ_0^i depends on whether the regulator uses taxes or quotas. (iv) In the special case where $E^i_t \theta_t = 0$ we can obtain the optimal control rule by solving a pair of recursive scalar fixed point problems, that do not depend on ρ or σ_μ^2 .*

(See Appendix 7.2 for the proof.) Proposition 1 tells us something about the outcome under passive learning, and this information gives us a clue about the outcome under active learning. The fact that *expected* trajectories are identical under taxes and quotas (under passive learning) suggests that they should be very similar under active learning – as indeed is the case for our simulations.

The control problem under active learning cannot be solved in closed form, and because the state has four dimensions, the “curse of dimensionality” is a serious issue. Any method that alleviates this numerical problem is useful. Proposition 1.i means that we get “closer to” a closed form solution under passive learning. We still need to approximate the unknown (scalar-vector- and matrix-valued) functions Ψ_j , $j = 0, 1, 2$. However, these functions depend on only two arguments, and the recursive relation that defines them is quite simple; in particular, it does not involve maximization. We can obtain an exact solution for these functions in the limiting (steady state) case where $\sigma_{g,t}^2 = 0$. Therefore, this numerical problem is much simpler than the problem of approximating the four-dimensional value function under active learning. Proposition 1.iv is important because it means that if we are interested in the case where $E^i_t \theta_t = 0$ (a restriction that is reasonable at least for the first period), we can obtain the

control rule by solving a *much* simpler numerical problem involving recursive scalar fixed point problems.

4 Calibration and Model Solution

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.

A characteristic of this model – in addition to the features described in the previous section – is its greater transparency. 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.

Table 1 contains the baseline parameter values. We discuss the main assumptions behind these values here, and provide the details in Appendix 7.3. We then describe the solution method, relegating details to Appendix 7.4.

4.1 Calibration

Perhaps the most controversial issue concerns the relation between carbon stocks and environmental damages. Calibration of the damage function requires three parameters, \bar{S} (the stock at which damages are 0), g , and σ_ω^2 . In addition, we need two state variables, the initial mean and variance g_1 and $\sigma_{g,1}^2$. We set \bar{S} equal to the pre-industrial level of stocks. The choice of the other four variables is less obvious.

Most readers would find it difficult to decide whether a particular value of g should be considered large or small. Therefore, we assume that stock related damages are proportional to Gross World Product (GWP), and we define ϕ as the percentage reduction of GWP due to a doubling of stocks from their pre-industrial level. We state our assumptions about model parameters in terms of ϕ , a parameter about which readers can form an opinion. We use a stationary model, so we treat GWP as constant.¹⁴ Nordhaus [28] surveys opinions of damages

¹⁴Since income will probably grow, this stationarity assumption means that our model understates true damages, if those damages are really proportional to income. Our other stationarity assumption is that abatement costs do

Parameter	Note	Value
β	a continuous discount rate of 3%	0.7408
Δ	an annual decay rate of 0.0083	0.9204
ρ	cost correlation coefficient	0.96
f	constant in the benefits, billion \$	-13089.03
a	intercept of the marginal benefit, \$/(ton of carbon)	224.26
b	slope of the marginal benefit, billion \$/(billion tons of carbon) ²	1.9212
$2e^g$	true slope of the marginal damage, billion \$/(billion tons of carbon) ²	0.0604
σ_μ	standard deviation of cost shock, \$/(ton of carbon)	5.5945
σ_ω^2	variance of ln(damage shock)	0.6349
\bar{S}	zero damage stock, billion tons of carbon	590
State Space:		
S	pollutant stock, billion tons of carbon	(781, 2190)
g	mean of the belief about g	(-6.3029, -2.0544)
	e^g : billion \$/(billion tons of carbon) ²	
σ_g^2	variance of the belief about g	(0, 0.6349)
$E\theta$	expectation about the cost shock, \$/(tons of carbon)	(-39.1615, 39.1615)
Initial State Variables:		
$[S, g, \sigma_g^2, E\theta] = [781, -4.8137, 0.6349, 0]$		

Table 1: Parameter values and region of interest in the state space.

associated with an estimated 3°C warming, a temperature change associated with a doubling of CO_2 stocks. The opinions about ϕ range from 0 to 21 percent of GWP with mean 3.6 and coefficient of variation 1.6 (Table 2 in Roughgarden [34]).

In order to make our model consistent with this survey, we assume that the true value of ϕ is 3.6, and that the coefficient of variation of damages is 1.6; this assumption gives us two pieces of information. We also assume that $\sigma_{g,1}^2$ equals the posterior after one observation, beginning with diffuse priors. This assumption implies that $\sigma_{g,1}^2 = \sigma_{\omega}^2$, so parameter uncertainty and the inherent randomness of damages contribute equally to the coefficient of variation of damages.¹⁵ With these three assumptions, we can assign values to g , σ_{ω}^2 and $\sigma_{g,1}^2$. Finally, we assume that the regulator's initial belief is that $\phi = 1.33$, a value used in previous numerical studies ([17] [20] [29]). This assumption implies a numerical value for the initial mean, g_1 .

Thus, our baseline assumes that the regulator currently underestimates the true level of damages. This case seems to be the most interesting, but we also studied scenarios in which the regulator correctly estimates or over-estimates the true damages.

We assume that the expected BAU level of emissions is constant, in order to use a stationary model. We choose this constant so that our model predicts a BAU level of CO_2 stocks of 1500 GtC in 2100, consistent with the IPCC IS92a scenario (IPCC [10], page 23). We then calibrate a quadratic abatement cost function that approximates (very closely) Nordhaus' [29] formula for expected abatement costs, for levels of abatement ranging from 0 to 75% of BAU emissions.

In our model, the actual BAU level of emissions is a random variable which is linearly related to the cost shock θ . We use 13 observations of historical emissions, at ten-year intervals, to estimate a detrended model of emissions. Using these estimates and the assumed relation between the innovation in emissions and the random variable θ , we obtain values of ρ and σ_{μ}^2 .

We set the length of a period equal to 10 years. This choice means that for a reasonable yearly discount rate, the single period discount factor is relatively small, making it easier to achieve convergence. It also implies that policy levels cannot change frequently.

not change over time.

The state vector in our model includes $\sigma_{g,t}^2$, which is deterministic and monotonic in time. Therefore, we could include a deterministic trend without increasing the dimension of our state vector. This change would require additional assumptions about growth trends, and in our view would reduce the transparency and usefulness of the model, without adding real information.

¹⁵Recall that a smaller value of σ_{ω}^2 implies that learning occurs more quickly. Given a particular level of the coefficient of variation of damages, an increase in the uncertainty about g (increase $\sigma_{g,1}^2$), requires an increase in the speed of learning.

4.2 Solving the model

In order to solve the model, we use neural networks to approximate the value function $J(\cdot)$. We also experimented with alternative functional families including ordinary polynomials and Chebyshev polynomials (Judd [12]). We found that for this problem the neural network approximation has a higher speed of convergence and less severe “curse of dimensionality”.

We divide the state space $[S, g, \sigma_g^2, E\theta]$ into $16 \cdot 12 \cdot 8 \cdot 8 = 12288$ points. The pollutant stock ranges from 781GtC (the current CO₂ concentration) to 2190GtC (the IPCC IS92e scenario). Assuming that ϕ , the percentage loss in GWP from CO₂ doubling, lies between 0.3 and 21 (the minimum non-zero and maximum expert opinions respectively in Nordhaus [28]), we obtain a range for the parameter g . The variance σ_g^2 lies between 0 and 0.6349, the posterior variance after one observation, beginning with a diffuse prior. We bound the cost shock θ using a 95% confidence interval, based on our estimation that uses historical emissions.

We find the optimal control for each grid point through numerical maximization. The state variables in the next period depend on the possible realization of damages D_t and the cost shock μ_t . In solving the dynamic program, we evaluate the expectations in equation (12) using 10 Monte Carlo simulations on random variables D_t and μ_t . The choice of ten simulations is a compromise between approximation time/complexity and approximation accuracy. In taking expectations, we use the regulator’s subjective beliefs – not the objective expectations, which require knowing the true value of g .

5 Results

We discuss the mean trajectories, the effect of learning, and then the comparative statics of the model. Our baseline case assumes that the true value of g is -3.5005 (corresponding to $\phi = 3.6$, the mean estimate in Nordhaus [28]) and that the initial belief about g is $g_1 = -4.8137$ with $\sigma_{g,1}^2 = 0.6349$ (corresponding to $\phi = 1.33$, a value used in previous studies ([17] [20] [29])). Recall that ϕ equals the percentage loss in GWP due to doubling of carbon stocks.

To simulate the outcome, we use this control rule and M draws of sequences of damage and abatement cost shocks; in some cases we use $M = 100$, and then to check accuracy we use $M = 1000$. These random sequences are drawn from the distributions given by equations (1) and (8). In each period, the realization of damages depends on the true value of g . Stocks and beliefs evolve as we described above. We average over the M realizations of trajectories in

order to obtain the mean trajectories for stocks, beliefs, and emissions.

5.1 Trajectory Simulation

Figure 1 shows the expected stock trajectory over 1000 years (100 periods). The right panel shows the trajectory under BAU and under four optimal trajectories corresponding to *known* values of ϕ . When ϕ is known, the Principle of Certainty Equivalence holds, so the expectation of the emission stock trajectories is the same under taxes and quotas. We solve the problem with known ϕ by solving a standard linear-quadratic control problem.

The high damage estimate ($\phi = 21$) causes optimal steady state stocks to reach only 39.91% of the BAU level. Steady state stocks are about 330.6 GtC (or 19.6%) lower when $\phi = 3.6$ rather than 1.33. Thus, our baseline case implies that in the absence of learning, the belief that $\phi = 1.3$ when it actually equals 3.6 would have serious consequences.

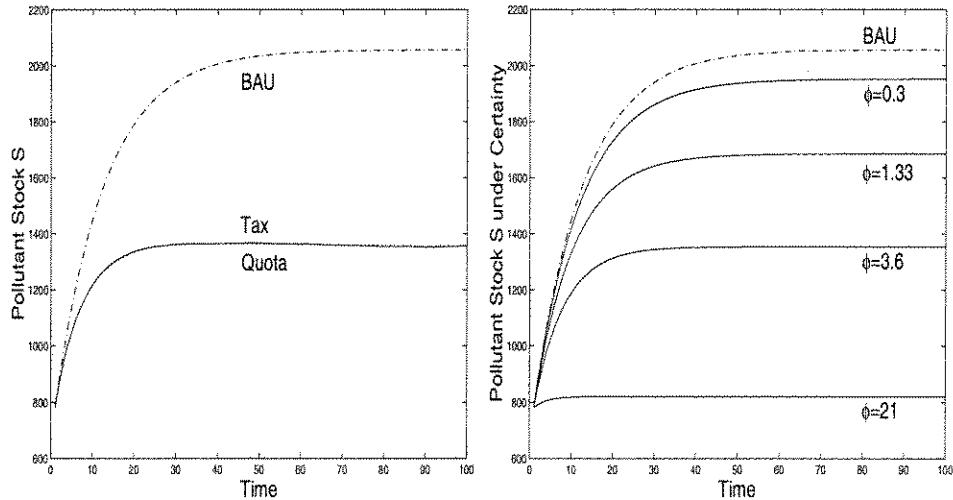


Figure 1: Pollutant stock over time.

The left panel of Figure 1 shows the expected stock trajectory when the Bayesian regulator begins by believing that $\phi = 1.33$, and the true value is 3.6. The expected trajectories under taxes and quotas are not exactly the same, but their difference is not visible at the scale used in Figure 1. Both the tax and quota trajectories converge to the same steady state as under certainty. The fact that the two policies have nearly the same expected trajectories was explained in Section 3.5. The fact that they converge to the steady state under certainty is not surprising,

since the subjective variance on g approaches 0.

We also compared the expected trajectory when the regulator is certain that $\phi = 3.6$ (i.e. the optimal trajectory under complete parameter certainty), with the expected trajectory when the regulator begins by thinking that $\phi = 1.33$ and its true value is $\phi = 3.6$. These two trajectories never differ by more than a couple of percent.

The two important results are: (i) the choice of taxes or quotas leads to essentially the same expected stock trajectory, and (ii) the expected stock trajectory when the Bayesian regulator begins by being much too optimistic about damages, remains close to the optimal trajectory when the regulator knows the true relation between stocks and damages.

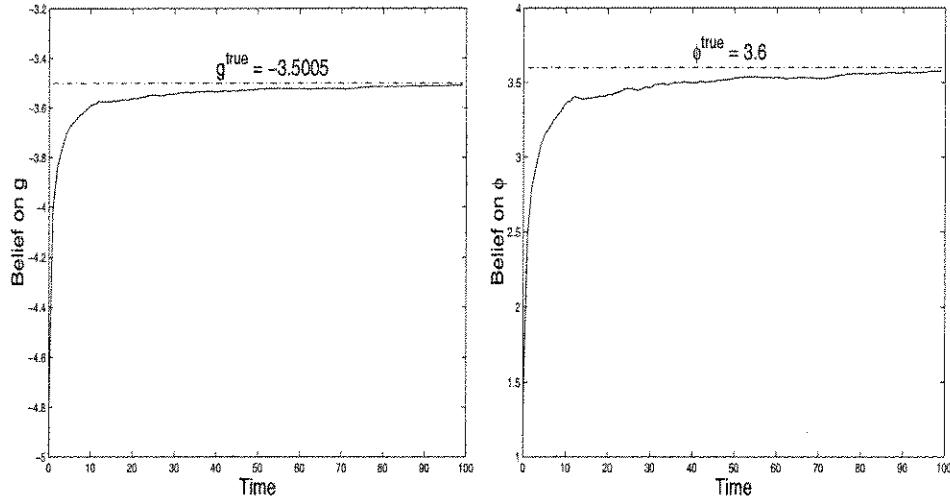


Figure 2: Evolution of belief on the unknown damage parameter over time.

The left panel of Figure 2 shows the expected trajectory for the subjective mean of g and the right panel shows the trajectory for the corresponding values of ϕ . The graphs under taxes and quotas are nearly identical, so we show only one graph. This similarity of mean beliefs is not surprising, since new information depends on the stocks, and the trajectory of stocks is similar under taxes and quotas. The subjective mean converges rather slowly to the true value. It takes five observations (50 years) for the subjective mean to travel approximately 80% of the distance between the initial mean and the true value, and 21 observations (210 years) to travel 95% of this distance.¹⁶

¹⁶Kelly and Kolstad [17] used a different criteria for convergence of beliefs: the expected amount of time it

As we noted above, the expected stock trajectory with learning remains close to the optimal trajectory under certainty. Thus, even though it appears that learning occurs slowly, it nevertheless occurs quickly enough to keep stocks from straying far from the optimal level. The explanation for this is simply that stocks also change quite slowly. For the first 50-70 years, the optimally regulated stocks are similar when the regulator is certain that $\phi = 1.33$ or certain that $\phi = 3.6$, although those trajectories subsequently diverge (see the right panel of Figure 1). Even though learning is slow, a substantial amount has occurred before the stock moves far from its optimal level under certainty.

We also conducted a number of simulations to compare the variability (rather than the mean) of emissions under taxes and quotas. Emissions under taxes are always more variable, because they respond to the actual value of the cost shock. An increase in ρ causes a substantial increase in the variability of emissions under taxes. The information set is much more variable when the regulator learns about the cost shock and these shocks are highly serially correlated.

5.2 The effect of anticipated learning

We emphasized the distinction between the intrinsic effect of parameter uncertainty (i.e., uncertainty – but no learning – about the stochastic relation between stocks and damages), and the effect of anticipated learning. Section 3.5 explains how this model can distinguish between the two. Here we begin with an assessment of the magnitude of parameter uncertainty. We then illustrate the intrinsic effect of parameter uncertainty, and the effect of anticipated learning.

Recall that in calibrating the model, we assumed that the coefficient of variation of damages (associated with doubling of greenhouse gas stock from the pre-industrial level) is 1.6, and that the initial estimate of the loss resulting from this doubling is $\phi = 1.33$ percent of GWP. We also assumed that parameter uncertainty and the random damage shock contribute equally to the uncertainty about damages, implying an initial value of $\sigma_{g,1}^2 = .64$.

There are several ways to assess whether our calibration implies a large or a small amount of parameter uncertainty. We can consider the range of possible values of ϕ – either before reaching the steady state or in the steady state; we can obtain an approximate 95% confidence interval for ϕ ; or we can consider the effect of $\sigma_{g,1}^2$ on the certainty equivalent value of ϕ .

Given our state space for g and σ_g^2 , the range of ϕ is [.3, 21]. If we take into account the would take the regulator to reject the hypothesis that the unknown parameter equals the initial prior. Using this criteria, they also found that convergence to the true belief was slow.

fact that in the steady state $\sigma_g^2 = 0$, the steady state range of ϕ is reduced to [.2, 15.3]. An approximate 95% confidence interval for ϕ is [.27, 6.6].¹⁷ The certainty equivalent damage coefficient is $G = \exp(g + \frac{1}{2}\sigma_g^2)$. We can calculate the value of g that corresponds to $\phi = 1.33$ under parameter certainty ($\sigma_g^2 = 0$). Using this same value of g , but now letting $\sigma_g^2 = 0.64$ (our baseline estimate) implies that $\phi = 1.83$, an increase of about 38%.

These calculations give somewhat different impressions regarding the magnitude of uncertainty about ϕ implicit in our calibration. However, they all suggest that uncertainty about ϕ is not negligible. We now consider whether this uncertainty is potentially important. Columns 3-6 of Table 2 show the optimal level of emissions and abatement in the first (ten-year) period under parameter certainty, for four (known) values of ϕ . This table and Figure 1 show that both the optimal level of first period emissions and the stock trajectories are quite sensitive to ϕ . This observation, and the fact that our model includes non-negligible uncertainty about ϕ , means that anticipated learning about ϕ is *potentially* important.

We emphasized in Section 3.5.1 that we can model parameter uncertainty without learning by solving a certainty equivalent problem, i.e. by changing the value of G_1 (which is linearly related to ϕ). Since optimal emissions are sensitive to ϕ , parameter uncertainty (without learning) is important. For example, increasing ϕ from 1.33 to 1.83 (a 38% increase) raises abatement from 9.8% to 13.11% (a 34% increase).

The last two columns of Table 2 show the emissions and abatement levels (in the first ten year period) under active and passive learning. The entries in these two columns should be compared to the entries in the column labelled $\phi = 1.33$. (Our baseline calibration used the assumption that the initial belief is $\phi = 1.33$.) This comparison shows that the anticipation of learning decreases abatement, but by a small amount.¹⁸

¹⁷We obtain this approximate confidence interval by taking two standard deviations around the initial point estimate of g , using the fact that G_1 is proportional to the initial expectation of ϕ (see equation (18) in Appendix 7.3). We set the damage shock ω equal to its expected value.

¹⁸The algorithm that we used to solve the problem under active learning uses neural networks, and is quite complex. Using Proposition 1, we were able to solve the passive learning problem using a simpler algorithm, described in more detail in Appendix 7.2. When we use different numerical methods to solve the two problems, the emissions levels are not directly comparable.

In order to obtain comparable numbers, we solved the passive learning model with the (complex) algorithm used for active learning. In this case, we find that the first period emissions under active and passive learning are virtually identical. Thus, the difference in the emissions levels reported in the last two columns of Table 2 is probably due to the difference in numerical methods, and not to the difference between active and passive learning. The fact that the two methods of solving the passive learning model lead to such similar control levels, inspires

	<i>BAU</i>	$\phi = 0.3$	$\phi = 1.33$	$\phi = 3.6$	$\phi = 21$	active learning	passive learning
emissions	116.73	113.96	105.26	89.4	26.06	105.8	105.54
abatement (% of BAU)		2.35%	9.8%	23.4%	77.7%	9.3%	9.6%
tax (\$/tC)		5.32	22.03	52.50	174.19	20.99	21.50

Table 2: First Period Emission, Abatement Level, and Tax

Under both active and passive learning, the current belief about the damage parameter is an unbiased estimate of future beliefs (Remark 1). This feature might appear to suggest that the regulator acts as if beliefs will not change – a conjecture that would explain the insensitivity of the optimal control to the amount of learning. This explanation is not correct, as it would imply that the certainty equivalent (no learning) model and the model of passive learning lead to exactly the same control rule. By comparing the functional equations that determine the control rules in the two cases (see Appendix 7.2), we can verify that the control rules are not identical. Anticipated learning does change optimal behavior – but not by very much.

The last row of Table 2 shows the tax level that supports the optimal (expected) level of emissions. Decanio [4] describes recent attempts to estimate the level of carbon tax that would be needed to achieve reductions in emissions. The Energy Modeling Forum at Stanford University estimated that a tax of between \$50 and \$260 per metric ton – with an average of \$170 – would induce a 20% emissions reduction, relative to 1990 levels. The Interagency Analytic Team of the US government estimated that a tax of between \$89 and \$160 per metric ton would stabilize emissions at 1990 levels, by the year 2010. Our point estimate of \$52.5 for a 23.4% reduction is lower than those estimates, but is of a similar order of magnitude.

5.3 Ranking taxes and quotas

We explained that the expected stock and emission trajectories are identical under taxes and quotas if the damage parameter is known with certainty, if it is unknown and there is no learning, or if it is unknown and there is passive learning. Even in these cases, the expected payoffs

confidence in our numerical results.

differ under the two policies, as explained in Section 2. Numerical experiments show how active learning affects the comparison of payoffs.

We compared the sensitivity of the policy ranking by varying one of the parameters or initial conditions, holding the others constant. In all cases, taxes dominate quotas, always by a fraction of one percent of 1998 GWP. For the baseline parameters, the difference between the value function under taxes and under quotas is approximately 0.13% of 1998 GWP. The largest difference in value functions (0.4% of 1998 GWP) occurred when we set the cost variance equal to its maximum allowable value. Thus, with a yearly discount rate of 0.03, our baseline estimate of the annual savings resulting from the use of taxes rather than quotas is $(0.13) \cdot 0.03 = 0.0039\%$ of 1998 GWP, or approximately \$113 million (1998 dollars). Our highest estimate of the annual savings from using taxes rather than quotas is about \$340 million.

Several previous studies also found a slight preference for taxes rather than quotas. In cases where there is a large welfare difference under the two policies, the model does not satisfy the Principle of Certainty Equivalence (Hoel and Karp [8] and Pizer [33]). Thus, a large difference in welfare under the two policies seems to require that the endogenous targets also be different under the two policies. To the extent that policy-makers even consider the choice between taxes and quotas, they probably want to hold the target level of emissions fixed.

Under active learning, a higher stock level decreases the preference for taxes. (Under passive learning or parameter uncertainty without learning, the Principle of Certainty Equivalence implies that the welfare comparison is independent of the stock level.) As the marginal environmental damage increases with the pollutant stock, it becomes more important to control emissions exactly (as occurs under a quota) rather than choosing only the mean of emissions (as occurs under a tax). The magnitude of the effect is small: doubling the stock causes only a 0.3% reduction in the difference in expected payoffs under taxes and quotas.

A higher expected value of g (g_1), corresponding to steeper expected marginal damages, favors the use of a quota. Section 2 explained this result. An increase in $\sigma_{g,1}^2$ favors the use of taxes. This result is surprising, because we know that a higher value of $\sigma_{g,1}^2$ has the same effect as a higher value of g_1 in the certainty equivalent version of the model. A possible explanation is that greater uncertainty about g makes marginal damages uncertain under either taxes or quotas, eroding the feature (described in Section 2) that favors the use of quotas.

Higher *objective* randomness of damages (higher σ_w^2) – unlike higher subjective uncertainty (higher $\sigma_{g,1}^2$) – favors quotas. An increase in objective randomness increases the probability of

a very bad damage shock. The increased danger of a bad cost shock, associated with higher value of σ_w^2 , makes it more important to be able to control the stock level. Since quotas enable the regulator to control stocks more precisely (relative to taxes), a higher value of σ_w^2 favors quotas.

As noted above, the informational difference between taxes and (non-tradable) quotas is an important reason that the emissions trajectories differ under the two policies. A larger variance in the cost shock (σ_μ^2) magnifies this informational difference. The payoff difference is much more sensitive to changes in σ_μ^2 , compared to changes in other parameters.

We also investigated the sensitivity of the policy ranking with respect to other parameter values. The comparative statics are the same as described in Section 2, so we do not repeat them here.

6 Conclusions

The high degree of uncertainty about the relation between environmental damages and stocks of greenhouse gasses is central to the debate concerning the optimal level and method of greenhouse gas abatement. The fact that we anticipate learning about this relation complicates the decision. This anticipation can be used as an excuse to delay action, in order to avoid unnecessary sacrifices, or as a reason to make additional efforts, in case we learn that the situation is more serious than we expected.

We used a simple model in order to study this dilemma. The model neglects many complexities of the science of global warming. However, it captures, in a nearly transparent manner, beliefs about probable orders of magnitude concerning abatement costs, environmental damages, and levels of uncertainty. In addition, the model allows for nearly catastrophic damages; it restricts damages to be positive regardless of the magnitude of uncertainty; and it implies that both the variance and the mean of damages increase with the stock. The model enables us to identify the effect of anticipated learning, as distinct from parameter uncertainty, and also to distinguish between active and passive learning. We obtained three main conclusions that have a direct bearing on the global warming debate.

The most important conclusion is that although anticipated learning (either active or passive) leads to slightly lower abatement (higher emissions), the effect is extremely small. Some policy discussions have emphasized that the anticipation of “better science” should influence our

current decisions. Our results suggest that the importance of this issue has been exaggerated. Environmentalists who favor current abatement, or those who oppose the sacrifices needed to achieve this abatement, should base their positions on their beliefs about the expected relative magnitudes of environmental damages and abatement costs. These considerations have an important effect on the optimal level of abatement, but the possibility of more precise knowledge in the future has a very small effect.

Since we assumed that damages are convex in the unknown parameter, uncertainty about that parameter (in the absence of learning) increases the optimal level of abatement. The amount of parameter uncertainty in our calibration is consistent with a 34% increase in first period abatement levels, suggesting that anticipated learning could be significant. The models under both active and passive learning are *qualitatively* different than the model with parameter uncertainty and no learning. However, for the parameters that are relevant for global warming, the *quantitative* difference in the model outcomes is small. In future work we intend to identify the region of parameter space for which the degree and type of learning is quantitatively important.

Our calibration implies that we will learn slowly about the true relation between stocks and damages. However, these stocks will also change slowly. Our second policy conclusion is that even if our current beliefs about global warming are too optimistic, we may be able to learn quickly enough to keep the stock level close to the full-information optimal level. This conclusion is based on a model that assumes a continuous relation between stocks and damages. If there is an unknown threshold level of stocks above which damages are truly catastrophic, our model is not appropriate. The conclusion might also be true but irrelevant, in the absence of political will to base abatement decisions on current science.

Third, for all “reasonable” parameter values, taxes are a better instrument than quotas. In our setting, much of this superiority is due to the fact that taxes provide more information about abatement costs than do quotas. Tradeable quotas eliminate this informational advantage. Previous work suggests that taxes would nevertheless remain superior to quotas. However, the magnitude of the payoff difference is small, when the target level of emissions does not depend on the policy tool. In other words, the policy tool used to achieve a particular target level of emissions may not matter much, although of course the target level is important.

7 Appendix: Technical Information

This appendix summarizes information about equations of motion and the expectation of abatement costs under taxes and quotas. We give proofs for Remark 1 and Proposition 1. We then provide the details of the model calibration and of the method of solution of the numerical problem.

7.1 Summary of equations of motion and expected benefits

Here we summarize the material from Section 3.3. The subjective mean satisfies the following equation of motion:

$$E^{Tax}(\theta_{t+1}|\theta_t) = E^{Tax}(\rho\theta_t + \mu_{t+1}|\theta_t)$$

$$= \rho\theta_t = \begin{cases} \rho\theta_0 = \rho\bar{\theta}_0 + \rho\mu_0, & t = 0 \\ \rho(\rho\theta_{t-1} + \mu_t) = \rho E^{Tax}(\theta_t|\theta_{t-1}) + \rho\mu_t, & t \geq 1 \end{cases}$$

under taxes, and

$$E^{Quota}(\theta_{t+1}|\bar{\theta}_0) = \rho E^{Quota}(\theta_t|\bar{\theta}_0)$$

under quotas. Under taxes, $\mu_0 \equiv \theta_0 - \bar{\theta}_0$ has mean 0 and variance σ_0^2 . μ_t ($t \geq 1$), as defined in equation (8), has mean 0 and variance σ_μ^2 .

The firm's emission response as a function of z_t is:

$$x_t^*(z_t, \theta_t) = z_t + \frac{1}{b} [\theta_t - E^{Tax}(\theta_t|\theta_{t-1})] = z_t + \frac{\mu_t}{b}.$$

The regulator's expectation of the firm's cost saving in period t is:

$$E^{Tax} \{ B[x_t^*(z_t, \theta_t), \theta_t] | \theta_{t-1} \}$$

$$= f + [a + E^{Tax}(\theta_t|\theta_{t-1})] z_t - \frac{b}{2} z_t^2 + \frac{Var^{Tax}(\theta_t|\theta_{t-1})}{2b}$$

under taxes, and

$$E^{Quota} [B(x_t, \theta_t) | \bar{\theta}_0] = f + [a + E^{Quota}(\theta_t|\bar{\theta}_0)] x_t - \frac{b}{2} x_t^2$$

under quotas. Under taxes, the firm adjusts emissions to the realized cost shock, increasing the expected cost saving by $\frac{Var^{Tax}(\theta_t|\theta_{t-1})}{2b}$ which equals $\frac{\sigma_0^2}{2b}$ when $t = 0$ and $\frac{\sigma_\mu^2}{2b}$ when $t \geq 1$.

1. However, taxes make the next-period pollutant stock stochastic, increasing expected future environmental damages (since the damage function is convex in the pollutant stock, by Jensen's Inequality).

7.2 Proofs

Proof. (Remark 1) Note that this proof also holds in the case of passive learning. We only need to show that $E_t G_{t+1} = G_t$. The Remark then follows from the rule of iterated expectations.

We find the mean and variance of the normally distributed posterior estimator g_{t+1} under active learning. Given the prior belief on g , the expected moment estimator (before observing current damages) under active learning is the prior:

$$\begin{aligned} E_t \hat{g}_t &= E_t \left\{ \ln \frac{D_t}{(S_t - \bar{S})^2} + \frac{\sigma_\omega^2}{2} \right\} = E_t \left\{ g + \ln \omega_t + \frac{\sigma_\omega^2}{2} \right\} \\ &= g_t - \frac{\sigma_\omega^2}{2} + \frac{\sigma_\omega^2}{2} = g_t. \end{aligned}$$

Thus the conditional mean of the posterior estimator g_{t+1} equals the prior:

$$E_t (g_{t+1}) = \frac{\sigma_\omega^2}{\sigma_\omega^2 + \sigma_{g,t}^2} g_t + \frac{\sigma_{g,t}^2}{\sigma_\omega^2 + \sigma_{g,t}^2} E_t (\hat{g}_t) = g_t.$$

The conditional variance of the posterior estimator g_{t+1} is

$$\begin{aligned} Var_t (g_{t+1}) &= \frac{\sigma_{g,t}^4}{(\sigma_\omega^2 + \sigma_{g,t}^2)^2} Var_t (\hat{g}_t) = \frac{\sigma_{g,t}^4}{(\sigma_\omega^2 + \sigma_{g,t}^2)^2} Var_t \left(g + \ln \omega_t + \frac{\sigma_\omega^2}{2} \right) \\ &= \frac{\sigma_{g,t}^4}{(\sigma_\omega^2 + \sigma_{g,t}^2)^2} (\sigma_\omega^2 + \sigma_{g,t}^2) = \frac{\sigma_{g,t}^4}{\sigma_\omega^2 + \sigma_{g,t}^2}. \end{aligned}$$

G_{t+1} is lognormally distributed with mean

$$\begin{aligned} E_t G_{t+1} &= E_t \exp \left(g_{t+1} + \frac{\sigma_{g,t+1}^2}{2} \right) = E_t \exp (g_{t+1}) \cdot \exp \left(\frac{\sigma_{g,t+1}^2}{2} \right) \\ &= \exp \left[E_t (g_{t+1}) + \frac{1}{2} Var_t (g_{t+1}) \right] \cdot \exp \left[\frac{\sigma_{g,t}^2 \sigma_\omega^2}{2 (\sigma_\omega^2 + \sigma_{g,t}^2)} \right] \\ &= \exp \left[g_t + \frac{\sigma_{g,t}^4}{2 (\sigma_\omega^2 + \sigma_{g,t}^2)} + \frac{\sigma_{g,t}^2 \sigma_\omega^2}{2 (\sigma_\omega^2 + \sigma_{g,t}^2)} \right] \\ &= \exp \left(g_t + \frac{\sigma_{g,t}^2}{2} \right) = G_t. \end{aligned}$$

■ **Proof.** (Proposition 1) To simplify notation, it is convenient to write the dynamic problem in terms of the states G_t and $\sigma_{g,t}^2$ rather than g_t and $\sigma_{g,t}^2$. Begin with a guess that the value function is linear-quadratic in $w_t \equiv (S_t, E_t^i \theta_t)'$:

$$\Psi_{0,t}^i + \underbrace{(v_{1,t} \ v_{2,t})}_{\Psi_1(G_t, \sigma_{g,t}^2)} w_t + \frac{1}{2} w_t' \underbrace{\begin{pmatrix} \psi_{11,t} & \psi_{12,t} \\ \psi_{12,t} & \psi_{22,t} \end{pmatrix}}_{\Psi_2(G_t, \sigma_{g,t}^2)} w_t.$$

We use time subscripts on the unknown functions to indicate the time subscript of arguments of the function. For example, $\psi_{11,t} \equiv \psi_{11}(G_t, \sigma_t^2)$.

Using the necessary condition (13), under passive learning the optimal control rule is linear in S_t and $E_t^i \theta_t$, and is identical under taxes and quotas if and only if the functions Ψ_2 and Ψ_1 are the same under taxes and quotas:

$$\begin{aligned} z_t^* = x_t^* &= \frac{[a + \beta E_t(v_{1,t+1})] + [1 + \beta \rho E_t(\psi_{12,t+1})] E_t^i \theta_t + \beta \Delta E_t(\psi_{11,t+1}) S_t}{b - \beta E_t(\psi_{11,t+1})} \quad (14) \\ &= \Phi_0 + \Phi_1 w_t. \end{aligned}$$

(Under active learning, the value function is not linear-quadratic in $w_t \equiv (S_t, E_t^i \theta_t)'$ and the optimal control is not linear in w_t , because of the dependence of G_{t+1} on S_t . Under parameter uncertainty without learning, the functions $v_{1,t+1}$, $\psi_{11,t+1}$ and $\psi_{12,t+1}$ are non-random constants. Under passive learning these functions depend on G_{t+1} , which is a random variable at time t .)

Substituting the optimal control under passive learning back into the DPE and equating the value function coefficients leads to a recursive system of nonlinear matrix equations:

$$\begin{aligned} \Psi_2(G_t, \sigma_{g,t}^2) &= F_2(G_t, E_t \Psi_2(G_{t+1}, \sigma_{g,t+1}^2)), \\ \Psi_1(G_t, \sigma_{g,t}^2) &= F_1(E_t \Psi_2(G_{t+1}, \sigma_{g,t+1}^2), E_t \Psi_1(G_{t+1}, \sigma_{g,t+1}^2)), \\ \Psi_0^i(G_t, \sigma_{g,t}^2) &= F_0(E_t \Psi_2(G_{t+1}, \sigma_{g,t+1}^2), E_t \Psi_1(G_{t+1}, \sigma_{g,t+1}^2), E_t \Psi_0^i(G_{t+1}, \sigma_{g,t+1}^2)). \end{aligned}$$

(The explicit system of matrix equations is cumbersome, so we do not produce it in its entirety here, but it is available on request.) For functions Ψ_j^i that satisfy these equations, the quadratic value function satisfies the DPE and the transversality condition, thus establishing Part (i) of the Proposition. This system of equations is recursive, and the equations for Ψ_2 and Ψ_1 are identical under taxes and quotas. The control rule (14) is therefore the same under taxes and quotas, establishing Part (ii). The equation for Ψ_0^i differs under taxes and quotas because of the

difference in the dynamic equation for $E_t^i \theta_t$ under taxes and quotas; this fact establishes Part (iii).

When $E_t^i \theta_t = 0$ the control rule (14) does not depend the function ψ_{12} , but it does depend on the functions ψ_{11} and v_1 . In order to establish Part (iv) we need the formulae that determine these two function. These equations are:

$$\psi_{11}(G_t, \sigma_{g,t}^2) = \frac{b\beta\Delta^2 E_t \psi_{11}(G_{t+1}, \sigma_{g,t+1}^2)}{b - \beta E_t \psi_{11}(G_{t+1}, \sigma_{g,t+1}^2)} - 2G_t, \quad (15)$$

$$v_1(G_t, \sigma_{g,t}^2) = \frac{\beta\Delta [bE_t v_1(G_{t+1}, \sigma_{g,t+1}^2) + aE_t \psi_{11}(G_{t+1}, \sigma_{g,t+1}^2)]}{b - \beta E_t \psi_{11}(G_{t+1}, \sigma_{g,t+1}^2)} + 2G_t \bar{S}. \quad (16)$$

These two functions do not depend on the functions ψ_{12}, ψ_{22} or v_2 or on the parameters ρ or σ_μ^2 . Equations (15) and (16) can be solved recursively. Thus, we can find the optimal control rule by solving two scalar fixed point problems. These remarks establish Part (iv). ■

Discussion. We solve the fixed point problems in equations (15) and (16) using the collocation method, described in Miranda and Fackler [24]. We use the following procedures from the toolbox that accompanies their book: FUNDEFN, FUNFITXY, FUNEVAL.

In the limiting case where there is no parameter uncertainty ($\sigma_g^2 = 0$), $G_{t+1} \equiv G_t$. In this case we have the standard linear-quadratic control problem. We merely remove the expectations operator in equations (15) and (16). For this limiting case we obtain a closed form (but complicated) expression for the control rule.

7.3 Model Calibration

Our discount factor for a ten-year period, $\beta = 0.7408$, implies an annual discount rate of 3% ([17] [20] [29]). Both costs and damages are measured in billions of 1998 US dollars.

CO₂ emissions and stock. The CO₂ atmospheric stock S_t is measured in billions of tons of carbon equivalent (GtC). The pre-industrial atmospheric stock is about 590GtC as estimated by Neftel et al. [25] and used in Kelly and Kolstad [17] and Pizer [32]. We take this level to be the steady state stock given a low level of economic activity. Let e_t be total anthropogenic CO₂ emissions in period t . Approximately 64% of these emissions contribute directly to the atmospheric stock ([20] [29]). Remaining emissions are absorbed by oceanic uptake, other terrestrial sinks, and the carbon cycle (IPCC [10]). The linear approximation of the evolution

of atmospheric stocks is

$$S_t - 590 = \Delta (S_{t-1} - 590) + 0.64e_t.$$

We take $x_t \equiv 0.64e_t$, the anthropogenic fluxes of CO_2 into the atmosphere, as the control variable and rewrite the above equation as

$$S_t = \Delta S_{t-1} + (1 - \Delta) 590 + x_t. \quad (17)$$

The estimate of the stock persistence is $\Delta = 0.9204$ (an annual decay rate of 0.0083 and a half-life of 83 years). ([17] [20] [29])

Environmental damage. There is a simple relation between ϕ , defined as the percentage reduction in Gross World Product (GWP) due to a doubling of CO_2 stocks, and the parameters of our model. In Nordhaus's survey [28] the expert opinions on ϕ range from 0 to 21 percent of GWP with mean 3.6 and coefficient of variation 1.6. Our calibration is consistent with these expert opinions.

The 1998 estimate of GWP is 29,185 billion dollars (International Monetary Fund [11]), for a 10 year estimate of GWP of 291,850. The estimated damages due to doubling of CO_2 stocks during this period is $291,850 \frac{\phi}{100}$. Equating this value to the expected damages given by equation (3) gives us one calibration equation:

$$291,850 \phi \frac{1}{100} = \exp(g_1 + \frac{1}{2} \sigma_{g,1}^2) (590)^2 \implies .0083841\phi = \exp(g_1 + \frac{1}{2} \sigma_{g,1}^2) = G_1. \quad (18)$$

(We have set the time index $t = 1$.) Equation (18) implies that if the true value of ϕ is 3.6 (and the regulator knows this, so that $\sigma_g^2 = 0$) then the true slope of marginal damages is $2(.0083841\phi) = 2e^g = 6.0366 \times 10^{-2}$.

We obtain our second calibration equation using the coefficient of variation of damages in Nordhaus' survey and equation (5):

$$CV(\phi) = 1.6 = [\exp(\sigma_{1,t}^2 + \sigma_\omega^2) - 1]^{\frac{1}{2}} \Rightarrow 3.56 = \exp(\sigma_{g,1}^2 + \sigma_\omega^2). \quad (19)$$

We need one more assumption to identify the model parameters. We assume that the regulator begins with diffuse priors ($\sigma_{g,0}^2 = \infty$) and has made one observation, so his posterior variance (using equation (7) is $\sigma_{g,1}^2 = \sigma_\omega^2$). Using this equation, we can solve equation (19) to obtain $\sigma_{g,1}^2 = \sigma_\omega^2 = .63488$. Using this value we can rewrite equation (18) as $g_1 = -.31744 + \ln(8.3841 \times 10^{-3}\phi)$. Thus, the value of g_1 corresponding to the belief that $\phi = 1.33$ and the level of uncertainty $\sigma_{g,1}^2 = .63488$ is $g_1 = -.31744 + \ln(8.3841 \times 10^{-3}(1.33)) = -4.8137$.

Abatement cost. Uncontrolled emissions are expected to rise over time, leading to more than doubling of carbon stocks. In order to retain a stationary model, we need to assume that expected BAU emissions, \bar{x} , are constant. Given the current atmospheric CO₂ concentration $S_0 = 781\text{GtC}$ (Keeling et al. [16]), using equation (17) the expected future BAU atmospheric CO₂ concentration is

$$S_t = \Delta^t S_0 + \frac{1 - \Delta^t}{1 - \Delta} [(1 - \Delta) 590 + \bar{x}].$$

We choose $\bar{x} = 116.73$ GtC so that our model is consistent with the IPCC IS92a scenario that projects CO₂ stocks at 1500 GtC in 2100 (IPCC [10], page 23).

We calibrate the abatement cost function as a quadratic approximation to Nordhaus' [29] formula, $A = 0.0686u^{2.887} \times 291,850$, where u is the fractional reduction in CO₂ emissions and A is the abatement cost. We draw 1000 realizations of u from a uniform distribution with support [0, 0.75] (the same support that Nordhaus [27] used) and calculate A using this formula. Each value of u implies a value of abatement, $\bar{x} - x = u\bar{x}$, with $\bar{x} = 116.73$. We regress A against $(u\bar{x})^2$ to obtain a quadratic function for abatement costs:

$$A(x_t) = 0.9606 (116.73 - x_t)^2.$$

The R^2 for this regression is 0.9762, implying that the quadratic function and the function in Nordhaus's formula are very similar, for reductions between 0 and 75% of emissions. The benefit function (the negative of abatement costs), including the additive cost shock is,

$$B(x_t, \theta_t) = -13089.03 + (224.26 + \theta_t) x_t - 0.9606 x_t^2$$

giving an estimated slope of marginal benefits of $b = 1.9212$ billion \$/GtC².

Cost correlation and uncertainty. In our model, the cost uncertainty is linearly related to the BAU level of emissions. We used data on actual emissions to estimate the variance and autocorrelation of the cost shock. Using maximum likelihood method on data from Marland et. al. [22] (total global carbon emissions over every 10 years during the period 1867-1996) we estimated the following model:

$$e_t = e_0 + \kappa t + \varepsilon_t, \quad \varepsilon_t = \rho \varepsilon_{t-1} + \nu_t, \quad \nu_t \sim iid N(0, \sigma_\nu^2).$$

(Since we have only 13 observations, we view this procedure as merely a means of calibration.) The estimates are $\rho = 0.96$ and $\sigma_\nu = 4.55$ GtC. We convert the emission uncertainty σ_ν into

cost uncertainty σ_μ by multiplying it by 0.64 (because $x_t \equiv 0.64e_t$), and then by the slope of marginal abatement cost $b = 1.9212$ (because $\theta_t \equiv b\tilde{\theta}_t$). The result is $\sigma_\mu = 4.55 \times 0.64 \times 1.9212 = 5.5945\$/\text{ton of carbon}$. For simplicity, we assume $\sigma_0 = \sigma_\mu$. (The regulator's subjective prior on θ_0 has the same variance as the subsequent cost shocks.)

7.4 Solving the DPE by Value Function Iterations using Neural Networks

Taylor and Uhlig [35] review a variety of methods to solve the DPE (12). The method of linear-quadratic approximations (Christiano [2], McGrattan [23]) is not applicable here because the environment in this problem is far from the steady state. We obtain an approximation of $J(\cdot)$ by value function iterations using a specific family of functions $\Phi(\cdot; \mathcal{B})$ with \mathcal{B} being the parameter vector. The objective is to find the parameter vector \mathcal{B}^* so that the approximated value function $\hat{J}(\cdot) \equiv \Phi(\cdot; \mathcal{B}^*)$ comes close to satisfying the DPE (12). We use neural networks to achieve this approximation. We experimented with alternative functional families including ordinary polynomials and Chebyshev polynomials (Judd [12]). We found that for this problem the neural network approximation has a higher speed of convergence and less severe “curse of dimensionality”.

We divide the state space $[S, g, \sigma_g^2, E\theta]$ into $16 \cdot 12 \cdot 8 \cdot 8 = 12288$ points. The range of the pollutant stock is from 781GtC (the current CO₂ concentration) to 2190GtC (the IPCC IS92e scenario). Assuming that ϕ is between 0.3 and 21 (the minimum non-zero and maximum expert opinions respectively in Nordhaus [28]), we have the parameter g between -6.3029 and -2.0544 . The variance σ_g^2 lies between 0 and 0.6349, the posterior variance after a diffuse prior. We bound the estimated cost shock $E\theta$ by $\pm \frac{1.96\sigma_\mu}{\sqrt{1-\rho^2}}$ which is $(-39.1615, 39.1615)$, the 95% confidence interval for the cost shock in the long-run.

We denote the vector of state variables $(S, g, \sigma_g^2, E\theta)'$ at grid point n by \vec{X}_n and the set of all grids by \vec{X} . Suppose there are $n = 1, \dots, N$ grid points in state space. With a specific family of functions $\Phi(\cdot; \mathcal{B})$, the method of value function iterations begins with an initial guess of the value $\hat{J}^{i-1}(\vec{X}_n)$ at each grid point. We then find the parameter vector \mathcal{B}^i that minimizes the sum of squared residuals over the set of grid points in the state space:

$$\mathcal{B}^i = \underset{\mathcal{B}}{\operatorname{argmin}} \sum_{n=1}^N \left\{ \left[\Phi(\vec{X}_n; \mathcal{B}) - \hat{J}^{i-1}(\vec{X}_n) \right]^2 \right\}. \quad (20)$$

Corresponding to the DPE (12), we denote the current expected payoff at each grid point by $F(\vec{X}_n; z_n)$ and growth equations of state variables by $G(\vec{X}_n, z_n; \mu_n, D_n)$. z_n is the optimal

control, which needs to be determined, for the specific grid point n . μ_n is the realized cost shock, and D_n is the realized environmental damages; both at grid point n . We find the optimal control z_n by solving

$$\underset{z_n}{\text{Max}} \left\{ F \left(\vec{X}_n; z_n \right) + \beta E_t \Phi \left[G \left(\vec{X}_n, z_n; \mu_n, D_n \right); \mathcal{B}^i \right] \right\} = \hat{J}^i \left(\vec{X}_n \right), \quad (21)$$

and get a new approximation to the optimal value, $\hat{J}^i \left(\vec{X}_n \right)$, at each grid point n .

Given a grid point (initial vector of state variables) \vec{X}_n , the vector of state variables in the next period is $G \left(\vec{X}_n, z_n; \mu_n, D_n \right)$; it depends on the optimal control z_n and realizations of the cost shock μ_n and the damage D_n . To evaluate the expectations in equation (21), we use 10 Monte Carlo simulations of the random variables μ_t and D_t .

We find the parameter vector \mathcal{B}^* by iterating steps (20) and (21) until the approximated value function converges:

$$\left\| \hat{J}^i \left(\vec{X} \right) - \hat{J}^{i-1} \left(\vec{X} \right) \right\| < 10^{-8}.$$

Figure 3 depicts a single hidden layer feedforward neural network. The input units correspond to the state variables at each grid point, $(x_1, x_2, x_3, x_4) \equiv (S, g, \sigma_g^2, E\theta) \equiv \vec{X}_n$. The output unit corresponds to the approximated value $\hat{J} \left(\vec{X}_n \right)$. Instead of a direct relation between the inputs and the output, the neural network assumes that there exists one layer of hidden units L_j ($j = 1, \dots, m$) between the input units and the output unit.

Each of the hidden units L_j ($j = 1, \dots, m$) receives a signal \tilde{L}_j that is the weighted sum of all inputs x_i ($i = 1, \dots, 4$), $\tilde{L}_j = w_{0j} + \sum_{i=1}^4 w_{ij} x_i$, and sends out a signal $L_j = \mathcal{H} \left(\tilde{L}_j \right)$. \mathcal{H} is a transfer function. Similarly, the output unit receives a signal \tilde{J} that is the weighted sum of signals from the hidden units, $\tilde{J} = r_0 + \sum_{j=1}^m r_j L_j$, and sends out a signal $\hat{J} = \mathcal{H} \left(\tilde{J} \right)$. The network is feedforward because signals flow in only one direction. Such a neural network mapping from inputs \vec{X}_n to the output $\hat{J} \left(\vec{X}_n \right)$ can be written as

$$\hat{J} \left(\vec{X}_n \right) = \mathcal{H} \left\{ r_0 + \sum_{j=1}^m r_j \left[\mathcal{H} \left(w_{0j} + \sum_{i=1}^4 w_{ij} x_i \right) \right] \right\} \equiv \Phi \left(\vec{X}_n; \mathcal{B} \right). \quad (22)$$

$\mathcal{B} = (w_{ij}, r_j : i = 0, \dots, 4; j = 0, \dots, m)$ is the vector of parameters in the neural network.

We use the logistic function $\mathcal{H}(\lambda) = \frac{1}{1+e^{-\lambda}}$, $\lambda \in \mathbb{R}$, as the transfer function.¹⁹ Here for each hidden unit, $\lambda = \tilde{L}_j$; and for the output unit, $\lambda = \tilde{J}$. Hornik et al. [9] proves the ability of

¹⁹We adopt the same transfer function for the hidden units and the output unit. The transfer functions for these two levels of units can be different. The usual requirement is that $\mathcal{H} : \mathbb{R} \rightarrow [0, 1]$, nondecreasing, $\lim_{\lambda \rightarrow \infty} \mathcal{H}(\lambda) = 1$, and $\lim_{\lambda \rightarrow -\infty} \mathcal{H}(\lambda) = 0$. (Hornik et al. [9])

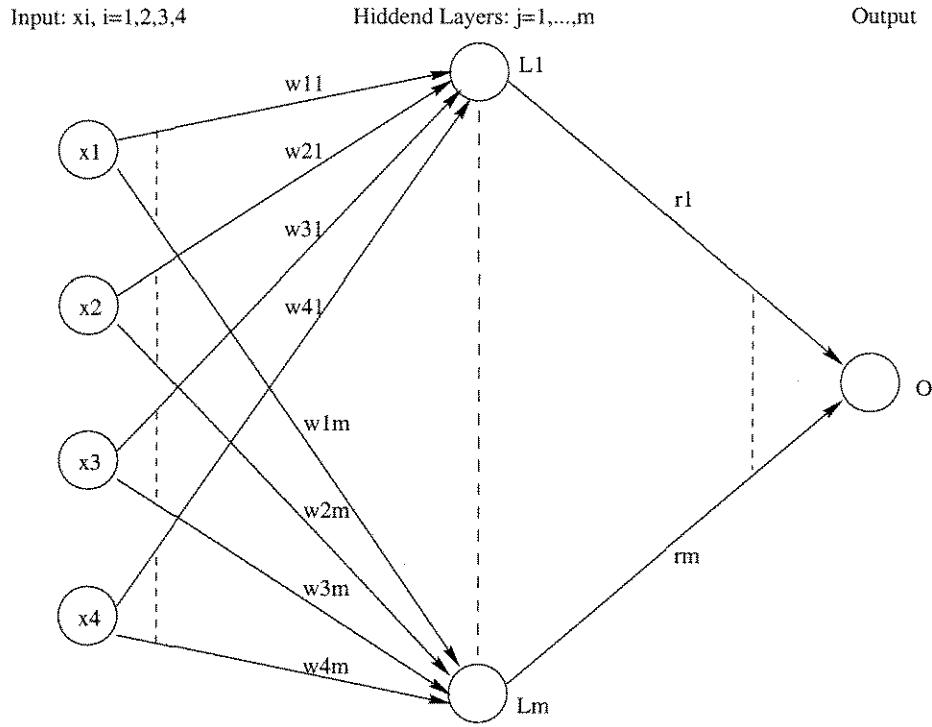


Figure 3: A single hidden layer feedforward neural network.

such a neural network to approximate an unknown mapping arbitrarily well, provided there are sufficient number of hidden units.

A possible way to choose the number of hidden units is cross-validation (White [38]). To balance the computation time/complexity and the approximation accuracy, we take 10 hidden units, resulting in 61 elements in \mathcal{B} .

The algorithm in optimizing (20) is Back-Propagation which implements a local gradient descent (see White [39] for technical details). The algorithm in optimizing (21) is the Quasi-Newton method with a mixed quadratic and cubic line searching. We implement both optimizations using MATLAB built-in routines.

References

[1] G. Chichilnisky and G. Heal. Global environmental risks. *Journal of Economic Perspectives*, 7(4):65–86, Fall 1993.

- [2] L. J. Christiano. Solving the stochastic growth model by linear-quadratic approximation and by value-function iteration. *Journal of Business and Economic Statistics*, 8(1):23–26, Jan. 1990.
- [3] C. Costello and L. Karp. Dynamic quotas with learning. Giannini Foundation of Agricultural Economics Working Paper No. 914, Nov. 2000.
- [4] S. Decanio. The economics of climate change. Redefining Progress, San Francisco, 1997.
- [5] A. Fisher and U. Narain. Global warming, endogenous risk, & irreversibility. Working Paper, Department of Agricultural and Resource Economics, University of California, Berkeley, Apr. 2000.
- [6] W. H. Greene. *Econometric Analysis*. Prentice-Hall, Inc., New Jersey, fourth edition, 2000.
- [7] M. Hoel and L. Karp. Taxes versus quotas for a stock pollutant. Giannini Foundation of Agricultural Economics Working Paper No. 855, Mar. 1999.
- [8] M. Hoel and L. Karp. Taxes and quotas for a stock pollutant with multiplicative uncertainty. *Journal of Public Economics*, forthcoming.
- [9] K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.
- [10] Intergovernmental Panel on Climate Change (IPCC), Cambridge, UK. *Climate Change 1995: The Science of Climate Change*, 1996. J.T. Houghton, L.G. Meira Filho, B.A. Callander, N. Harris, A. Kattenberg, and K. Maskell, editors.
- [11] International Monetary Fund (IMF), Washington, DC. *World Economic Outlook*, Oct. 1999.
- [12] K. L. Judd. Projection methods for solving aggregate growth models. *Journal of Economic Theory*, 58(2):410–452, Dec. 1992.
- [13] L. Karp and J. Zhang. Regulation of stock externalities with learning. Giannini Foundation of Agricultural Economics Working Paper No. 892, Sept. 1999. Revised January, 2001.

- [14] L. Karp and J. Zhang. Regulating global climate change with bayesian learning about damages. Working Paper, Department of Agricultural and Resource Economics, University of California, Berkeley, May 2000.
- [15] L. Karp and J. Zhang. Controlling a stock pollutant under asymmetric information and endogenous investment. Working Paper, Department of Agricultural and Resource Economics, University of California, Berkeley, July 2001.
- [16] C. Keeling and T. Whorf. Atmospheric CO₂ records from sites in the SIO air sampling network. In *Trends: A Compendium of Data on Global Change*. Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, U.S. Department of Energy, Oak Ridge, Tennessee, 1999.
- [17] D. L. Kelly and C. D. Kolstad. Bayesian learning, growth, and pollution. *Journal of Economic Dynamics and Control*, 23(4):491–518, Feb. 1999.
- [18] P. W. Kennedy. Learning about environmental damage: Implications for emissions trading. *Canadian Journal of Economics*, 32(5):1313–1327, Nov. 1999.
- [19] C. D. Kolstad. Fundamental irreversibilities in stock externalities. *Journal of Public Economics*, 60(2):221–233, May 1996.
- [20] C. D. Kolstad. Learning and stock effects in environmental regulation: The case of greenhouse gas emissions. *Journal of Environmental and Economic Management*, 31(1):1–18, July 1996.
- [21] A. S. Manne and R. G. Richels. *Buying Greenhouse Gas Insurance: The Economic Costs of Carbon Dioxide Emission Limits*. MIT Press, Cambridge, Massachusetts, 1992.
- [22] G. Marland, T. Boden, R. Andres, A. Brenkert, and C. Johnston. Global, regional, and national CO₂ emissions. In *Trends: A Compendium of Data on Global Change*. Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, U.S. Department of Energy, Oak Ridge, Tennessee, 1999.
- [23] E. R. McGrattan. Solving the stochastic growth model by linear-quadratic approximation. *Journal of Business and Economic Statistics*, 8(1):41–44, Jan. 1990.

[24] M. J. Miranda and P. L. Fackler. *Applied Computational Economics and Finance*. MIT Press, Cambridge, Massachusetts, forthcoming.

[25] A. Neftel, H. Friedli, E. Moor, H. Lötscher, H. Oeschger, U. Siegenthaler, and B. Stauffer. Historical CO₂ record from the Siple Station ice core. In *Trends: A Compendium of Data on Global Change*. Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, U.S. Department of Energy, Oak Ridge, Tennessee, 1999.

[26] R. Newell and W. Pizer. Stock externality regulation under uncertainty. Resources for the Future Working Paper 99-10, Dec. 1998.

[27] W. D. Nordhaus. To slow or not to slow: The economics of the greenhouse effect. *Economic Journal*, 101(407):920–937, July 1991.

[28] W. D. Nordhaus. Expert opinion on climate change. *American Scientist*, pages 45–51, January–February 1994.

[29] W. D. Nordhaus. *Managing the Global Commons: The Economics of Climate Change*. MIT Press, Cambridge, Massachusetts, 1994.

[30] W. D. Nordhaus and D. Popp. What is the value of scientific knowledge? an application to global warming using the PRICE model. *Energy Journal*, 18(1):1–45, 1997.

[31] S. C. Peck and T. J. Teisberg. Global warming uncertainties and the value of information: An analysis using CETA. *Resource and Energy Economics*, 15(1):71–97, Mar. 1993.

[32] W. A. Pizer. The optimal choice of climate change policy in the presence of uncertainty. *Resource and Energy Economics*, 21(3-4):255–287, Aug. 1999.

[33] W. A. Pizer. Prices vs. quantities revisited: The case of climate change. *Journal of Public Economics*, forthcoming.

[34] T. Roughgarden and S. H. Schneider. Climate change policy: Quantifying uncertainties for damages and optimal carbon taxes. *Energy Policy*, 27(7):415–429, July 1999.

[35] J. B. Taylor and H. Uhlig. Solving nonlinear stochastic growth models: A comparison of alternative solution methods. *Journal of Business and Economic Statistics*, 8(1):1–17, Jan. 1990.

- [36] A. Ulph and D. Ulph. Global warming, irreversibility and learning. *Economic Journal*, 107(442):636–650, May 1997.
- [37] M. L. Weitzman. Prices versus quantities. *Review of Economic Studies*, 41(4):477–491, Oct. 1974.
- [38] H. White. Connectionist nonparametric regression: Multilayer feedforward networks can learn arbitrary mappings. *Neural Networks*, 3(5):535–550, 1990.
- [39] H. White. *Artificial Neural Networks: Approximation and Learning Theory*. Blackwell Publishers, Cambridge, Massachusetts, 1992. with A. Ronald Gallant, Kurt Hornik, Maxwell B. Stinchcombe, and Jeffrey M. Wooldridge.