AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

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.*

*Department of Agricultural &
Resource Economics, UCB*
CUDARE Working Papers
(University of California, Berkeley)

Year 2000

Paper 870

Taxes and Quotas for a Stock Pollutant
with Multiplicative Uncertainty

Michael Hoel *

Larry Karp †

*University of Oslo

†University of California, Berkeley and Giannini Foundation

This paper is posted at the eScholarship Repository, University of California.

http://repositories.cdlib.org/are_uch/870

Copyright ©2000 by the authors.

Taxes and Quotas for a Stock Pollutant with Multiplicative Uncertainty*

Michael Hoel[†] and Larry Karp[‡]

Abstract

We compare taxes and quotas when firms and the regulator have asymmetric information about the slope of firms' abatement costs. Damages are caused by a stock pollutant. We calibrate the model using cost and damage estimates of greenhouse gasses. Taxes dominate quotas, as with additive uncertainty. This model with multiplicative uncertainty allows us to compare expected stock levels under the two policies, and to investigate the importance of stock size and the magnitude of uncertainty on the policy ranking.

Key Words: Pollution control, asymmetric information, taxes and quotas, stochastic control, global warming, multiplicative disturbances

JEL Classification Numbers: H21, Q28

---

*Giannini Foundation Working Paper 870. An earlier version of this paper was distributed as Nota di Lavoro 15.99, Fondazione Eni Enrico Mattei. We thank three referees for their comments, but claim responsibility for any remaining errors.

[†]Univ.of Oslo, PO Box 1095 Blindern, N-0317 OSLO Norway; email: michael.hoel@econ.uio.no

[‡]Corresponding Author. Department of Agriculture and Resource Economics, 207 Giannini Hall, Berkeley CA 94720; email: karp@are.berkeley.edu

1 Introduction

Most important pollutants persist in the environment for a non-negligible period. For these pollutants, the environmental damage at a point in time depends primarily on the existing stock, rather than the flow. The current flow changes future stocks and therefore affects future damages, but in many cases the current flow does not contribute directly to current damages. For example, the danger of global warming (possibly) depends on stocks of greenhouse gasses. The flow of these gasses may also be associated with health costs, but their major social cost arises through the increased future stock of greenhouse gasses.

When the firms that produce emissions have better information about abatement costs than do regulators, control of pollution becomes more difficult. The regulator needs to choose both the type of policy (e.g., a tax or quota) and the level of the policy. Asymmetry of information can occur regardless of whether environmental damages are associated with flows or stocks. A static model can be used to represent the problem of a regulator facing a flow pollutant, but a stock pollutant requires a dynamic model. Most of the theoretical literature on pollution control under asymmetric information has focused on flow pollutants. Here we contribute to the literature on the control of stock pollutants under asymmetric information about abatement costs.

Many ingenious ways of controlling pollution under asymmetric information have been devised using principal-agent models. In practice, most regulations rely on either quantity restrictions (quotas) or taxes. An important literature, beginning with Weitzman (1974), compares social welfare under these two policies. Subsequent contributions that study the

static model include Malcomson (1978), Roberts and Spence (1976), Stavins (1996), Watson and Ridker (1984) and Yohe (1977).

Recently, several papers compare taxes and quotas under asymmetric information for a stock pollutant. Staring (1995) studies the simplest open-loop model. Hoel and Karp (1998) emphasize the importance of the flexibility of firms and regulators, using both an open-loop and feedback model. Newell and Pizer (1998) and Karp and Zhang (1999) include correlated shocks, as discussed in Section 2.4. All these papers conclude that (i) a steeper marginal environmental damage curve, or a flatter marginal abatement cost curve favor the use of quotas, and (ii) a higher discount factor or a lower decay rate – factors which make stocks more important – favor the use of quotas. The first conclusion echoes Weitzman’s main result: a higher ratio of the slope of marginal damages relative to the slope of marginal abatement costs favor quotas. The second result is specific to the dynamic setting. All of these dynamic models assume that damages and abatement costs are quadratic functions, and that the equation of motion for the stock is linear. These functional assumptions have three advantages: (i) The static models also use quadratic approximations, so the static and dynamic results can be compared in a straightforward manner. (ii) The linear-quadratic assumption leads to analytic results. (iii) The models can be easily calibrated.

The previous dynamic papers assume that the regulator is imperfectly informed about the *intercept* of the marginal abatement costs; uncertainty is additive. This assumption – which was also used in most of the static models – leads to clear analytic results. It also has three important implications. It implies that (i) the variance of the random term has no

effect on the optimal policy choice or on the level of the policy; (ii) the ranking of policies is independent of the level of the existing stock of pollution, and (iii) the expectations of the optimal trajectories under the two policies are identical.

Here we assume that the regulator and firms have asymmetric information about the *slope* of marginal abatement costs. Uncertainty is multiplicative rather than additive. There is no reason to think that one form of uncertainty is more realistic than the other. The previous emphasis on additive uncertainty was due to its mathematical convenience, not its realism. Our model allows us to check whether, in a specific application (i.e., for specific parameter values) the policy ranking is robust to the type of uncertainty.

The model with multiplicative uncertainty also sheds light on three questions: How does uncertainty affect the optimal policy choice and the optimal policy level? Does a higher stock of the pollutant favor the use of a particular policy? How does the choice of policy affect the level of the expected stock trajectory? With additive uncertainty, the answers to these questions (given above) are clear but vacuous, in the sense that they follow trivially from the mathematical assumptions. Those answers are also implausible; it seems that the optimal tax or quota might depend on the magnitude of the uncertainty, that the policy choice might depend on the level of existing stock, and that different policies might lead to different stock trajectories. All of these possibilities arise with multiplicative uncertainty.

With multiplicative uncertainty the answers to these (and other) questions are ambiguous. However, given specific parameter values we can *easily* answer the three questions, and also compare the conclusions about policy ranking under additive and multiplicative

uncertainty. The numerical results, together with the closed form solutions, provide a basis for intuition about the problem.

We use previous estimates of the magnitudes of abatement costs and environmental damages to calibrate a linear-quadratic model of global warming. The linear-quadratic model cannot capture the enormous complexity of the problem, but it has the virtue of *parsimony and simplicity*.¹ Differences of opinion about fundamental issues, such as whether abatement costs and environmental damages are large or small, are captured in two or three parameters.

Under a range of parameter values that are consistent with the published range of opinion about the magnitude of damages, we find that taxes are preferred to quotas under both additive and multiplicative uncertainty. However, even though the two types of uncertainty lead to the same policy ranking, they lead to different conclusions about the importance of making the right choice. We also find that the expected stock trajectory is higher under optimal taxes compared to quotas. A larger stock favors the use of quotas. A higher variance of the cost shock decreases the expected stock trajectory.

2 The model

We first describe the elements of the model: the cost of abatement function, the damage function, and the equation of motion. We then present the solution to the optimization

---

¹ Our formulation assumes a direct relation between carbon stocks and environmental damages. Most models assume that carbon stocks increase future global temperature and that this increase in temperature causes environmental damages. In this case there is a lag between changes in stocks and changes in damages. The linear-quadratic model can incorporate this modification by including an additional state variable. Footnote 9 discusses how this change would affect our results.

problem. The third subsection provides the analysis and intuition for the comparison of taxes and quotas. The final subsection explains the reason for assuming that the random cost shocks are uncorrelated.

2.1 Elements of the model

Pollution emissions and costs are defined as flows, e.g. billions of tons per unit of time or billions of dollars per unit of time. Variables are constant within a period and each period lasts for h units of time. We can think of h as being the amount of time between the arrival of new information, or the amount of time during which decisions are held fixed. Thus, h can be viewed as a measure of the flexibility of decision-makers. For smaller values of h , information arrives and policies are changed more frequently. The comparison between taxes and quotas depends on many parameters, including the length of each period.²

The random variable $\frac{1}{\theta}$ enters the cost function multiplicatively. When the (representative) firm's actual flow of pollution is x , the cost of abatement is $-(f + ax - \frac{b}{2\theta}x^2)$, where f, a and b are parameters.³ In the absence of regulation, the firm minimizes costs by emitting the random "business-as-usual" flow $x^* \equiv \frac{a\theta}{b}$.

At the beginning of each period the firm, but not the regulator, observes the current value of θ . If the regulator chooses a tax p per unit of pollution, the firm minimizes the

---

² The reader should mentally set $h = 1$ in this paper. Hoel and Karp (1998) discuss the parameter h extensively. Since it plays the same role in the model with multiplicative disturbances, we do not repeat the discussion here. It is important to retain h in our formulae so that these can be used by other researchers.

³ In his reply to Malcomson (1978), Weitzman (1978) models multiplicative disturbances by having the random variable divide rather than multiply the slope of marginal costs. We adopt this formulation because it leads to a slight simplification in the derivations. Of course the two formulations are equivalent, since one is obtained from the other merely by re-defining the random variable. An (equivalent) alternative, replacing $\frac{bx^2}{\theta}$ with $(b + \theta)x^2$, can also be obtained by redefining variables.

sum of abatement cost and tax payments.⁴ The first order condition to the firm's cost minimization problem implies that the flow of pollution is

$$x = \left(\frac{a-p}{b} \right) \theta = z\theta, \quad (1)$$

which uses the definition $z \equiv \frac{a-p}{b}$. We can think of the regulator as choosing the variable z rather than the tax p . Under a tax, the flow of pollution in any period is a random variable.

We assume that θ is independently and identically distributed – an assumption we discuss below. The first and second moments of θ are $E\theta = \bar{\theta}$ and $E\theta^2 = \gamma$, so $\text{var}(\theta) = \sigma^2 = \gamma - \bar{\theta}^2$. The regulator knows the probability distribution of θ . If we normalize by setting $\bar{\theta} = 1$, then $\text{var}(\theta) = \gamma - 1$. With this normalization, z is the expected flow of pollution when the regulator uses the tax $p = a - bz$. We provide the formulae for general values of $\bar{\theta}$ and then specialize to $\bar{\theta} = 1$.

If the regulator chooses a quota, x , we assume that it is binding. With a quota, the flow of pollution in a period is deterministic. We assume that the quota is allocated efficiently, e.g. by means of tradable permits.

The flow of damages resulting from the stock of pollution, S , is $cS + \frac{g}{2}S^2$, where c and g are parameters. The firm ignores these costs, but the regulator cares about them. The flow of payoff (the negative of abatement costs minus damages) for the regulator is

---

⁴ We interpret θ as a sector-wide shock, or as the shock to a “representative firm”. This firm solves a sequence of static problems because it behaves non-strategically with respect to the regulator and takes the future trajectory of taxes or quotas as given. If the firm made investment decisions which affect future abatement costs (and therefore affect future policies) the firm would have a dynamic problem. In this case, we need to solve a dynamic game in order to determine the equilibrium policy (Karp and Zhang 2000).

$$\left(f + ax - \frac{b}{2\theta}x^2\right) - \left(cS + \frac{g}{2}S^2\right). \quad (2)$$

Using (1) and (2), the regulator's expected flow of payoff in a period in a tax regime (denoted by T) is

$$\lambda(t; T) \equiv \lambda(z_t, S_t; T) = f + \left(az_t - \frac{b}{2}z_t^2\right)\bar{\theta} - \left(cS_t + \frac{g}{2}S_t^2\right). \quad (3)$$

Under a quota regime (denoted by Q), the regulator's expected flow of payoff in a period is

$$\lambda(t; Q) \equiv \lambda(x_t, S_t; Q) = f + ax_t - \frac{b}{2}x_t^2 E\left(\frac{1}{\theta}\right) - \left(cS_t + \frac{g}{2}S_t^2\right). \quad (4)$$

Since each period lasts for h units of time, and since all variables are constant within a period, the regulator's expected payoff for a period is $\lambda(t; i)h$, for $i = T, Q$ (taxes, quotas).

The equation of motion for the stock of pollutant, S , is

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

where the fraction of stock that persists until the next period is $\Delta = e^{-\delta h}$; δ is the continuous time decay rate.

With a discount factor $\beta = e^{-rh}$ and an initial value of the stock S_0 , the maximized present discounted value of the regulator's payoff (i.e., the value function) is

$$J(S_0; i) = \max E \sum_{t=0}^{\infty} \beta^t \lambda(t; i) h$$

subject to (5), for $i = T, Q$.

2.2 Solution to the optimization problems

Under quotas, the regulator has a standard deterministic linear-quadratic problem. The term $E\left(\frac{1}{\theta}\right)$ enters the payoff as a constant which multiplies b (see 4); uncertainty has no other effect on the problem. Under taxes, the regulator has a stochastic control problem with multiplicative disturbances. In this case, the “Principle of Certainty Equivalence”⁵ does not apply: the state contingent optimal control rule depends on both the mean and the variance of θ . However, the control rule is still linear in the stock, and the value function is still quadratic, (as is the case for both the deterministic problem and the problem with additive uncertainty). Thus, we can use standard methods to solve the control problem under taxes. We first solve this problem, and then obtain the solution to the problem under quotas.

With taxes, we take the control variable to be z . Once we know the optimal value of z we use equation (1) to obtain the tax, $p = a - bz$. In the appendix we show that under taxes the optimal value of z is linear in the pollution stock

$$z = -\frac{(a + \beta(\rho_1 + \rho_2\Delta S))\bar{\theta}}{\beta\rho_2 h\gamma - b\bar{\theta}} \quad (6)$$

and the value function is quadratic in S : $J(S; T) = \rho_0 + \rho_1 S + \frac{\rho_2}{2} S^2$. The parameters of this value function satisfy

$$\rho_2 = -gh + \beta\rho_2\Delta^2 - \frac{(\beta\rho_2\Delta)^2 h\bar{\theta}^2}{\beta\rho_2 h\gamma - b\bar{\theta}} \quad (7)$$

---

⁵ The Principle of Certainty Equivalence (which holds only under restrictive conditions) says that the optimal policy is unchanged if random variables are replaced by their expected values. See Bertsekas (1995).

$$\rho_1 = -ch + \beta\rho_1\Delta - \frac{(a + \beta\rho_1)\rho_2\Delta h\bar{\theta}^2}{\beta\rho_2h\gamma - b\bar{\theta}} \quad (8)$$

$$\rho_0 = fh + \beta\rho_0 - \frac{(a + \beta\rho_1)^2 h\bar{\theta}^2}{2(\beta\rho_2h\gamma - b\bar{\theta})}. \quad (9)$$

Given that the value function is bounded above (for $\beta < 1$) and that it is quadratic, it must be the case that $\rho_2 \leq 0$. Equation (7) is quadratic in ρ_2 . We obtain the (unique) negative root of this equation, and then solve the linear equations (8) and (9) recursively. Given the values of ρ_i we use equation (6) to obtain the optimal expected level of emissions under taxes.

We obtain the value function under quotas, $J(S; Q) = \rho_0 + \rho_1 S + \frac{\rho_2}{2} S^2$, and the optimal quota rule by specializing the solution given above for taxes. Using the normalization $\bar{\theta} = 1$, we simply replace $\bar{\theta}$ and γ by 1 (so that $var(\theta) = 0$) and replace b by $bE\left(\frac{1}{\theta}\right)$ in the control rule and in the equations for ρ_i . The optimal quota rule and the equations that define the parameters of the value function under quotas are

$$x = -\frac{(a + \beta(\rho_1 + \rho_2\Delta S))}{\beta\rho_2h - bE\left(\frac{1}{\theta}\right)} \quad (10)$$

$$\rho_2 = -gh + \beta\rho_2\Delta^2 - \frac{(\beta\rho_2\Delta)^2 h}{\beta\rho_2h - bE\left(\frac{1}{\theta}\right)} \quad (11)$$

$$\rho_1 = -ch + \beta\rho_1\Delta - \frac{(a + \beta\rho_1)\rho_2\Delta h}{\beta\rho_2h - bE\left(\frac{1}{\theta}\right)} \quad (12)$$

$$\rho_0 = fh + \beta\rho_0 - \frac{(a + \beta\rho_1)^2 h}{2(\beta\rho_2h - bE\left(\frac{1}{\theta}\right))}. \quad (13)$$

The value functions under taxes and quotas are both quadratic in the state, but the coefficients ρ_i , $i = 0, 1, 2$, are different under the two policies. For the remainder of this paper, we set $h = 1$.

2.3 Analysis and Intuition

The characteristics of the static and dynamic models with additive uncertainty provide a basis for intuition. In both of those models, a larger value of b favors taxes and a larger value of g favors quotas. The advantage of taxes – the fact that under taxes abatement is negatively correlated with costs – increases when costs are more sensitive to output, i.e. when b is larger. The advantage of quotas – the fact that under quotas the *stock* is nonstochastic – increases when the damage function is more convex, i.e. when g is larger. (This fact is a straightforward application of Jensen’s inequality.) To paraphrase: considerations of stock effects (damages) favor the use of quotas, and consideration of flow effects (abatement costs) favor the use of taxes.

Under quotas (but not under taxes) the regulator in the current period can choose the exact level of emissions, and thus choose the exact level of stock in the next period. The value of this ability increases with β and Δ in the dynamic problem. A larger value of β increases the importance of the future damages associated with current emissions, and a larger value of Δ increases the effect of current emissions on future stock and thus on future damages. Therefore, increases in these two parameters favor the use of quotas ((Hoel and Karp 1998),(Newell and Pizer 1998)).

The complexity of the dynamic model with multiplicative uncertainty precludes general analytic results. However, we can consider limiting cases by studying the unique negative root of equations (7) and (11) for limiting parameter values. For example, as $g \rightarrow 0$ then $\rho_2 \rightarrow 0$ (under both policies). In this case it is easy to show that taxes are preferred to quotas

regardless of the other parameter values. If $b = 0$ the firm's decision rule, equation (1), is not lower semi-continuous (it is a correspondence for $p = a$). However, for b arbitrarily close to 0, the firm's decision rule is well-defined and we can compare the payoffs by considering their limiting values. For $x = z$ and $b = 0$ the single period payoffs under taxes and quotas are equal. (See equations (3) and (4).) However, the evolution of the stock is stochastic under taxes and remains deterministic under quotas. Thus for $b = 0$ the only difference between the two policies is that quotas enable the regulator to exactly control the evolution of the state, whereas taxes enable the regulator to choose only the mean of the evolution of the state. In this case, we expect the payoff to be higher under quotas. In view of the continuity of all the functions that define ρ_i , we conclude that taxes dominate quotas for sufficiently small g , and we expect that quotas dominate taxes for sufficiently small b . For the simulations reported in the next section, an increase in $\frac{b}{g}$ favors the use of taxes.

The most important difference between the dynamic models with additive and with multiplicative uncertainty is that the Principle of Certainty Equivalence holds in the former, but not in the latter. This Principle implies that (with additive uncertainty) the expectations of the optimal trajectories are identical under taxes and quotas and are independent of the variance, and that the ranking of policies does not depend on the stock size. These conclusions are an artifact of the assumption of additive uncertainty, and are not particularly plausible. None of these conclusions hold with multiplicative uncertainty. Inspection of the optimal tax and quota rules given above shows that these depend on the variance. Similarly, the parameters ρ_2 and ρ_1 under taxes and quotas depend on the variance, and the form of this

dependence is different under taxes and quotas. (See equations (7), (8), (11) and (12)). We can show, for example, that the value of ρ_2 is different under taxes and quotas. Therefore, the expectations of the stock and flow trajectories are different under taxes and quotas, and the difference in payoffs, $J(S; T) - J(S; Q)$, depends on S .

Our simulation results (described in the next section) lead to three qualitative conclusions. Although we cannot prove these results analytically, we attempt to provide some intuition for them here. The three qualitative results are:

Result (i) A larger initial stock favors the use of quotas.

Result (ii) A larger variance in the cost shock reduces the (expected) flow of emissions and steady state stock under both policies.

Result (iii) For long-lived stocks (i.e., for Δ close to 1) (a) the effect of a larger variance in the cost shock is more pronounced under quotas, and (b) the use of taxes leads to a higher steady state stock of pollution, compared to quotas.

Result (i) is intuitive. The marginal damage increases linearly with the stock. Marginal damages are high when the stock is large. With large stocks, it is important to reduce the stock. Under quotas, the evolution of the stock is deterministic, so the regulator can reduce it with certainty. Under taxes, on the other hand, the regulator can influence only the expected trajectory of the stock. For large stocks, it is important to obtain the certain reduction, and therefore quotas are preferred.⁶

---

⁶ We solve the regulator's problem assuming that at the initial time she has to decide to use either taxes or quotas forever. Therefore, it is possible that the choice of instruments is time-inconsistent. For example, given the constraint that she sticks with a policy instrument, if the initial stocks are low she might choose to use taxes. If the stock grows along the equilibrium trajectory, and if the regulator were able to revise her choice of instruments at some time in the future, she might want to switch to quotas. If this occurs, her

The fact that the difference in payoffs depends on stock size under multiplicative but not under additive uncertainty has an important practical implication. For a range of parameter values and a range of stock size, the policy ranking may be the same for both types of uncertainty. In this case, the choice of policies does not depend on the form of uncertainty. However, as the stock varies within this range, the difference in the magnitude of the preference can change dramatically. In this case, the importance of making the right choice depends on the form of uncertainty and on the stock size. Our simulations illustrate this possibility.

Result (ii) is also straightforward. It is not surprising that greater uncertainty makes the regulator more cautious, which in this setting means that the (expected) emissions and (expected) trajectories are lower. Uncertainty affects the payoffs under both taxes and quotas. With quotas, greater uncertainty has a *direct* effect on the current period payoff (see equation (4)): it decreases the expected marginal value of emissions in the current period, and therefore decreases the optimal quota.⁷ With taxes, the variance has no direct effect on the payoff in the current period (equation (3)). However, a larger variance of θ increases the variance of S_{t+1} (which equals $z^2 \text{var}(\theta)$, using equations (1) and (5)) and

---

initial choice of instruments is time-inconsistent. However, in our simulations, this time-consistency never arises. The threshold stock size (at which the choice of policy instruments changes) is far above any of the stock levels ever reached in equilibrium.

We could solve a more complicated problem in which the regulator has the option of switching from one policy to another. This option has no value under the parameters in our simulations. More generally, if the regulator would want to switch policies once, the option to do so would increase the preference for the policy that would have been chosen in the absence of the option.

⁷ Since $\frac{1}{\theta}$ is a convex function, an increase in the variance of θ increases $E\frac{1}{\theta}$ and decreases the expected marginal value of an addition unit of emissions, $a - bx E\frac{1}{\theta}$.

therefore decreases future payoffs.⁸ The variance of S_{t+1} can be reduced by lowering expected emissions, z . Thus, for both quotas and taxes, a larger variance of θ creates an incentive to lower (expected) emissions.

Result (iii) consists of two parts. However, part (b) is an immediate consequence of Result (ii) and Result (iia). When the variance is 0, taxes and quotas are equivalent. Therefore, if it is true that an increase in the variance always decreases expected emissions, and if it is also true that the effect is more pronounced under quotas, then expected emissions are higher under taxes than under quotas at a given stock. Consequently, the steady state value of the stock of pollution must be higher under taxes.

We pointed out above that the variance has a direct effect on the current payoff under quotas, whereas under taxes the variance affects future payoffs, via the randomness of the state. In this sense, the effect of uncertainty under taxes is indirect, i.e. it is “filtered” through the state variable S . It is plausible that the variance has a greater affect where its influence is direct (under quotas) rather than indirect (under taxes). The meaning of “indirect” here is related to the longevity of the stock, i.e. the magnitude of Δ . When $\Delta \approx 1$, the flow x is small relative to the stock, S (at least in the neighborhood of the steady state). At an expected steady state S where expected emissions under taxes are $z = (1 - \Delta) S$, the variance of the stock in a period is $z^2 \text{var}(\theta) = (1 - \Delta)^2 S^2 \text{var}(\theta)$. For a given $\text{var}(\theta)$, the variance of the stock goes to 0 as Δ approaches 1. For large Δ , a change in the current tax has a small effect on the variance of S_{t+1} , so the “indirect effect” is small.

---

⁸ Since $\rho_2 < 0$, the present discounted value of future payoffs, $J(S_{t+h}; T)$ is a concave function of the stock and therefore (by Jensen’s inequality) a decreasing function of the *variance* of the stock.

At the other extreme, where $\Delta \approx 0$, the model is very similar to the static model (except for discounting). In that case, the stock in the next period is approximately the same as current emissions and there is no significant distinction between a direct and an indirect effect.

We can improve upon this explanation by considering the differential effect of uncertainty on expected marginal costs and benefits, under taxes and quotas. To achieve simplicity and clarity, we concentrate on explaining Result (iiib) rather than Result (iiia). Denote the optimal deterministic steady state under quotas as S^Q . Imagine that the regulator has been using quotas and is at the steady state, and for some reason begins to use taxes. We want to explain why the regulator would increase the stock when she switches to taxes. At S^Q the expected marginal costs and marginal benefits of emissions, appropriately discounted, are equal under quotas. We need to explain how those marginal costs and benefits change under taxes.

The marginal benefit side is simple. Evaluated at the same expected level of emissions (i.e., $z = x$) the difference in expected marginal benefits of emissions (the reduction in abatement costs) under taxes and quotas is $a - bz - (a - xbE_\theta^1) = xb(E_\theta^1 - 1) > 0$. This difference is increasing in the variance of θ . Under taxes, emissions are high when abatement costs are high. The regulator chooses the current tax without knowing the realization of the current cost shock. A larger cost shock therefore has no effect on the current tax, but it causes the firm to reduce abatement, i.e. to increase emissions. In this sense, emissions, are “arbitraged” across states of nature. The fact that emissions and abatement costs are

positively correlated under taxes is always valuable, and makes the (expected) marginal benefit of emissions higher under taxes than under quotas. Note that the difference in expected marginal benefits is independent of Δ and is non-negligible for positive variance. Thus, the switch from quotas to taxes in our thought experiment increases the expected marginal benefit of emissions by a non-negligible amount.

The marginal cost side is more complicated. The fact that the stock is stochastic under taxes and that damages are convex in the stock, makes the expected marginal cost of emissions higher under taxes. However, at the steady state the randomness of the stock is not “very important” provided that the stock is long-lasting ($\Delta \approx 1$). For a steady state (expected) stock of S , the steady state (expected) flow is $(1 - \Delta) S$. For $\Delta \approx 1$ the flow is small relative to the stock in the steady state. An unexpectedly large flow in one period (caused by a large cost shock) causes a small unexpected change in the stock. That change can be offset by choosing smaller expected flows in subsequent periods. In other words, fluctuations about the steady state are small when $\Delta \approx 1$.⁹

These considerations provide a basis for understanding why the steady state stock of pollution is likely to be higher under taxes than under quotas, for long-lived stocks. Under taxes, both the expected marginal benefits and marginal costs increase. The previous paragraphs explain why the former is likely to change more than the latter. Thus, when

---

⁹ In footnote 2 we mentioned that the model could be modified by assuming that damages are caused by changes in temperature, which respond to changes in carbon stock with a lag. This modification means that cost shocks are “filtered” through two state variables before affecting damages. This delay and smoothing of cost shocks further reduces the disadvantage of taxes, while not affecting their advantage. Consequently, we conjecture that this modification would favor the use of taxes. We thank an anonymous referee for suggesting this modification to our model, and the conjecture.

the regulator switches to taxes she has an incentive to increase emissions, moving the stock to a higher expected steady state.

The above explanation of the characteristics of tax and quota policies is based on simulation results as well on the closed form solutions, so it is speculative. We are able to obtain one additional analytic result by considering the limiting case where $\Delta = 0$, i.e. emissions in the current period cause damages only in the next period. The stock in period $t + 1$ is x_t . The values of ρ_1 and ρ_2 are the same under taxes and quotas, so a comparison of the policies requires only a comparison of the values of ρ_0 . It is straightforward to show that for $\Delta = 0$ taxes dominate quotas if and only if

$$\Phi \equiv \frac{E_{\theta}^1 - 1}{\text{var}(\theta)} > \beta \frac{g}{b}. \quad (14)$$

Setting $\beta = 1$ we have the same criterion for ranking policies as in the static setting.¹⁰

This expression enables us to see clearly the importance of the distribution of the random variable. A larger value of Φ favors the use of taxes; Φ depends on the distribution of the random variable. This dependence also holds in the dynamic setting.

Under the assumption that the distribution of θ is sufficiently regular that we can interchange the order of integration and differentiation, we obtain the second order approximation of Φ for small variance (equation (22) in the appendix). If θ has a symmetric distribution, the first order approximation is $\Phi = 1$. Thus, for very small variance, the policy-ranking criteria under additive and multiplicative uncertainty are virtually the same (as Weitzman

---

¹⁰ As $\beta \rightarrow 1$ the present value of the difference between payoffs under taxes and quotas becomes unbounded. However, for $\beta \leq 1$ equation (14) determines the sign of the difference in payoffs.

(1978) previously pointed out). The second order approximation shows that if the distribution of θ is symmetric or skewed to the left, then $\Phi > 1$. In this case, Φ is an increasing function of the variance (for small variance); an increase in the variance makes taxes more attractive in the static model, and in some cases reverses the policy ranking.¹¹

In the dynamic simulations reported in the next section, where Δ is close to 1, an increase in the variance of the shock favors whichever policy is preferred. For “reasonable” stock size, where a tax is preferred, a higher variance favors taxes. For “extremely large” stocks, where under some circumstances a quota is preferred, a higher variance favors quotas. In our simulations the level of uncertainty does not reverse the policy ranking, but only changes the magnitude of the preference. However, the previous comments imply this reversal can occur in the static model, which is a limiting case of the dynamic model; therefore, there exist parameter values and stock levels such that an increase in uncertainty can reverse the policy ranking in the dynamic model. Under additive uncertainty, on the other hand, we previously showed (Hoel and Karp 1998) that the variance affects neither the policy ranking nor the critical ratio $\frac{g}{b}$.

2.4 The assumption of i.i.d. shocks

We assume that the shocks to cost are serially uncorrelated. Here we briefly discuss an alternative assumption, and explain our reason for adopting the simpler formulation.

---

¹¹ However, if the distribution is skewed to the right, Φ is first decreasing and then increasing in the variance. In this case, a small increase in the variance (beginning with an initial variance close to 0) favors quotas. This ambiguity in the relatively simple static setting illustrates the difficulty of obtaining a complete set of intuition for the much more complicated dynamic model.

If the cost shocks follow first-order serial correlation, then $\theta_t = \alpha\theta_{t-1} + \omega_t$, where α is a parameter and ω is i.i.d. with mean 0. For example, if $\alpha = 1$, cost shocks follow a random walk. Our formulation above takes $\alpha = 0$. For $\alpha \neq 0$ we need to add another state variable to the regulator's control problem. With taxes, the regulator can infer θ_{t-1} after observing x_{t-1} , using equation (1). The value of θ_{t-1} affects the optimal choice of the tax at time t and must therefore be included in the regulator's information set.

Using a model with additive errors, Newell and Pizer (1998) study the role of serial correlation in an open-loop equilibrium (i.e., no learning) under the assumption that the regulator's priors equal their stationary level. Karp and Zhang (1999) study the role of serially correlated shocks in the open-loop equilibrium with arbitrary priors, and in the feedback model (i.e., with learning). If firms are not permitted to trade quotas, or if they have identical shocks so that they have no incentive to trade, then the regulator learns nothing about their cost shocks even in the feedback equilibrium (under the assumption that quotas are binding with probability 1). However, taxes enable the regulator to learn the cost shock with a one-period lag. This informational advantage increases the preference for taxes, and may swamp other considerations. If firms are heterogeneous and are able to trade quotas, this informational advantage vanishes.

The objective of the modelling exercise determines the importance of the assumption $\alpha = 0$. If we wanted to model cost dynamics and regulatory learning, our model with $\alpha = 0$ would be inadequate. However, we are not studying cost dynamics and regulatory learning about costs. Our focus is the effect of persistent asymmetric information with a

stock externality; we want to depart from previous work by studying a model that does not presuppose that the expectations of trajectories are independent of the policy choice and of the variance of the cost shock.

In the real world, cost functions are not stationary. In our model, privately observed shocks are the only source of changes in cost functions. These two facts have lead some readers to conclude that realism requires that our model have serially correlated cost shocks. This conclusion conflates two distinct issues: cost dynamics and private information. Consider a more complicated model with known deterministic technical change and/or serially correlated *publicly observed* cost shocks. In such a model, with cost dynamics, realism does not require that the firm's *private information* is also serially correlated. Since more generality is usually better than less, the model with $\alpha \gtrless 0$ is better than the model with $\alpha = 0$ even if there is another source of cost dynamics. However, if $\alpha \neq 0$ is mistakenly used to capture cost dynamics that are in fact due to publicly observed information, it might be more accurate to simply ignore the cost dynamics.

We recognize that nonstationarity of cost functions is a feature of the real world. However, there is no compelling reason to think that *private information* is serially correlated. Our stationary model provides a clearer view of the role of asymmetric information.

3 An Application to Global Warming

Under the assumption that uncertainty about marginal abatement costs is additive, Hoel and Karp (1998) compare taxes and quotas as a means of controlling global warming. With

additive uncertainty we need estimates of the ratio $\frac{g}{b}$ and of the parameters r, δ , and h . Even if the largest available estimate of $\frac{g}{b}$ understates the true ratio by a factor of 1000, taxes dominate quotas for reasonable values of r, δ and h . The robustness of the comparison suggests that in fact taxes are likely to yield a higher payoff than quotas. Here we want to determine if this conclusion holds when the slope of abatement costs is uncertain. We also want to examine how the choice of policies affects the long run stock levels, and how the amount of uncertainty and the initial stock level affect the ranking of policies. With multiplicative disturbances we also need estimates of the intercepts a and c and information about $var(\theta)$ and $E\frac{1}{\theta}$. We discuss the calibration and then the numerical results.

3.1 Calibration

In order to obtain estimates of the parameters of the damage and abatement cost functions, we use estimates of the absolute levels of damages and abatement costs. Our unit of time is years and we set $h = 1$, so one period equals one year. We measure costs in billions of 1990 dollars, and the stock of carbon in billions of tons. The estimated stock in 1990 was 800 billion tons ((Falk and Mendelsohn 1993), (Reilly 1992)), and the estimated Gross World Product (GWP) was 22,000 billion dollars ((Manne 1993), (OECD 1992)).

We assume that the cost of a stock higher than 800 is $\frac{g(S-800)^2}{2}$, so that the parameter c is given by $c = -800g$. We define the parameter ϕ as the annual percentage reduction in GWP due to doubling the world atmospheric stock of carbon. The parameters ϕ and g satisfy the relation

$$\frac{\phi 22000}{100} = \frac{g 800^2}{2} \implies g = 6.875 \times 10^{-4} \phi. \quad (15)$$

A high estimate for the annual cost resulting from a doubling of the stock of carbon is 400, implying $\phi = 1.8$; many other estimates are approximately half of that magnitude.¹² We consider three damage functions that correspond to three values of ϕ : $\phi = 1$ (a conservative damage estimate); $\phi = 5$ (a high damage estimate) and $\phi = 30$ (an extremely high estimate). For the conservative estimate, $g = 6.875 \times 10^{-4}$ and for the extremely high estimate, $g = 2.0625 \times 10^{-2}$.

There are a range of estimates of annual emissions in 1990. We adopt a “moderate” estimate of 6 billion tons per year ((Manne 1993), (OECD 1992),). There also exist a range of estimates of the absolute costs of reducing emissions. These estimates vary according to country and time period. It is cheaper to reduce emissions slowly, because of the lower adjustment costs (assuming that these are convex in the rate of adjustment) and because of technological improvements. A “moderate” estimate is that a 50% reduction in emissions leads to a 1% loss in GWP, or 220 billion 1990 dollars (Nordhaus 1991).¹³ We take $E_{\theta}x^* = 6$, the expected business-as-usual level of emission. We assume that at the expected value of the cost shock (i.e. at $\theta = \bar{\theta} = 1$), decreasing emissions (x) below 6 results in abatement costs of $\frac{b}{2}(6 - x)^2$ so our parameters for the cost function are $a = 6b$ and $f = 18b$. The moderate estimate (1% loss of GWP due to a 50% reduction of emissions) implies that $b = 48.9$, $a = 293.3$, and $f = 879.9$.

---

¹² For a range of estimates see: (Barns, Edmonds and Reilly 1993), (Bruce, Lee and Haites 1996), (Cline 1992), (Falk and Mendelsohn 1993), (Fankhauser 1995), (Maddison 1995), (Nordhaus 1991), (Nordhaus 1993), (Nordhaus 1994), (OECD 1992), (Tol 1995).

¹³ This estimated cost refers to a reduction that is phased in over decades, not an instantaneous reduction. Here we ignore adjustment costs. Footnote 4 explains that including adjustment costs causes a qualitative change in the model.

The conservative estimates for damages ($\phi = 1$) and the moderate estimate of abatement costs imply that the ratio $\frac{g}{b} = 1.4062 \times 10^{-5}$. This ratio is critical in ranking taxes and quotas. Since the ratio has little intrinsic economic interpretation, we perform sensitivity studies by varying the parameter ϕ (the percentage loss in GWP due to doubling the stock of carbon from 800 to 1600). Equation (15) shows that g (and thus the ratio $\frac{g}{b}$) is proportional to ϕ .

To compare the two policies we also need assumptions about the random variable. We have two free parameters, $E\frac{1}{\theta}$ and $var(\theta)$. In order to reduce the dimensionality of parameter space, we assume that θ is uniformly distributed with support $[1 - \epsilon, 1 + \epsilon]$, i.e.

$$\theta \sim U[1 - \epsilon, 1 + \epsilon]. \quad (16)$$

This distribution implies

$$E\theta = 1, \quad E\theta^2 \equiv \gamma = \frac{3 + \epsilon^2}{3}, \quad var(\theta) = \frac{\epsilon^2}{3} \quad (17)$$

$$E\frac{1}{\theta} = \frac{1}{2\epsilon} \left[\ln\left(\frac{1}{1 - \epsilon}\right) - \ln\left(\frac{1}{1 + \epsilon}\right) \right] = \frac{\ln(1 + \epsilon) - \ln(1 - \epsilon)}{2\epsilon}. \quad (18)$$

$$\Phi(\epsilon) \equiv \frac{E\frac{1}{\theta} - 1}{var(\theta)} = \frac{3[\ln(1 + \epsilon) - \ln(1 - \epsilon) - 2\epsilon]}{2\epsilon^3}. \quad (19)$$

The function $\Phi(\epsilon)$ is strictly increasing. Based on the intuition from the limiting case $\Delta = 0$ we therefore expect that a larger value of ϵ increases the difference between payoffs under the two policies.

Our estimate of $\delta = 0.005$ ((Falk and Mendelsohn 1993),(Nordhaus 1993))¹⁴ implies $\Delta = 0.995$ when $h = 1$. For greenhouse gasses Δ is not small, and therefore we need a complete solution in order to compare the policies. We assume that the continuous discount rate is $r = 0.03$.

3.2 Numerical results

Figures 1 - 3 graph the difference between the present discounted value of payoffs under taxes and quotas ($J(S; T) - J(S; Q)$) for S in the interval $[800, 2000]$. We use nine combinations of parameter values¹⁵ : $\phi \in (1, 5, 30)$ and $\epsilon \in (0.2, 0.4, 0.6)$. The value $\epsilon = 0.2$ implies a standard deviation (which equals the coefficient of variation) of 0.149, and $\epsilon = 0.6$ implies a standard deviation of 0.258.

For the conservative estimates ($\phi = 1$ and $\epsilon = .2$), when $S = 800$ taxes dominate quotas by about 450 (billion 1990 dollars – see Figure 1), approximately 2% of GWP. This amount is twice the estimate of the annual loss in GWP due to a 50% reduction in annual emissions.¹⁶

This difference decreases slightly with the stock size. Tripling the value of ϵ leads to nearly an eight-fold increase in the advantage of taxes when $S = 800$. Greater uncertainty increases the magnitude of the preference for whichever policy is optimal. For the high estimate of damages ($\phi = 5$) taxes still dominate quotas (Figure 2). This difference decreases when

---

¹⁴ This estimate may be low, thus biasing our results in favor of quotas. Other estimates suggest a value of $\delta = .0083$ (Reilly 1992).

¹⁵ The other parameter values are: $h = 1, r = .03, \delta = .005, b = 48.9, a = 293.3$ (which correspond to the estimate that a 50% reduction in emissions leads to a 1% fall in GWP).

¹⁶ Note that this amount is the difference in the present discounted value of the stream of payoffs under the two policies. With a yearly discount factor of $e^{-.03} = .97$, a present discounted value of 450 implies an annual flow of 13.3.

Table 1: (Expected) steady state stocks under taxes and quotas

	$\epsilon = .2$	$\epsilon = .4$	$\epsilon = .6$
$\phi = 1$	1173, 1158	1173, 1112	1173, 1025
$\phi = 5$	1087, 1076	1087, 1043	1087, 978
$\phi = 30$	914, 909	913, 898	913, 874

the damage parameter ϕ is larger. In addition, the difference becomes more sensitive to the stock size. For the extremely high estimate of damages ($\phi = 30$) taxes continue to dominate quotas for moderate stock levels, but by a smaller amount. However, for stock levels between 1350 and 2000, the payoff is higher under quotas.

Table 1 reports the expected steady state stock under taxes (the first entry) and the steady state under quotas (the second entry) for the nine sets of parameter values. The expected steady state under taxes is always larger than under quotas.

In all cases, the (expected) steady state stock decreases with the severity of damages and with the magnitude of uncertainty. Increasing ϵ leads to a small fall in the steady state under quotas, but a scarcely perceptible fall under taxes.¹⁷ Since we have no estimates of the actual magnitude of ϵ , this insensitivity is encouraging. Not surprisingly, increasing damages (ϕ) decreases the (expected) steady state.

Perhaps the most striking feature of Table 1 is the similarity of the steady states. The largest number (1173) is only 34% larger than the smallest number (874), despite a considerable range in parameters (300% for ϵ and 3000% for ϕ). If we think that $\phi = 5$ and $\epsilon = 0.4$ are reasonable upper bounds for the parameters, and $\phi = 1$, $\epsilon = 0.2$ are reasonable

---

¹⁷ We rounded to whole numbers, so the slight decrease of the expected steady state under taxes, due to an increase in ϵ , is usually not apparent in the table.

lower bounds, the results suggest that a target steady state of carbon stock between 1045 and 1175 is optimal. In other words, we should attempt to keep the stock below 150% of its 1990 level. The optimal steady states under severe damages ($\phi = 30$) are well below the level (1350) beyond which quotas dominate taxes. Thus, the policy choice is always time-consistent for these simulations.

In the absence of regulation, $x = 6$ and the steady state is $S = 1203$. This non-intervention steady state is approximately 2.5% higher than the largest steady state in Table 1, and 37% higher than the smallest steady state. Whether regulation leads to a large change in the outcome depends on the regulator's beliefs about the magnitude of damages and abatement costs.

We pointed out above that for our range of parameter estimates, taxes dominate quotas regardless of whether uncertainty is additive or multiplicative. However, the magnitude of the preference for a particular policy may be sensitive to the form of uncertainty. The simplest way to measure this sensitivity is to take the ratio of the payoff differences under the two types of uncertainty:

$$R(S) \equiv \frac{J(S; T) - J(S; Q)}{D}.$$

where D is defined as the difference in payoffs under taxes and quotas with additive uncertainty. The functions $J(S; i)$ and D depend on the parameters of the problem (including the variance σ^2), but D is independent of the stock of pollution.

If, for example, $\phi = 30$, we know (from Figure 3) that the numerator of R is approximately 0 for stocks slightly less than 1300. In this case, $R \approx 0$; here the gain from choosing the

right policy (taxes) is of a higher order of magnitude when uncertainty is additive rather than multiplicative.

Since $J(S; T) - J(S; Q)$ is decreasing in S , this conclusion can be reversed for low stock levels. In the appendix we explain how to calculate D and discuss the adjustment that we need to make in the variance in order to be able to compare additive and multiplicative uncertainty. We find that for reasonable parameter values and stock levels, R is a large number. For example, for $\epsilon = .2, \phi = 1, S = 800$, we have $R = 2700$. Here, the gain from choosing the right policy (taxes) is much greater when uncertainty is multiplicative.

We also noted that in order to reverse the preference for taxes over quotas under additive uncertainty, the ratio $\frac{a}{b}$ would have to be approximately 1000 times larger than a “reasonable” point estimate. Under multiplicative disturbances, on the other hand, a thirty-fold increase in this ratio can reverse the preference for taxes, if stocks are also large.

4 Conclusion

There has been great interest in the effect of economic activity on stocks of greenhouse gasses, and in the relation between these stocks and global warming. There is a growing consensus that limiting the stock of greenhouse gasses is important to human welfare, but there has been little research on the best means of achieving such a limit. A large body of literature compares taxes and quotas in the presence of asymmetric information between regulators and firms, but assumes that damages are related to emissions rather than stocks. This literature is not directly applicable to the problem of controlling greenhouse gasses.

Previous research that examined the relative merits of the two policies for stock pollutants assumes that the random variable affects the intercept but not the slope of abatement costs. We extend this literature by allowing the random variable to enter multiplicatively. This extension allows us to see how the stock size and the amount of uncertainty affects the relative payoffs under taxes and quotas. We are also able to see how the optimal (expected) stock trajectory responds to the magnitude of uncertainty, and to determine how the policy choice affects the level of the stock – in particular, the steady state level. In addition, this model enables us to determine whether the policy ranking depends on the manner in which uncertainty affects costs (additively or multiplicatively).

We used the closed form expressions and simulations to develop intuition about the characteristics of the policies. With multiplicative uncertainty, a greater variance tends to reduce the expected trajectories. Higher stocks favor the use of quotas. Taxes lead to higher steady state stocks.

We calibrated the model using published estimates of the magnitude of environmental damages of greenhouse gasses and of the abatement costs of limiting carbon emissions. In all of our simulations, we found that taxes dominate quotas. We had previously reached this conclusion using the model of additive uncertainty. Therefore, this conclusion appears (quite) robust. Although the ranking of policies is the same, the magnitude of the importance of making the right choice is much greater under multiplicative uncertainty for current stock levels.

The linear-quadratic model is obviously very special; any one of its assumptions can be

disputed. The advantage of this model, however, is that it allows a transparent calibration and a simple solution – one which is just a step above a back-of-the-envelope calculation. This transparency and simplicity is extremely useful for policy discussions, where there is likely to be considerable disagreement about the magnitude of environmental damages and abatement costs. The linear-quadratic model with multiplicative uncertainty enables us to check the robustness of the additive uncertainty model and also to ask a larger set of questions. We have used these two models to study the problem of global warming, but the same models will be useful in any situation where environmental damages depend on stocks.

5 Appendix: Technical details

In this section we: (i) derive the parameters of the value function and control rule, (ii) derive equation (14), and (iii) show how to determine whether the magnitude of the preference for a policy depends on the form of uncertainty.

5.1 The value function

Since we have a linear-quadratic control problem, the value function is quadratic: $J(S; T) = \rho_0 + \rho_1 S + \frac{\rho_2}{2} S^2$ for some parameters ρ_0, ρ_1, ρ_2 . Using this functional form, we can write the regulator's dynamic programming equation under taxes as:

$$\begin{aligned} \rho_0 + \rho_1 S + \frac{\rho_2}{2} S^2 &= \max_z \{ \lambda(z, S; T) h \\ &\quad + \beta E_\theta \left[\rho_0 + \rho_1 (\Delta S + z\theta h) + \frac{\rho_2}{2} (\Delta S + z\theta h)^2 \right] \} \\ &= \max_z \left\{ \alpha_0 + \alpha_1 \bar{\theta} h z + \frac{\alpha_2 h}{2} z^2 \right\}, \end{aligned} \tag{20}$$

which uses the definitions

$$\begin{aligned}\alpha_0 &= \left(f - cS - \frac{gS^2}{2}\right)h + \beta \left(\rho_0 + \rho_1 \Delta S + \frac{\rho_2}{2} (\Delta S)^2\right) \\ \alpha_1 &= a + \beta (\rho_1 + \rho_2 \Delta S) \\ \alpha_2 &= \beta \rho_2 h \gamma - b \bar{\theta}.\end{aligned}$$

The optimal control rule is obtained by performing the maximization in equation (20):

$$z^* = -\frac{\alpha_1 \bar{\theta}}{\alpha_2}. \quad (21)$$

This equation is reproduced as equation (6) in the text.

Substituting equation (21) into (20) and equating coefficients of powers of S gives the equations for ρ_i in the text.

5.2 The function Φ

We define $\theta = 1 + kw$, where w is an arbitrary random variable with mean 0, variance σ^2 , and the distribution of w is sufficiently regular that we can interchange the order of differentiation and integration. Here $\text{var}(\theta) = k^2 \sigma^2$; as $k \rightarrow 0$, $\text{var}(\theta) \rightarrow 0$. Substituting $\theta = 1 + kw$ into equation (14), we have $\Phi = -E \left[\frac{w}{(1+kw)} \right] \frac{1}{k\sigma^2}$. Using L'Hospital's Rule, we obtain $\Phi \rightarrow 1$ as $k \rightarrow 0$ (i.e., $\text{var}(\theta) \rightarrow 0$). Taking the derivative of Φ with respect to k gives $\Phi'(k) = \frac{E\Omega'(k)}{\sigma^2}$ with $\Omega \equiv -\frac{w}{k(1+kw)}$. We have

$$\frac{d\Omega}{dk} = \frac{w \frac{1+2kw}{(1+kw)^2}}{k^2}$$

Using L'Hospital's rule a second time gives

$$\Phi'(0) = \frac{-Ew^3}{\sigma^2}$$

The second derivative of Φ is

$$\Phi'' = \frac{E \left[-2w \frac{3kw+3k^2w^2+1}{(1+kw)^3} \right]}{k^3 \sigma^2}$$

Applying L'Hospital's rule a third time gives

$$\Phi''(0) = \frac{2Ew^4}{\sigma^2}$$

The second order approximation of Φ is

$$\Phi(k) \approx 1 - \frac{Ew^3}{\sigma^2}k + \frac{Ew^4}{\sigma^2}k^2. \quad (22)$$

Our description of Φ in the text follows from equation (22).

5.3 The magnitude of the preference for a policy

Here we explain how to calculate D , the difference in payoffs under additive uncertainty, and how to adjust the variance so that the payoffs under additive and multiplicative disturbances are comparable.

If abatement costs are $-(f + (a + \theta^*)x - \frac{b}{2}x^2)$, where θ^* is a zero mean, i.i.d. random variable with variance $var(\theta^*)$, uncertainty is additive. Using results from Hoel and Karp (1998), the difference between the present discounted value of payoff under taxes and quotas is

$$D = \frac{var(\theta^*) \left(1 + \frac{\beta \rho_2}{b}\right)}{(1 - \beta)2b}. \quad (23)$$

The parameter ρ_2 can be computed using equation (7) with $\gamma = 1$.

Under a tax z , the variance in the additional pollution in a period is $(z)^2 var(\theta)$ when uncertainty is multiplicative, and the variance is $var(\theta^*)$ when uncertainty is additive. Therefore, if we want approximately the same order of magnitude of the uncertainty in the flows

in the two cases, we must to set $var(\theta^*) = z^2 var(\theta)$. The equilibrium z is a (decreasing) function of S , but $z < 6$ for any $S \geq 800$. We therefore obtain an upper bound for the difference in the value functions under additive uncertainty, which we can compare with the difference in value functions under multiplicative uncertainty, by setting $var(\theta^*) = 6^2 var(\theta)$.

For example at $\epsilon = .2, \phi = 1$, this upper bound is .16608, or 166 million 1990 dollars. We noted that for $\epsilon = .2, \phi = 1$, the payoff under taxes exceeds the payoff under quotas by 450 billion dollars when $S = 800$. In this case, $R(800) = \frac{450}{.166} = 2711$.

References

- Barns, D., Edmonds, J. and Reilly, J.: 1993, “The Edmonds-Reilly-Barns model”, *The Cost Cutting Carbon Emission: Results from Global Models*, OECD, Paris.
- Bertsekas, D.: 1995, *Dyanamic Programming and Optimal Control, Vol I and II*, Athena Scientific Press, Belmont MA.
- Bruce, J., Lee, H. and Haites, E. (eds): 1996, *Climate Change 1995: Economic and Social Dimensions of Climate Change*, Cambridge University Press.
- Cline, W.: 1992, *The Economics of Global Warming*, Institute for International Economics, Washington D.C.
- Falk, I. and Mendelsohn, R.: 1993, The economics of controlling stock pollutants: An efficient strategy for greenhouse gasses, *Journal of Environmental and Economic Management* **25**(1), 76–133.
- Fankhauser, S.: 1995, *Valuing Climate Change: The Economics of the Greenhouse*, Earthscan, London.
- Hoel, M. and Karp, L.: 1998, Taxes versus quotas for a stock pollutant, *Fondazione Eni Enrico Mattei, Nota di Lavoro, 29.98*, <http://are.Berkeley.EDU/karp/>.
- Karp, L. and Zhang, J.: 1999, Regulation of stock externalities with learning, Working Paper No. 892, Department of Agricultural and Resource Economics, University of California, Berkeley; <http://are.Berkeley.EDU/karp/>.
- Karp, L. and Zhang, J.: 2000, Stock externality regulation with endogenous technical change and learning, Department of Agricultural and Resource Economics, University of California, Berkeley; <http://are.Berkeley.EDU/karp/>.
- Maddison, D.: 1995, A cost-benefit analysis of slowing climate change, *Energy Policy* **23**(4-5), 337–46.
- Malcomson, J.: 1978, Prices vs. quantities: A critical note on the use of approximations, *Review of Economic Studies* **45**(1), 203–07.
- Manne, A.: 1993, “Global 2100: Alternative scenarios for reducing carbon emissions”, *The Cost Cutting Carbon Emission: Results from Global Models*.
- Newell, R. and Pizer, W.: 1998, Stock externality regulation under uncertainty, Resources for the Future Working Paper.
- Nordhaus, W.: 1991, To slow or not to slow: The economics of the greenhouse effect, *The Economic Journal* **101**(407), 920–37.

- Nordhaus, W.: 1993, Rolling the 'dice': An optimal transition path for controlling greenhouse gases, *Resource and Energy Economics* **15**(1), 27–50.
- Nordhaus, W.: 1994, *Managing the Global Commons: The economics of the greenhouse effect*, Cambridge, MA: MIT Press.
- OECD: 1992, Global warming: The benefits of emission abatement, OECD, Paris.
- Reilly, J.: 1992, "Climate-change damage and the trace-gas-index issue", in J. Reilly and M. Anderson (eds), *Economic Issues in Global Climate Change: Agriculture, Forestry, and Natural Resources*, Boulder and Oxford: Westview Press, pp. 72–88.
- Roberts, M. and Spence, A.: 1976, Effluent charges and licenses under uncertainty, *Journal of Public Economics* **5**(3-4), 193–208.
- Staring, K.: 1995, Regulating av forurensing under usikkerhet - en dynamisk modell - - regulation of pollution under uncertainty - a dynamic model, Master degree thesis, Department of Economics, University of Oslo.
- Stavins, R.: 1996, Correlated uncertainty and policy instrument choice, *Journal of Environmental Economics and Management* **30**(2), 218–232.
- Tol, R.: 1995, The damage costs of climate change: Towards more comprehensive calculations, *Environmental and Resource Economics* **5**(4), 353–74.
- Watson, W. and Ridker, R.: 1984, Losses from effluent taxes and quotas under uncertainty, *Journal of Environmental Economics and Management* **11**(4), 310–26.
- Weitzman, M.: 1974, Prices versus quantities, *Review of Economic Studies* **41**(4), 477–91.
- Weitzman, M.: 1978, Reply to "Prices vs. quantities: A critical note on the use of approximations" by James M. Malcomson, *Review of Economic Studies* **45**(1), 209–10.
- Yohe, G.: 1977, Comparisons of price and quantity controls: A survey, *Journal of Comparative Economics* **1**(3), 213–33.

Figure 1: Payoff Difference Between Taxes and Quotas, $\phi = 1$

Figure 2: Payoff Difference Between Taxes and Quotas, $\phi = 5$

Figure 3: Payoff Difference Between Taxes and Quotas, $\phi = 30$