The role of stocks and shocks concepts in the debate over price vs. quantity

John E. Parsons and Luca Taschini

The Role of Stocks & Shocks Concepts
in the Debate Over Price vs. Quantity[*]

John E. Parsons[a][†]   Luca Taschini[b][‡]

[a]*Massachusetts Institute of Technology*
[b]*London School of Economics*

October 26, 2011

Abstract

Many economists and policy makers have long favored the use of a price instrument to control greenhouse gases because they are a stock pollutant and as such the marginal benefit of abatement is relatively flat. While the early literature on the problem is consistent with this view, the later literature is less categorical. It showed that the choice between a price or quantity control depends, in part, upon the assumption on the dynamic structure of cost uncertainty. Temporary shocks to abatement cost favors the use of a price control, while permanent shocks favor a quantity control. Unfortunately, the importance of this assumption to the optimal choice has not yet received wide currency among economists. We analyze the problem in an alternative setting and reproduce the result that temporary shocks favor use of a price control while permanent shocks favor use of a quantity control. Our contribution is the simplicity of the model and the accessibility of the results, which reinforce the critical role played by the assumed structure of uncertainty.

Keywords: Cap & trade, Permanent Shocks, Tax, Transitory Shocks.

JEL Classifications: H23, Q28, Q50, Q58.

1 Introduction

Two classic alternatives for regulating stock pollutants, such as greenhouse gas emissions, are a cap & trade system or a carbon tax. Economists refer to the former as a quantity control and the latter as a price control. While a cap & trade system yields a price, this is a secondary result of regulating the quantity. Correspondingly, a carbon tax effectively reduces the quantity of emissions, but as a secondary result of setting a price. Under idealized circumstances the two methods are equivalent. If the parameters of the underlying economy are well known, then there is a simple duality in the problem and it doesn't matter whether it is the price or the quantity which is fixed directly.

Of course, circumstances are never ideal. Considering a static model of uncertainty about the cost curve, Weitzman (1974) showed that a flat expected marginal benefit function, relative to marginal cost, favours a price control, while a steep marginal benefit function favours a quantity control. Intuitively, a flat marginal benefit function implies a constant benefit per unit of pollution abated, suggesting that a tax could best correct the externality. In contrast, a steep marginal benefit function implies a hazardous threshold that should be avoided and that is efficiently enforced by a quantity control.

The early literature extends Weitzman's static model to a long term horizon. This literature emphasizes the fact that for a stock pollutant the marginal benefit function for abatement within any single period is flat while the marginal cost function slopes sharply, so that the Weitzman logic argues in favor of a price control within each period –see, for example, Nordhaus (1994) (Ch. 8, fn. 4), Pizer (1999) and Hoel and Karp (2002). Although the choice ultimately depends upon the empirical parameters of the problem, the stock pollutant nature of greenhouse gases appeared to make the outcome obvious.

Later research on the dynamic problem pointed out that the early literature had made a narrow assumption about the dynamic character of the uncertainty on cost through time, and solved a more general case. This later literature solved a more general case and showed that uncorrelated uncertainty on cost across periods (temporary uncertainty) led to a preference for using a price control in each period, while correlated uncertainty on cost (permanent uncertainty) led to a preference for using a quantity control –see, for example, Newell and Pizer (2003) and Karp and Zhang (2005). Consequently, whether a price control or a quantity control is preferred depends crucially on ones estimate of the degree of correlation in costs across periods, together with the other parameters of the problem, including the stock pollutant nature of greenhouse gases. Assessing the relative strength

of the different factors is more difficult. While the stock pollutant nature of greenhouse gases clearly and dramatically flattens the marginal benefit function for emissions within a single period, the correlation of costs across periods, depending upon how strong it is, could operate on a comparable scale, making the choice between a quantity control or a price control more contentious than previously understood.

The original claim from the early literature asserting that a price control is superior for stock pollutants like greenhouse gases achieved wide currency among economists and policy makers. The later revision, noting the important role of the assumption about the dynamics of cost uncertainty, is less widely appreciated. For example, the Stern Review (2006) reports that the stock pollutant character of the greenhouse gas problem unambiguously argues in favor of a price control.

This paper attempts to address the lag in digesting the later research results by reproducing the results in a slightly different model with some useful features. The model is relatively simple, so that the optimal dynamic policy is easily derived and can be handily simulated. The match between alternative extreme assumptions on the dynamic structure on costs and the preference for either a price or cost control is stark, so that the underlying relationship between the assumption and the result is highlighted. As always, reproducing a fundamental result in a different context helps clarify the essential relationship between the assumptions and the result.

The bottom line is that the early stock pollutant story oversimplifies the problem. In imagining an extension of the Weitzman model from a single period to a multiple one, the early stock pollutant story makes a strong implicit assumption that the relevant uncertainty involves purely transitory shocks to the cost function. The logic is inconsistent with uncertainty about permanent or even lasting shocks. When the shocks are permanent, the stock pollutant problem looks exactly like the one period problem so that which form of control is optimal – prices or quantities – is a difficult empirical problem. The fact that greenhouse gases are a stock pollutant does not alter the situation at all if the important cost uncertainties are about permanent shocks to the cost function.

In the next section of the paper, we neglect benefits and concentrate the analysis on the cost-effectiveness problem. Assuming a given cap on aggregate emissions, we present a model of dynamic abatement cost uncertainty which can be easily solved to yield the cost minimizing emissions policy. The model can be calibrated to the extreme case in which all cost uncertainty is temporary or to the other extreme case in which all cost uncertainty is permanent. These are two extreme special cases of a more general specification used,

for instance, in Karp and Zhang (2005), and correspond to the extreme special cases of perfectly correlated or perfectly uncorrelated uncertainty specification. In the case of permanent uncertainty, the cost minimizing emissions policy does not need to be contingent on knowing the realization of the cost parameter at all. So no matter the information available to the regulator about the realization of the cost parameter, the cost minimizing emissions policy can be implemented using a quantity control. In this case, using a price control will always be suboptimal (in cost-effective terms) with respect to minimizing the cost of achieving a fixed level of emissions. This result is consistent with Newell and Pizer (2003) and Karp and Zhang (2005) findings: the more correlated the shocks (ie. shocks have a permanent impact), the more likely is the optimal policy to be a quantity control. The case of temporary uncertainty, shows just the opposite, that the cost minimizing policy does require adjusting the emissions level, and that this is equivalent to setting a price schedule each period independent of the realization of the cost parameter. Therefore, a regulator who is unable to observe the cost parameter for any period of time can more closely attain the cost minimizing emissions policy by means of a price instrument as opposed to a quantity instrument.

In the third section, we show that the insights from the model of cost effectiveness extend to the fuller problem of weighing costs and benefits. We do this by solving a pair of two-stage decision problems –one of completely temporary uncertainty and one of completely permanent uncertainty. In the temporary uncertainty case, the original Weitzman model together with the assumption that greenhouse gasses are a stock pollutant, combine to suggest that a price control is the preferred regulatory action. In the permanent uncertainty case, we show that the stock pollutant character of greenhouse gases is irrelevant to the problem, and we are thrown back onto the original Weitzman problem where the preference for price or quantity control is an empirical question.

2   Temporary&Permanent Shocks to Abatement Costs

We analyze emissions within a time horizon divided into N periods indexed by $\{i, i = 1, 2, \ldots, N\}$. Emissions in each period are denoted q_i, which is a control variable that can be adjusted without constraint. Costs are incurred at a rate that is a function of emissions and a cost parameter, θ_i,

$$c_i(q_i, \theta_i) = e^{\theta_i - q_i}.$$

With this form, marginal costs are the negative of the cost,

$$\frac{\partial c_i(q_i, \theta_i)}{\partial q_i} = -e^{\theta_i - q_i}.$$

Higher emissions lower cost. Cutting emissions –abatement– increases cost. Increasing the parameter θ_i shifts cost up while also steepening the cost curve, so that both the cost of abatement and the marginal cost of abatement increases. We select this form for the cost function because it allows us to conveniently handle a multi-period optimization problem with discounting. It has the disadvantage that no matter how large the emissions, there is some positive cost. Nevertheless, it is useful for unpacking the issues at hand in this paper.

We evaluate two contrasting specifications of cost uncertainty. These are two extreme special cases of a more general specification, and correspond to the extreme special cases of perfectly correlated or perfectly uncorrelated uncertainty in the specification used in Newell and Pizer (2003) and Karp and Zhang (2005). In the first specification, shocks to the cost parameter are completely temporary or transitory. A shock affects the cost parameter in that period, but has no impact on the cost parameter in any future period. Under the second specification, shocks to the cost have a permanent impact. A shock affects the cost parameter in that period, and the expected cost in all future periods is incremented by the same amount, too.

The first specification of the cost parameter θ_i is:

$$\theta_i = \theta_0 + i\nu + \sigma\epsilon_i, \tag{1}$$

where θ_0 is the starting cost parameter, ν is the constant expected growth rate in the mean cost parameter, and ϵ_i are independent standard normal random variables, i.e. the shocks to the cost parameter. This defines a process that is white noise around a linearly increasing trend. It is comparable to a mean reverting process with full reversion to the growing mean within the period.

The second specification of the cost parameter θ_i is:

$$\theta_i = \theta_{i-1} + \mu + \sigma\epsilon_i. \tag{2}$$

where μ is the constant expected growth in the cost parameter. This process is a random walk with trend. It is often said that the random walk has "infinite memory". It is in this sense that we say the shocks have a permanent impact on the cost parameter. A very simple and intuitive presentation of the difference between temporary and permanent uncertainty can be found in the Appendix.

We concentrate our attention on the cost-effectiveness solution neglecting benefits and assuming a fixed target of aggregate emissions, $\overline{q}$, over the N periods so that

$$\sum_{i=1}^{N} q_i \leq \overline{q}.$$

Considering the fixed aggregate emission constraint $\overline{q}$, we ask what is the cost minimizing dynamic emissions policy in light of the stochastic evolution of the cost parameter. Our purpose is to show how the degree to which uncertainty is temporary or permanent shapes the cost minimizing emissions policy. To anticipate our results, we compare the variability in quantity and price under the cost minimizing emissions policy. We show that when uncertainty is temporary, and θ_i follows Equation (1), most of the period-by-period variability in the cost parameter translates into variability in the quantity of emissions. Price –actually, marginal cost– is relatively constant. We show that, in contrast, when uncertainty is permanent, and θ_i follows Equation (2), all of the period-by-period variability in the cost parameter translates into variability in the price of emissions. Quantity is constant.

We solve for the cost minimizing policy using backward induction. We first set up the general solution format we use, and then we solve each of the cases. A dynamic emissions policy will set emissions in each period conditional on some function of past aggregate emissions and on the current value of the cost parameter. We denote the remaining allowed emissions as we arrive in period i by $\bar{q}_i$. Analytically, $\bar{q}_{i+1} = \bar{q}_i - q_i$ for $i = 1, \ldots, N-1$. The choice of emissions will also depend upon the level of the cost parameter, θ_i, and so we write emissions as a function of these two parameters, $q_i(\bar{q}_i, \theta_i)$. We denote the value function in period $i = 1, \ldots, N$, as V_i. Since the target of aggregate emissions is fixed, in the final period when $i = N$, the value function is simply the total cost of remaining emissions to abate:

$$V_N(\bar{q}_N, q_N, \theta_N) \equiv c\Big(q_N(\bar{q}_N, \theta_N), \theta_N\Big).$$

The cost minimizing emissions level, $q_N^*(\bar{q}_N, \theta_N)$, is the solution of the following problem

$$\min_{q_N} V_N(\bar{q}_N, q_N, \theta_N)$$

subject to the constraint $\sum_{i=1}^{N} q_i \leq \bar{q}$. Given the cost function, the solution is trivially:

$$q_N^*(\bar{q}_N, \theta_N) = \bar{q}_N.$$

We denote the optimized value function as V_N^*. It is a function of the remaining allowed emissions coming into the period and the realized cost parameter in the period:

$$
\begin{aligned}
V_N^*(\bar{q}_N, \theta_N) &\equiv V_N\Big(\bar{q}_N, q_N^*(\bar{q}_N, \theta_N), \theta_N\Big) \\
&= c\Big(q_N^*(\bar{q}_N, \theta_N), \theta_N\Big) = c(\bar{q}_N, \theta_N).
\end{aligned}
$$

We will also want to take note of the marginal cost of emissions under this optimal policy which is:

$$p_N^*(\bar{q}_N, \theta_N) \equiv -\frac{\partial c(q_N^*(\bar{q}_N, \theta_N), \theta_N)}{\partial q_N} = c(q_N^*(\bar{q}_N, \theta_N), \theta_N).$$

where $p_N^*(\bar{q}_i, \theta_i)$ represents the shadow price and is defined as the negative of marginal cost. For convenience of comparison, we will generally focus on the log of the marginal cost, $\ln(p_i^*) = \theta_i - q_i^*$.

In earlier periods, when $1 \leq i < N$, the value function is the total cost of current period emissions plus the discounted expectation of the value function in the subsequent period:

$$V_i\Big(\bar{q}_i, q_i, \theta_i\Big) \equiv c\Big(q_i(\bar{q}_i, \theta_i), \theta_i\Big) + \mathbb{E}_{\theta_i}\Big[V_{i+1}^*\Big(\bar{q}_{i+1}(\bar{q}_i, q_i), \theta_{i+1}\Big)\Big].$$

The uncertainty about θ_i is resolved at the end of each period, therefore the expectation operator is conditioned on the realisation of the cost parameter. The cost minimizing emissions level $q_i^*(\bar{q}_i, \theta_i)$ solves

$$\min_{q_i} V_i(\bar{q}_i, q_i, \theta_i).$$

The optimized value function is:

$$V_i^*(\bar{q}_i, \theta_i) \equiv V_i\Big(\bar{q}_i, q_i^*(\bar{q}_i, \theta_i), \theta_i\Big).$$

The marginal cost of emissions is:

$$p_i^*(\bar{q}_i, \theta_i) \equiv -\frac{\partial c(q_i^*(\bar{q}_i, \theta_i), \theta_i)}{\partial q_i} = c(q_i^*(\bar{q}_i, \theta_i), \theta_i).$$

The sequence of emissions functions, $q_i^*(\bar{q}_i, \theta_i)$, form the cost minimizing dynamic emissions policy. The sequence of price functions, $p_i^*(\bar{q}_i, \theta_i)$, form the price corresponding to the cost minimizing dynamic emissions policy.

We now turn to examining the solution to this problem under different circumstances. We begin by solving the certainty case, since this provides useful intuition for the uncertainty cases. We then solve the two contrasting uncertainty cases and show how the degree to which uncertainty is temporary or permanent shapes the cost minimizing emission policy.

Certainty Case For the certainty case we have $\sigma = 0$, so that the dynamics of θ_i reduces to

$$\theta_i = \theta_0 + i\nu \quad \text{for} \quad i = 1, \ldots, N.$$

As shown in the Appendix, the cost minimizing emissions path has a conveniently simple general form:

$$q_i^* = \frac{1}{N - i + 1}\bar{q}_i - \frac{1}{2}(N - i)(\nu - r), \tag{3}$$

and

$$V_i^*(\bar{q}_i, \theta_i) = ie^{\theta_i - q_i^*} = ie^{\theta_i - \frac{1}{N-i+1}\bar{q}_i + \frac{1}{2}(N-i)(\nu-r)}. \tag{4}$$

The log marginal cost of emissions is:

$$\ln\Big(p_i^*(\bar{q}_i, \theta_i)\Big) = \theta_i - q_i^*(\bar{q}_i, \theta_i) = \theta_i - \frac{1}{N - i + 1}\bar{q}_i + \frac{1}{2}(N - i)(\nu - r). \tag{5}$$

The expressions in Equations (3)–(5) are contingent on whatever may be the endowment of remaining allowed emissions coming into the period, $\bar{q}_i$, and they are expressed in terms

8

of the remaining number of periods. Therefore, it is not immediately clear from Equation (3) how the emissions in different periods compare to one another. In the certainty case, we can readily translate back to an equation that describes emissions in different periods as a function of the total allowed emissions, $\bar{q}$, the total number of periods, N, the rate of growth in the cost parameter, ν, and the discount rate, r :

$$q_i^*(\bar{q}, N, \nu, r) = \frac{1}{N}\bar{q} + (i - \frac{N}{2})(\nu - r). \tag{6}$$

Transforming the price –the negative of the marginal cost– in log terms for the ease of exposition, p_i^* can be expressed as:

$$\ln\left(p_i^*(\bar{q}, N, \nu, r)\right) = \theta_0 - \frac{1}{N}\bar{q} + \frac{N}{2}(\nu - r) + ir. \tag{7}$$

To understand the optimal emissions policy in Equation (3) begin by assuming that $\nu = r$, so that the cost parameter is growing at a rate equal to the discount rate. In that case, the optimal policy is to allocate to period i a pro rata share of the remaining allowed emissions, $\frac{1}{N-i+1}\bar{q}_i$. Application of this policy to successive periods means that emissions are equal in every period, $\frac{\bar{q}}{N}$. The (log) marginal cost of emissions rises, $\theta_0 - \frac{\bar{q}}{N} + ir$, but at a rate equal to the discount rate, so that the discounted marginal cost is equal across all periods. When $\nu \neq r$ the optimal policy is to adjust the pro-rate allocation in period i to reflect the differential between the growth rate on the cost parameter and the discount rate. The adjustment assures that emissions increase linearly through time at the rate $\nu - r$, as seen in Equation (6) and shown in Figure 1(a), so that the marginal cost of abatement grows at the discount rate, r, as seen in Equation (7) and reported by Figure 1(b). If $\nu \geq r$, and the underlying cost parameter is growing at a rate greater than the discount rate, then this adjustment leads to reducing the rate of emissions now, in period i, increasing the realized marginal cost today, so as to preserve allowed emissions for the later periods when the cost parameter is higher, thus reducing the growth rate in the realized marginal cost to equal the discount rate.

An important feature to take note of in the solution to this certainty case is that the optimal emissions level, q_i^*, is independent of the realized cost parameter, θ_i. The cost minimizing emissions path is fully determined by (i) the quantity of emissions being targeted relative to the time remaining, and (ii) the rate of growth in the cost parameter

9
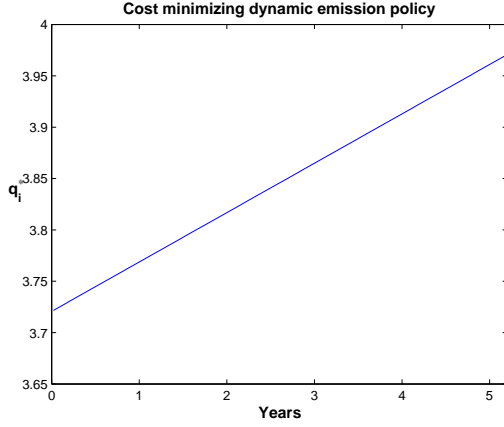
(a) Optimal emission policy q_i^*.
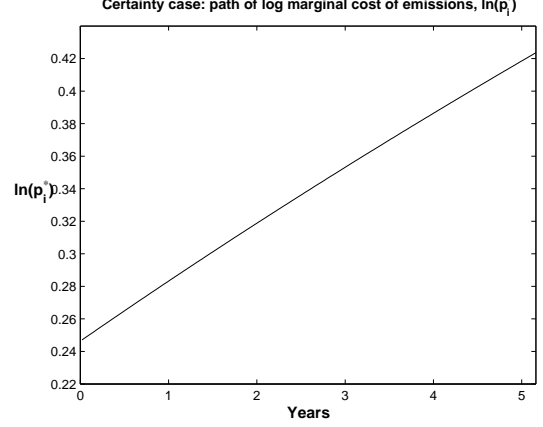
(b) Log of marginal cost $\ln\left(\frac{\partial c(q_i^*, \theta_i)}{\partial q_i^*}\right)$.

Figure 1: 5-year evolution of the cost minimizing dynamic emission policy, q_i^*, (left) and the log of marginal cost of emissions, $\ln\left(\frac{\partial c(q_i^*, \theta_i)}{\partial q_i^*}\right)$, (right) for the certainty case. In this example $\bar{q} = 1,000$, $\theta_0 = 5$, $\nu = 0.1$, $\sigma = 0$, $r = 0.05$, and the time step corresponds to a week.

relative to the discount rate. The level of the cost parameter does not enter the equation. If we change the current value of the cost parameter, we don't change the cost minimizing emissions policy![1] This fact significantly aids our solution of the cost minimizing policy when the evolution of the cost parameter is uncertain, i.e., when we allow $\sigma > 0$, whether for the case of temporary or permanent shocks.

Temporary Uncertainty Case As shown in the Appendix, the general form of the optimal dynamic policy is:

$$q_i^* = \frac{1}{N - i + 1}\bar{q}_i - \frac{1}{2}(N - i)(\nu - r) - A_i\sigma^2 + \frac{N - i}{N - i + 1}\sigma\epsilon_i \tag{8}$$

where

$$A_i = \frac{N - i}{N - i + 1}\left(A_{i+1} + \frac{1}{2(N - i)^2}\right) \quad \text{for} \quad i = 1, \ldots, N - 1 \quad \text{and} \quad A_N = 0$$

---

[1]Of course, if we were solving for the optimal emissions path, trading off costs and benefits, then we would consider the level of the costs. But we would also be comparing the aggregate benefits over the full horizon against the aggregate cost minimizing emissions policy over the full horizon.

The general form of the optimized value function is:

$$V_i^*(\bar{q}_i) = ie^{\theta_i - q_i^*} = ie^{\theta_i - \frac{1}{N-i+1}\bar{q}_i + A_i\sigma^2 + \frac{1}{2}(N-i)(\nu-r) - \frac{N-i}{N-i+1}\sigma\epsilon_i}.$$

The log price is:

$$\ln(p_i^*) = \theta_0 + i\nu - \frac{1}{N-i+1}\bar{q}_i + A_i\sigma^2 + \frac{1}{2}(N-i)(\nu-r) - \frac{1}{N-i+1}\sigma\epsilon_i$$

The optimal emissions policy in Equation (8) is similar to the certainty case in two of the components: the pro rata share of the remaining allowances, $\frac{1}{N-i+1}\bar{q}_i$, and the linear growth factor, $\frac{1}{2}(N-i)(\nu-r)$. In addition, there is a deduction in the current emissions level, $A_i\sigma^2$, which is tied to the overall volatility of emissions. This is an adjustment to the inter-temporal allocation of emissions necessitated by the increasing variability of emissions through time. Finally, there is the component of emissions that fluctuates with the current realization of costs: $\frac{N-i}{N-i+1}\sigma\epsilon_i$. If the remaining number of periods is large, then the coefficient is close to 1, which means that all of the shock in the cost parameter is absorbed in adjustment to the current level of emissions. This adjustment keeps the current level of marginal cost approximately constant. As the remaining number of periods declines, the coefficient on the quantity adjustment falls, so that only a portion of the shock in the cost parameter is absorbed in adjustment to the current level of emissions. This is because of the fixed aggregate emissions constraint. Any adjustment in the current level of emissions must be compensated for with an opposite adjustment in emissions over all of the remaining periods. The coefficient on the quantity adjustment, $\frac{N-i}{N-i+1}$, results in all periods sharing equally in the increased or decreased marginal cost so as to minimize the aggregate cost impact. When there are fewer remaining periods to share the remaining costs, a larger fraction must be absorbed in the current period. Consequently, as the final period approaches, price reflects a larger and larger portion of the shock of the cost parameter.

These points can be formalized by showing the formula for the variance of emissions and of the log of price. The variance of emissions and log price one period ahead are:

$$Var_{i-1}(q_i^*) = \frac{N-i}{N-i+1}\sigma, \quad \text{and} \quad Var_{i-1}(\ln(p_i^*)) = \frac{1}{N-i+1}\sigma.$$

The variance of the forecasted emissions and log price at any period, i, relative to the starting period, $i = 0$, are:

$$Var_0(q_i^*) = \sqrt{\sum_{h=1}^{i-1} \left(\frac{1}{N-h+1}\right)^2 + \left(\frac{N-i}{N-i+1}\right)^2} \, \sigma.$$

and

$$Var_0(\ln(p_i^*)) = \sqrt{\sum_{h=1}^{i} \left(\frac{1}{N-h+1}\right)^2} \, \sigma.$$

Figure 2(a) shows a pair of one-standard deviation confidence bounds around the expected path of the optimal quantity of emissions through time. Figure 2(b) shows a pair of one-standard deviation confidence bounds around the expected path of the log of marginal cost through time.

(a) Expected path and confidence bounds for q_i^*.

(b) Expected path and confidence bounds for the log of marginal cost $\ln\left(\frac{\partial c(q_i^*,\theta_i)}{\partial q_i^*}\right)$.

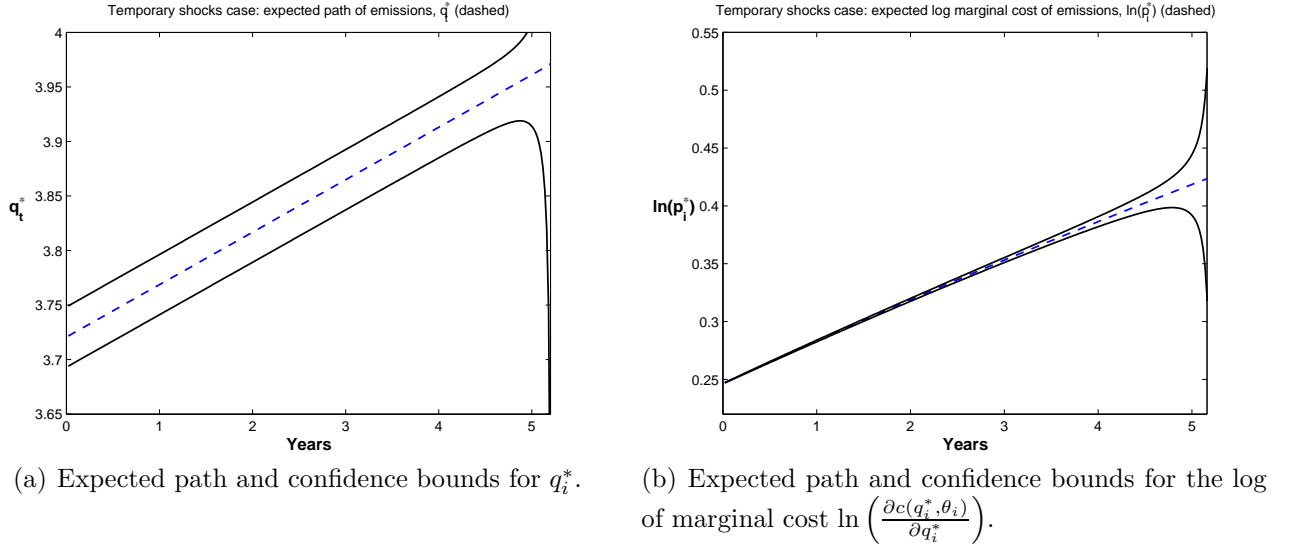Figure 2: One-standard deviation confidence bounds around the expected path of the optimal quantity of emissions, q_i^*, (left) and around the expected path of the log of marginal cost, $\ln\left(\frac{\partial c(q_i^*,\theta_i)}{\partial q_i^*}\right)$, (right). In this example $\bar{q} = 1,000$, $\theta_0 = 5$, $\nu = 0.1$, $\sigma = 0.2$, $r = 0.05$, 5-year horizon, and the time step corresponds to a week.

Permanent Uncertainty Case As shown in the Appendix, the general form of the optimal dynamic policy is:

$$q_i^* = \frac{1}{N-i+1}\bar{q}_i - \frac{1}{2}(N-i)(\mu - r), \tag{9}$$

and

$$V_i^*(\bar{q}_i, \theta_i) = ie^{\theta_i - q_i^*} = ie^{\theta_i - \frac{1}{N-i+1}\bar{q}_i + \frac{1}{2}(N-i)(\mu-r)}.$$

The log price is:

$$\ln(p_i^*) = \theta_{i-1} + \mu + \sigma\epsilon_i - \frac{1}{N-i+1}\bar{q}_i + \frac{1}{2}(N-i)(\mu - r) \tag{10}$$

With a permanent shock, the cost in all periods is equally higher or lower, so that an adjustment in the current period emissions will not help. In fact, the distribution of emissions over time has no reason to change. The optimal emissions policy in Equation (9) and represented in Figure 3(a) is identical to the certainty case. Emissions in each period are a proportional fraction of the remaining available allowances as determined by the remaining number of periods over which those allowances must be shared. Emissions are adjusted for an allowance for growth in emissions to match the growth rate in the cost parameter and are independent of the cost parameter itself. Emissions are entirely unresponsive to shocks to the cost parameter. Since none of the cost uncertainty is absorbed by the quantity of emissions, all of the cost uncertainty must be absorbed by the price as shown in Equation (10).

These points can be formalized by showing the formula for the variance of emissions and of the log of price. The variance of emissions and log price one period ahead are:

$$Var_{i-1}(q_i^*) = 0, \quad \text{and} \quad Var_{i-1}(\ln(p_i^*)) = \sigma.$$

The variance of the forecasted emissions and log price at any period, i, relative to the starting period, $i = 0$, are:

$$Var_0(q_i^*) = 0, \quad \text{and} \quad Var_0(\ln(p_i^*)) = \sqrt{i}\sigma.$$

13

Permanent shocks case: a collection of sample paths of emissions, q_i^*

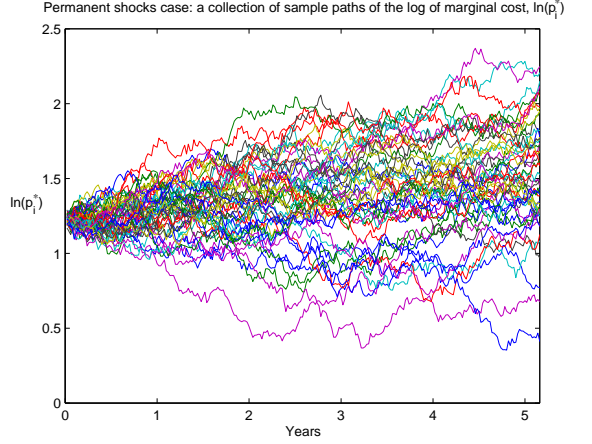Permanent shocks case: a collection of sample paths of the log of marginal cost, $\ln(p_i^*)$
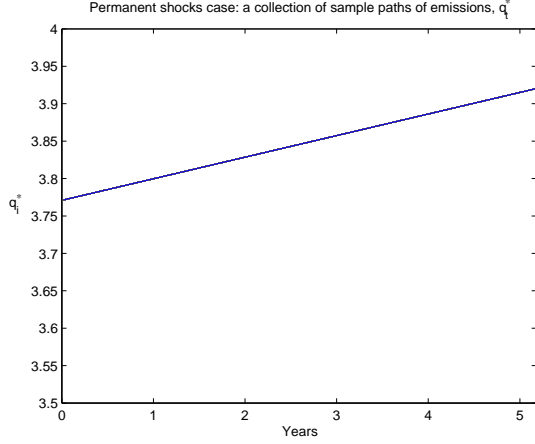
(a) Optimal emission policy q_i^*.

(b) Collection of sample paths of the log of marginal cost of emissions $\ln\left(\frac{\partial c(q_i^*,\theta_i)}{\partial q_i^*}\right)$

Figure 3: 5-year evolution of the cost minimizing dynamic emission policy, q_i^*, (left) and the log of marginal cost of emissions, $\ln\left(\frac{\partial c(q_i^*,\theta_i)}{\partial q_i^*}\right)$, (right) for the permanent shock case. In this example $\bar{q} = 1,000$, $\theta_0 = 5$, $\mu = 0.08$, $\sigma = 0.2$, $r = 0.05$, the time step corresponds to a week, and we consider a set of 50 sample paths.

Figure 3(a) shows the deterministic emissions policy in the face of permanent uncertainty in the cost parameter. Figure 3(b) shows the log of the marginal price.

These two cases provide sharp insight into the different impact that uncertainty in cost should have upon the cost minimizing emissions path depending upon whether it is a temporary uncertainty or a permanent uncertainty. In the case of temporary uncertainty, it is the quantity of emissions that absorbs shocks to the cost parameter, while the price of emissions is relatively constant. In the case of permanent uncertainty, quantity is constant and it is price that absorbs shocks to the cost parameter.

3 A Discrete Time Pair of Examples

In the previous section we fixed a cumulative emission target and solved for the cost-effective emissions policy. We did not weigh the costs against the benefits. In particular, in the case of permanent uncertainty, we showed that quantity was entirely invariant with respect to shocks to the cost parameter. This would not be true if one were weighting costs against benefits. Quantity would respond to shocks to the cost parameter, even if only to a small degree. In this section we show how the intuition developed above nevertheless

14

extendeds to the case of a complete weighing of costs against benefits.

A careful modeling of costs and benefits in a dynamic context like the one above is complex. In order to simplify things and for the sake of illustration, we construct two extremely stylized examples of temporary and permanent uncertainty within a simple two-stage decision model of costs and benefits. We assume N discrete periods, with no discounting. Emissions in each period are $q_i \geq 0$, with $i = 1, \ldots, N$. Aggregate emissions are $Q_n = \sum_{i=0}^{n} q_i$. The benefits function is a sort of "settling up" at the end based on the total stock of emissions over the full N periods:

$$B(Q_N) = -\frac{b}{2} Q_N^2,$$

where $b > 0$ is a parameter. Using the simple sum of emissions is equivalent to assuming that there is no decay in the accumulated stock. Benefits would be maximized by setting $Q_N = 0$. Higher emissions lower the benefits by progressively larger amounts: $B_{Q_N}(0) = 0$, and $\forall\ Q_N > 0$ we have $B_{Q_N}(Q_N) = -b \cdot Q_N < 0$ and also $B_{Q_N Q_N}(Q_N) = -b < 0$, where $B_x(x)$ and $B_{xx}(x)$ are the first and the second derivative with respect to x. Costs are a function of per period emissions, and controlling emissions is costly. Per period cost as a function of emissions is written as,

$$C(q_i, \theta_i) = \frac{c}{2}(q_i - \bar{q})^2 - \theta_i(q_i - \bar{q}), \tag{11}$$

where $c > 0$ is a fixed parameter, $\bar{q}$ is a reference level of emissions and θ_i is a non-negative random variable.[2] Costs in a given period are minimized at the adjusted reference level $\bar{q} + (\theta_i/c)$. Emissions less than the adjusted reference level cost progressively more. $\forall q_i \leq \bar{q} + (\theta_i/c)$ we have $C_{q_i}(q_i, \theta_i) = c(q_i - \bar{q}) - \theta_i < 0$ and $C_{q_i q_i}(q_i, \theta_i) = c > 0$, where $C_x(x, y)$ and $C_{xx}(x, y)$ are the first and the second derivative, respectively, with respect to the first component.

We consider a discrete-time problem in which either a quantity or a price constraint is established at the start of each period, then the uncertain parameter for that period is realized, and then producers choose their action given the regulatory constraint. At the end of the period the realization of that period's parameter is common knowledge and

---

[2]The variable θ_i takes truly random variables only at time $i = 1$. For $i = 2, ..., N$, θ_i is either a given value, θ, or equal to the value taken at time $i = 1$, $\tilde{\theta}_1$.

a new quantity or price constraint must be set. We use the two-stage framework just described to construct a pair of very special cases with a simple uncertainty structure that dramatically reduces the complexity of the problem and yet nevertheless exposes the key feature to which we wish to call attention. In both cases, all of the true uncertainty is embedded in the value taken by the random cost parameter θ_i at the first period, i.e. $\tilde{\theta}_1$. This is stage one. The first case captures the situation in which the shock in period 1 is purely transitory, and the second case captures the situation in which the shock in period 1 is permanent. In the first case, the cost parameters for the remaining periods (stage two) are known ex ante –i.e., not uncertain– and therefore independent of the realization $\tilde{\theta}_1$. We assume the values are identical across years, $\theta_i = \theta$ for $i = 2, \ldots, N$. In the second case, the cost parameters for the remaining periods (stage two) exactly equal the realization of the first period cost parameter, so that resolution about the first period cost resolves all the uncertainty about future costs, $\theta_i = \tilde{\theta}_1$ for $i = 2, \ldots, N$.

We solve the model by backward induction. In both cases, whatever uncertainty existed has been resolved at the conclusion of the first stage. Therefore, the optimal level of emissions in every future period can be calculated and readily enforced. Since all of the remaining periods are identical in their cost functions, and since we have no discounting, the optimal level of emissions will be identical across these subsequent periods, $q_i^* = q^*$ for $i = 2, \ldots, N$. In the first case, these optimal outputs will be independent of the realization of θ_i at time $i = 1$, while in the second case they will be a function of $\tilde{\theta}_1$. In both cases they will be a function of the choice of first period emissions, q_1^*. Given these optimal outputs, we write the value function at the conclusion of period 1 as the (deterministic) sum of the benefits and the remaining costs:

$$V(q_1^*, \tilde{\theta}_1) = \max_{q(q_1^*, \tilde{\theta}_1)} B\Big(q_1^* + (N-1)q\Big) - (N-1)C(q, \theta).$$

The first period problem can be modeled as the maximization of the expected difference between this value function and the first period cost:

$$\max_{q_1(\cdot)} \mathbb{E}_{\tilde{\theta}_1}\Big[V(q_1(\cdot), \tilde{\theta}_1) - C(q_1(\cdot), \tilde{\theta}_1)\Big]. \tag{12}$$

We have written this generally, without being clear about whether the first period output is a function of the cost parameter. In the first best (in the absence of uncertainty), it clearly will be: q_1 is a function of the realization $\tilde{\theta}_1$ such that the marginal value and

marginal cost are equal for each realization of θ_i at $i = 1$. In the second best, à la Weitzman, either (i) q_1, or (ii) fix a price, p_1, must be set before observing the realization $\tilde{\theta}_1$. In (i) the output will not be a function of the cost parameter. In (ii) it will vary with the realization θ_i at $i = 1$, but not necessarily according to the first best optimal schedule.

Weitzman asked which was better, the quantity or price control, in a setting with just one period. We, too, focus on whether the quantity or price control is better for regulating output in this one period. But the problem is posed and evaluated in a multi-period context as demanded by the analysis of a stock pollutant and framed in a two-stage decision model. In stage one θ_i is uncertain. In stage two the realization of the cost parameter and the resulting first period emissions, q_1^*, has been made, and then optimal level of emissions in subsequent periods are chosen. Before solving our problem we first present in Figure 4 a graph like those that are often presented in expositing the difference between price and quantity controls – see, for example, the Stern Report (2006, Box 14.1) among many others. It contains a graph of the marginal benefit and the marginal cost of alternative levels of period 1 emissions. Recall that $q_i = q$ for $i = 2, \ldots, N$. The marginal benefit function graphed is:

$$\frac{\partial B(Q_N)}{\partial q_1} = -b\Big(q_1 + (N-1)q\Big), \tag{13}$$

where the value for q is taken as fixed and independent of q_1 and θ_1. The marginal cost function graphed is:

$$\frac{\partial C(q_1, \theta_1)}{\partial q_1} = c\Big(q_1 - \bar{q}\Big) - \theta_1. \tag{14}$$

Since it is most common in the literature to graph the marginal benefit and marginal cost of abatement, we have done so as well in Figure 4. Abatement is just the difference between actual emissions and some benchmark level of emissions.

Three separate cases of the marginal cost function are shown, corresponding to a high and low realization of θ_i at $i = 1$ and to the mean value: θ_1^H, θ_1^L, and θ_1^M. A high realization means a higher marginal cost of abatement (a lower marginal cost of emissions) and corresponds to the higher of the three lines. The quantity $\hat{q}_1$ corresponds to the intersection of the marginal benefit function with the marginal cost function for the mean value of θ_i at $i = 1$, θ_1^M.[3] Suppose that the government constrains period 1 emissions to this level, $\hat{q}_1$. If

---

[3]In a single period framework, $\hat{q}_1$ would be the optimal ex ante quantity constraint given the uncertainty and inability to directly condition on it. In the multi-period framework, this is not exactly correct, since the marginal benefit function as written above does not properly reflect the possibilities for adaptation in future periods to the realizations in the first period uncertainty. But this complication will not concern us here.
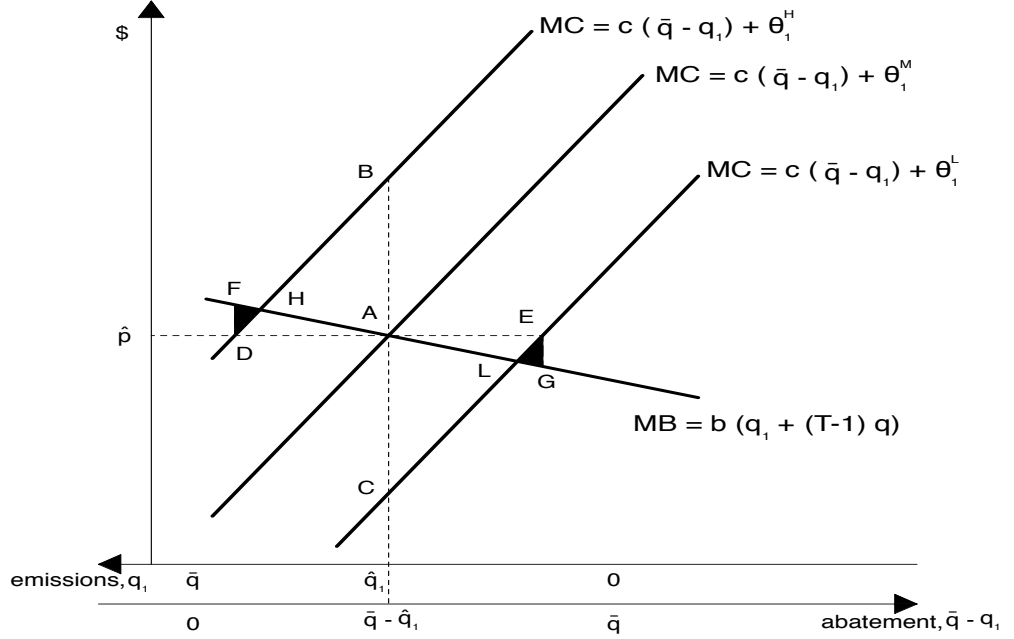
Figure 4: **Marginal costs and marginal benefits of abatement.** It is most common to show the marginal benefits and marginal costs as a function of abatement, which is the difference between emissions and a benchmark. This is equivalent to charting the negative of the marginal benefit and marginal cost functions from Equations (13) and (14), and reversing the direction of the horizontal axis, as is done here. The marginal benefit of abatement is decreasing, while the marginal cost of abatement is increasing. MC and MB represents marginal costs and marginal benefits, respectively.

the realized cost parameter is θ_1^H, then the economy will bear the marginal cost as marked by point B in the figure. The deadweight cost of producing the pre-specified quantity given this marginal cost curve is shown by the triangle ABH. If the realized cost parameter is θ_1^L, then the economy will bear the marginal cost as marked by point C in the figure. The deadweight cost of producing the pre-specified quantity given this marginal cost curve is shown by the other triangle ACL.

The price corresponding to this quantity constraint $\hat{q}_1$, and to the mean cost parameter, θ_1^M, is $\hat{p}$ which is marked on the figure. Suppose, that instead of fixing the quantity constraint, $\hat{q}_1$, the government had fixed the price to be $\hat{p}$. In that case, the quantity of

18

emissions would vary with each realization of θ_i at $i = 1$ as determined by the intersection of the price line and the marginal cost curve associated with the realization. These quantities are also shown in the figure. For the high realization of the cost parameter θ_1^H, the quantity of emissions corresponds to point D in the figure. For the low realization θ_1^L, the quantity of emissions corresponds to point E. The deadweight cost of producing the resulting quantity for the high and low realizations of the cost parameter are shown by the respective DFH and EGL triangles which are very, very small.

Clearly for this drawing of the graphs the solid black regions are smaller than the empty regions, so that the price control is preferred. Were the relative slopes of the marginal benefit and marginal cost functions reversed, quantity controls would be preferred. But, the argument goes, because greenhouse gases are a stock pollutant, the marginal benefit function is virtually flat and clearly less sharply sloped than the single period marginal cost function. Since the range of variation of output in a single year is small compared to the anticipated accumulation over the relevant horizon, the slope of the marginal benefit function must be nearly flat. In contrast, the marginal cost of adjusting emissions within the year curves sharply. The argument that it is better to regular a stock pollutant using a price control hinges firmly on this assumption of the different time scales: a steep rise of the marginal cost curve for a variation in emissions within a single year, and a gradual rise of the marginal benefit curve for this exact same quantity of emissions as a fraction of the centuries long total level of emissions creating the climate change problem.

The problem is that the marginal benefit function written in Equation (13) and shown in the figure is not the correct marginal benefit function for the first period optimization shown in Equation (12). The correct marginal benefit function is:

$$\frac{\partial B(Q_N)}{\partial q_1} = -b\Big(q_1 + (N-1)q^*(q_1, \theta_1)\Big) \cdot \Big(1 + (N-1)\frac{\partial q^*}{\partial q_1}\Big),$$

which recognizes as well how optimal outputs in the $N-1$ future years are set conditional on the first period cost realization and the first period choice of quantity. Therefore, the correct first order condition for the optimization is:

$$
\begin{aligned}
\frac{\partial V(q_1, \theta_1)}{\partial q_1} \;=\; & -b\Big(q_1 + (N-1)q^*(q_1, \theta_1)\Big)\Big(1 + (N-1)\frac{\partial q^*}{\partial q_1}\Big) \\
& -(N-1)c(q^*(q_1,\theta_1) - \bar{q})\frac{\partial q^*}{\partial q_1} + (N-1)\theta\frac{\partial q^*}{\partial q_1}.
\end{aligned}
\tag{15}
$$

19

The first order condition on q^* implies:

$$-b\Big(q_1+(N-1)q^*(q_1,\theta_1)\Big)(N-1)\frac{\partial q^*}{\partial q_1}-(N-1)c(q^*(q_1,\theta_1)-\bar{q})\frac{\partial q^*}{\partial q_1}+(N-1)\theta\frac{\partial q^*}{\partial q_1}=0, \quad (16)$$

so that by substituting (16) into (15) we have:

$$\frac{\partial V(q_1,\theta_1)}{\partial q_1} = -b\Big(q_1+(N-1)q^*(q_1,\theta_1)\Big). \quad (17)$$

The evaluation of Equation (17) depends upon the form of $q^*(q_1,\theta_1)$. For the first case, q^* is independent of the realization of θ_t at $t=1$, and – glossing over the dependence on q_1, which is likely to be small – Equation (17) reduces to Equation (13) so that Figure 4 is approximately correct.

However, for the second case Figure 4 is entirely inappropriate. In the second case the realization $\tilde{\theta}_1$ affects the cost functions in years $i=2,\ldots,N$, so that q^* is not fixed and independent of the cost parameter realization at $i=1$. It is not appropriate to ignore the dependence on $\tilde{\theta}_1$ as we ignored the dependence on q_1. Assuming that N is large, the output in a single year, q_1, will have a small impact on the optimal output in subsequent years. But the realization $\tilde{\theta}_1$ is a different sort of variable, which is why it is multiplied by $N-1$. The scale of the impact of a variation in θ_i at $i=1$ is of the very same order as the time scale of the stock pollutant problem. This is the crux of the problem. Therefore, it is not correct to show a single marginal benefit function in Figure 4. A change in the realization of the cost parameter $\tilde{\theta}_1$ actually shifts the marginal benefit function, and does so at a large scale! In the second case we have:

$$\frac{d}{d\theta_1}\Big(\frac{\partial V(q_1,\theta_1)}{\partial q_1}\Big) = -b(N-1)\frac{\partial q^*}{\partial\theta_1},$$

which is inconsistent with Figure 4. Figure 5 shows the actual situation for the second case. Even when the ceteris paribus marginal benefit function appears flat, the relevant relationship for comparing deadweight costs is not what this would seem to imply. Different realizations of the cost parameter change the presumed baseline emissions in later periods and therefore shift the marginal benefit function appropriate for evaluating a change in period 1 emissions. The effect is comparable to what Stavins (1996) illustrates in the case

of correlation between cost and benefit uncertainty. It is entirely possible that a quantity control is preferred, despite the apparently flat marginal benefit function.
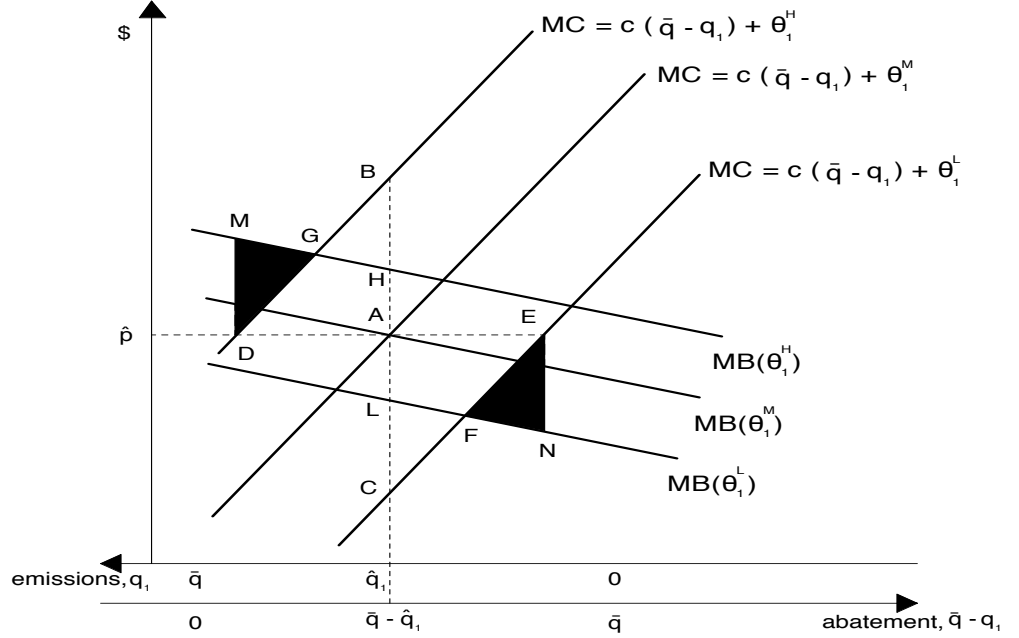

Figure 5: **Marginal costs and marginal benefits for the second case.** MC and MB represents marginal costs and marginal benefits, respectively.

In the second case, the preference for quantity or price controls depends upon the relative steepness of the marginal benefit function against the steepness of the marginal *aggregate cost* function. Equation (11) is a per period cost function. The aggregate cost function, $D(Q_N)$, is the result of allocating total emissions efficiently across years:

$$D(Q_N) = \min_{q_1,\ldots,q_N} \sum_{i=1}^{N} c(q_i, \theta_i),$$

where $\sum_{i=1}^{N} q_i = Q_N$.

Therefore in the second case the argument about stock pollutants loses its force entirely.

21

There is no basis for arguing that the marginal cost is necessarily more sharply sloping than the marginal benefit function. In the case of greenhouses gases, the assessment of these aggregate benefit and aggregate cost functions is itself a matter of great uncertainty and debate, see Hepburn (2005) and references therein.

4 Conclusions

We constructed a pair of simple examples that help clarify the role of uncertainty in the choice of a price or quantity instrument for controlling a stock pollutant. Our contribution is the simplicity of the model and the sharpness of our results. In particular, we fully characterize the nature of the dynamic cost effective emission policy responding to shocks to the cost of abatement, whether those shocks temporarily change the cost or permanently change it. If the shocks to cost are exclusively temporary, then it is cost effective to adjust the quantity period-by-period, keeping the marginal cost relatively constant. If the shocks are exclusively permanent, then it is cost effective to keep the quantity fixed, letting the marginal cost vary as the cost varies.

This result provides insight into the optimal instrument choice for a regulator attempting to control emissions by private agents with better information on costs, although we do not explicitly model this problem. In the case of temporary uncertainty, when the dynamic cost effective emissions policy is to adjust the quantity period-by-period, keeping the marginal cost relatively constant, it is most likely optimal to employ a price instrument. However, in the case of permanent uncertainty, a price instrument will clearly not be optimal since it is price that ought to absorb all of the shocks to cost. Instead, since quantity should be invariant to the cost shocks period-by-period, a quantity instrument is likely to be optimal. Our results, therefore, reinforce the findings of Newell and Pizer (2003) and Karp and Zhang (2005) that the choice between using a quantity or price instrument hinges on whether or not the uncertainty in costs is correlated across periods.

Our model also provides insight into the operation of a cap-and-trade system that allows banking and borrowing of allowances across periods. The dynamic cost effective emissions policy we derived is also an optimal dynamic allocation of allowances across the periods included under a cap. Therefore, cap-and-trade with banking and borrowing implements the dynamically cost effective emission policy, regardless of the sort of uncertainty. If the cap-and-trade system faces temporary uncertainty in costs, then it will be the period-by-

period quantity of emissions that will fluctuate under the cap-and-trade system, and the price will be relatively constant. If the cap-and-trade system faces permanent uncertainty in costs, then it will be the period-by-period price that will fluctuate under the cap-and-trade system, and the quantity of emissions in each period will not be stochastic, but rise deterministically at the rate of growth in costs less the interest rate.
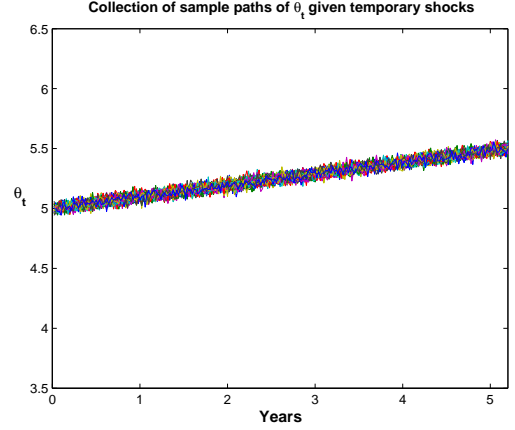
Appendix: Temporary & Permanent Shocks

To grasp the difference between temporary and permanent uncertainty, it may help to observe how shocks can affect the evolution of the cost parameter θ_i. Figure A-1 shows a simulation of the cost parameter driven by temporary uncertainty and can be contrasted with Figure A-2 which shows a simulation of the cost parameter driven by permanent uncertainty. For this pair of simulations we generate a single set of sample paths for the random errors, $\epsilon_i, i = 1, \ldots, N$. We use this one set of random errors to generate a set of sample paths for the cost parameter that follows Equation (1), and to generate a set of sample paths for the cost parameter that follows Equation (2). We use the same initial cost parameter, θ_0, and the same volatility parameter, σ. Also, we set the two drift parameters so that the expected cost at the conclusion of the simulation are also approximately the same: this requires that $\mu = \nu - 1/2\sigma^2$. Therefore, both sets of sample paths for the cost parameter have the exact same absolute volatility within each single period. However, the volatility's impact along a sample path is different, and the simulation helps one to visualize this difference. Figure A-1(a) shows a single sample path of the cost parameter when θ_i follows Equation (1). Figure A-1(b) shows the set of sample paths. Because each shock has a purely transitory impact on the cost parameter, successive values of the parameter are close to the original forecasted value, and vary from it only by the size of the most recent shock. Therefore the path of the cost parameter is tight around the forecasted path and remains tight at all horizons. The confidence interval for a forecast of the cost parameter is constant at every forecasting horizon. Figure A-2(a) shows the corresponding single sample path of the cost parameter when θ_i follows Equation (2). Figure A-2(b) shows the corresponding set of sample paths. Because each shock has a purely permanent impact on the cost parameter, successive values of the parameter may wander further and further from the original forecasted value. Therefore the confidence interval for a forecast of the cost parameter grows with the forecast horizon. The contrast between Figure A-1(b) and Figure A-2(b) is the critical points of contrast between purely temporary and purely

permanent shocks in this paper.
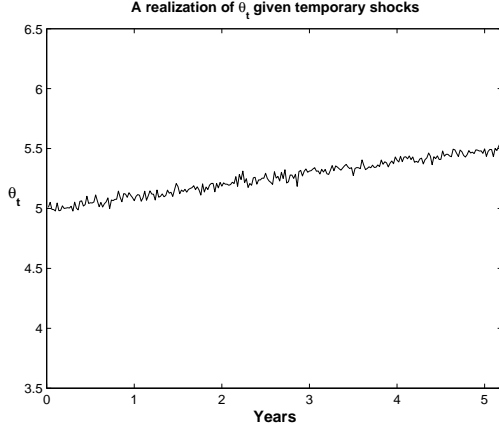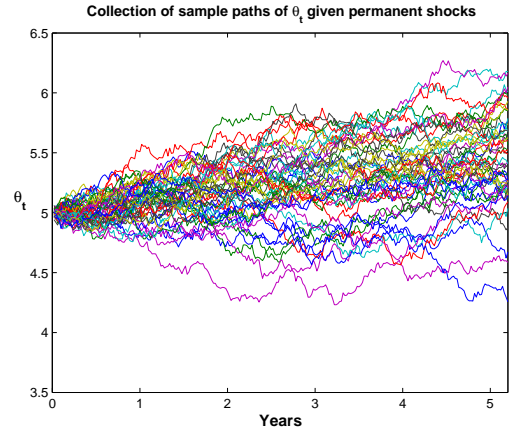

(a) A path given purely temporary shocks

(b) Collection of sample paths given purely temporary shocks

Figure A-1: 5-year evolution of θ_i given purely temporary shocks. In this example, $\theta_0 = 5$, $\nu = 0.10$, $\sigma = 0.2$, and the time step corresponds to a week.

(a) A path given purely permanent shocks

(b) Collection of sample paths given purely permanent shocks

Figure A-2: 5-year evolution of θ_i given purely permanent shocks. In this example, $\theta_0 = 5$, $\nu = 0.10$, $\sigma = 0.2$, and the time step corresponds to a week.

Appendix

We solve the optimal pollution control problem using recursive substitution. For the sake of exposure, we count periods backwards from the endpoint using the index j to denote

periods from the endpoint, with $j = N, ..., 2, 1$. Recall also that the allowed emissions remaining in the subsequent period is a function of the emissions chosen in the current period, $\bar{q}_{j-1} := \bar{q}_j - q_j$. We begin by solving the certainty case, since this provides useful intuition for the uncertainty cases. We then solve the uncertainty case when per period's shock is purely temporary and purely permanent, respectively.

Certainty case

When $\sigma = 0$, we have the certainty case and the cost parameter follows the dynamics:

$$\theta_{j-1} = \theta_j + \nu = \theta_0 + (N - j + 1)\nu.$$

Solving for $j = 1$, we have $q_1^*(\bar{q}_1, \theta_1) = \bar{q}_1$. Therefore, the value function in the last period $V_1^*(\bar{q}_1, \theta_1) = c(\bar{q}_1, \theta_1) = e^{\theta_1 - q_1^*}$. For $j = 2$, we have

$$
\begin{aligned}
V_2(\bar{q}_2, q_2, \theta_2) &= \mathbb{E}_{\theta_2}\left[c(q_2, \theta_2) + e^{-r}V_1^*(\bar{q}_1(\bar{q}_2, q_2), \theta_1)\right] \\
&= e^{\theta_2}\left[e^{-q_2} + e^{-(r-\nu)}e^{-(\bar{q}_2 - q_2)}\right].
\end{aligned}
$$

Solving the first order condition for the cost minimizing emissions gives us:

$$q_2^* = \frac{1}{2}\bar{q}_2 - \frac{1}{2}(\nu - r). \tag{A-1}$$

Substituting this into the value function gives the optimized value function:

$$V_2^*(\bar{q}_2, \theta_2) = 2e^{\theta_2 - q_2^*} = 2e^{\theta_2 - \frac{1}{2}\bar{q}_2 + \frac{1}{2}(\nu - r)}. \tag{A-2}$$

For $j = 3$, we have

$$
\begin{aligned}
V_3(\bar{q}_3, q_3, \theta_3) &= \mathbb{E}_{\theta_3}\left[c(q_3, \theta_3) + e^{-r}V_2^*(\bar{q}_2(\bar{q}_3, q_3), \theta_2)\right] \\
&= e^{\theta_3}\left[e^{-q_3} + 2e^{-\frac{\bar{q}_3 - q_3 - 3\nu + 3r}{2}}\right].
\end{aligned}
$$

Solving the first order condition for the cost minimizing emissions gives us:

$$q_3^* = \frac{1}{3}\bar{q}_3 - \frac{2}{2}(\nu - r)$$

Substituting this into the value function gives the optimized value function:

$$V_3^*(\bar{q}_3, \theta_3) = 3e^{\theta_3 - q_3^*} = 3e^{\theta_3 - \frac{1}{3}\bar{q}_3 + (\nu - r)}. \tag{A-3}$$

For $j = 4$, we have

$$
\begin{aligned}
V_4(\bar{q}_4, q_4, \theta_4) &= \mathbb{E}_{\theta_4}\Big[c(q_4, \theta_4) + e^{-r}V_3^*(\bar{q}_3(\bar{q}_4, q_4), \theta_3)\Big] \\
&= e^{\theta_4}\Big[e^{-q_4} + 3e^{-\frac{\bar{q}_4 - q_4 - 6\nu + 6r}{3}}\Big].
\end{aligned}
$$

Solving the first order condition for the cost minimizing emissions gives us:

$$q_4^* = \frac{1}{4}\bar{q}_4 - \frac{3}{2}(\nu - r).$$

Substituting this into the value function gives the optimized value function:

$$V_4^*(\bar{q}_4, \theta_4) = 4e^{\theta_4 - q_4^*} = 4e^{\theta_4 - \frac{1}{4}\bar{q}_4 + \frac{3}{2}(\nu - r)} \tag{A-4}$$

Continuing the substitution iteratively, we obtain the general form of the optimal dynamic policy:

$$q_j^* = \frac{1}{j}\bar{q}_j - \frac{1}{2}(j - 1)(\nu - r). \tag{A-5}$$

and in calendar period i, it is:

$$q_i^* = \frac{1}{N - i + 1}\bar{q}_i - \frac{1}{2}(N - i)(\nu - r). \tag{A-6}$$

The general form of the optimized value function is:

$$V_j^*(\bar{q}_j, \theta_j) = je^{\theta_j - q_j^*} = je^{\theta_j - \frac{1}{j}\bar{q}_j + \frac{1}{2}(j-1)(\nu - r)} \tag{A-7}$$

and in calendar period i, it is:

$$V_i^*(\bar{q}_i, \theta_i) = ie^{\theta_i - q_i^*} = ie^{\theta_i - \frac{1}{N-i+1}\bar{q}_i + \frac{1}{2}(N-i)(\nu - r)} \tag{A-8}$$

Temporary shock case

When the per period's shock is temporary, the cost parameter follows the dynamics:

$$\theta_j = \Theta_j + \sigma\epsilon_j, \quad \text{where} \quad \Theta_j \equiv \Theta_{j+1} + \nu.$$

Solving for $j = 1$, we have $q_1^*(\bar{q}_1, \theta_1) = \bar{q}_1$. Therefore, the value function in the last period $V_1^*(\bar{q}_1, \theta_1) = c(\bar{q}_1, \theta_1) = e^{\theta_1 - q_1^*}$. For $j = 2$, we have

$$\begin{aligned}
V_2(\bar{q}_2, q_2, \theta_2) &= \mathbb{E}_{\theta_2}\left[c(q_2, \theta_2) + e^{-r}V_1^*(\bar{q}_1(\bar{q}_2, q_2), \theta_1)\right] \\
&= \left[e^{\theta_2 - q_2} + e^{-r}\mathbb{E}_{\theta_2}\left[e^{\theta_1 - (\bar{q}_2 - q_2)}\right]\right] \\
&= e^{\Theta_2}\left[e^{\sigma\epsilon_2 - q_2} + e^{-(r-\nu-\frac{\sigma^2}{2})}e^{-(\bar{q}_2 - q_2)}\right]
\end{aligned}$$

Solving the first order condition for the cost minimizing emissions gives us:

$$q_2^* = \frac{1}{2}\bar{q}_2 - \frac{1}{2}(\nu - r) - \frac{1}{4}\sigma^2 + \frac{1}{2}\sigma\epsilon_2 \tag{A-9}$$

Substituting this into the value function gives the optimized value function:

$$V_2^*(\bar{q}_2, \theta_2) = 2e^{\theta_2 - q_2^*} = 2e^{\theta_2 - \frac{1}{2}\bar{q}_2 + \frac{1}{2}(\nu - r) + \frac{\sigma^2}{4} - \frac{1}{2}\sigma\epsilon_2} \tag{A-10}$$

For $j = 3$, we have

$$
\begin{aligned}
V_3(\bar{q}_3, q_3, \theta_3) &= \mathbb{E}_{\theta_3}\Big[c(q_3, \theta_3) + e^{-r}V_2^*(\bar{q}_2(\bar{q}_3, q_3), \theta_2)\Big] \\
&= \Big[e^{\theta_3 - q_3} + e^{-r}\mathbb{E}_{\theta_3}\big[2e^{\theta_2 - \frac{\bar{q}_3 - q_3 - \nu + r - \sigma^2/2 + \sigma\epsilon_2}{2}}\big]\Big] \\
&= e^{\Theta_3}\Big[e^{\sigma\epsilon_3 - q_3} + 2e^{-\frac{\bar{q}_3 - q_3 - 3\nu + 3r - 3\sigma^2/4}{2}}\Big]
\end{aligned}
$$

Solving the first order condition for the cost minimizing emissions gives us:

$$
q_3^* = \frac{1}{3}\bar{q}_3 - \frac{2}{2}(\nu - r) - \frac{1}{4}\sigma^2 + \frac{2}{3}\sigma\epsilon_3
$$

Substituting this into the value function gives the optimized value function:

$$
V_3^*(\bar{q}_3, \theta_3) = 3e^{\theta_3 - q_3^*} = 3e^{\theta_3 - \frac{1}{3}\bar{q}_3 + (\nu - r) + \frac{1}{4}\sigma^2 - \frac{2}{3}\sigma\epsilon_3} \tag{A-11}
$$

For $j = 4$, we have

$$
\begin{aligned}
V_4(\bar{q}_4, q_4, \theta_4) &= \mathbb{E}_{\theta_4}\Big[c(q_4, \theta_4) + e^{-r}V_3^*(\bar{q}_3(\bar{q}_4, q_4), \theta_3)\Big] \\
&= \Big[e^{\theta_4 - q_4} + e^{-r}\mathbb{E}_{\theta_4}\big[3e^{\theta_3 - \frac{\bar{q}_4 - q_4 - 3\nu + 3r - 3\sigma^2/4 - 2\sigma\epsilon_3}{3}}\big]\Big] \\
&= e^{\Theta_4}\Big[e^{\sigma\epsilon_4 - q_4} + 3e^{-\frac{\bar{q}_4 - q_4 - 6\nu + 6r - 11\sigma^2/12}{3}}\Big]
\end{aligned}
$$

Solving the first order condition for the cost minimizing emissions gives us:

$$
q_4^* = \frac{1}{4}\bar{q}_4 - \frac{3}{2}(\nu - r) - \frac{11}{48}\sigma^2 + \frac{3}{4}\sigma\epsilon_4
$$

Substituting this into the value function gives the optimized value function:

$$
V_4^*(\bar{q}_4, \theta_4) = 4e^{\theta_4 - q_4^*} = 4e^{\theta_4 - \frac{1}{4}\bar{q}_4 + \frac{3}{2}(\nu - r) + \frac{11}{48}\sigma^2 - \frac{3}{4}\sigma\epsilon_4} \tag{A-12}
$$

Continuing the substitution, the general form of the optimal dynamic policy is:

$$q_j^* = \frac{1}{j}\bar{q}_j - A_j\sigma^2 - \frac{1}{2}(j-1)(\nu - r) + \frac{j-1}{j}\sigma\epsilon_j \qquad \text{(A-13)}$$

where
$$A_j = \frac{j-1}{j}\left(A_{j-1} + \frac{1}{2(j-1)^2}\right) \text{ for } j = 2,\ldots,N, \text{ and } A_1 = 0. \qquad \text{(A-14)}$$

Rewriting in calendar period i, we have:

$$q_i^* = \frac{1}{N-i+1}\bar{q}_i - A_i\sigma^2 - \frac{1}{2}(N-i)(\nu - r) + \frac{N-i}{N-i+1}\sigma\epsilon_j \qquad \text{(A-15)}$$

where

$$A_i = \frac{N-i}{N-i+1}\left(A_{i+1} + \frac{1}{2(N-i)^2}\right) \text{ for } i = 1,\ldots,N-1, \text{ and } A_N = 0. \qquad \text{(A-16)}$$

The general form of the optimized value function is:

$$V_j^*(\bar{q}_j) = je^{\theta_j - q_j^*} = je^{\theta_j - \frac{1}{j}\bar{q}_j + A_j\sigma^2 + \frac{1}{2}(j-1)(\nu-r) - \frac{j-1}{j}\sigma\epsilon_j} \qquad \text{(A-17)}$$

and in calendar time:

$$V_i^*(\bar{q}_i) = ie^{\theta_i - q_i^*} = ie^{\theta_i - \frac{1}{N_i+1}\bar{q}_i + A_i\sigma^2 + \frac{1}{2}(N-i)(\nu-r) - \frac{N-i}{N-i+1}\sigma\epsilon_i}. \qquad \text{(A-18)}$$

Permanent shock case

When the per period's shock is permanent, the cost parameter follows the dynamics

$$\theta_i = \theta_{i-1} + \mu + \sigma\epsilon_i.$$

Solving for $j = 1$, we have $q_1^*(\bar{q}_1, \theta_1) = \bar{q}_1$. Therefore, the value function in the last period $V_1^*(\bar{q}_1, \theta_1) = c(\bar{q}_1, \theta_1) = e^{\theta_1 - q_1^*}$. For $j = 2$, we have

$$
\begin{aligned}
V_2(\bar{q}_2, q_2, \theta_2) &= \mathbb{E}_{\theta_2}\left[c(q_2, \theta_2) + e^{-r}V_1^*(\bar{q}_1(\bar{q}_2, q_2), \theta_1)\right] \\
&= \left[e^{\theta_2 - q_2} + e^{-r}\mathbb{E}_{\theta_2}\left[e^{\theta_1 - (\bar{q}_2 - q_2)}\right]\right] \\
&= e^{\theta_2}\left[e^{-q_2} + e^{-(r - \mu + \frac{1}{2}\sigma^2)}e^{-(\bar{q}_2 - q_2)}e^{\frac{1}{2}\sigma^2}\right].
\end{aligned}
$$

Solving the first order condition for the cost minimizing emissions gives us:

$$
q_2^* = \frac{1}{2}\bar{q}_2 - \frac{1}{2}(\mu - r). \tag{A-19}
$$

Substituting this into the value function gives the optimized value function:

$$
V_2^*(\bar{q}_2, \theta_2) = 2e^{\theta_2 - q_2^*} = 2e^{\theta_2 - \frac{1}{2}\bar{q}_2 + \frac{1}{2}(\mu - r)}. \tag{A-20}
$$

For $j = 3$, we have

$$
\begin{aligned}
V_3(\bar{q}_3, q_3, \theta_3) &= \mathbb{E}_{\theta_3}\left[c(q_3, \theta_3) + e^{-r}V_2^*(\bar{q}_2(\bar{q}_3, q_3), \theta_2)\right] \\
&= \left[e^{\theta_3 - q_3} + e^{-r}\mathbb{E}_{\theta_3}\left[2e^{\theta_2 - \frac{\bar{q}_3 - q_3 - \mu + r}{2}}\right]\right] \\
&= e^{\theta_3}\left[e^{-q_3} + 2e^{-\frac{\bar{q}_3 - q_3 - 3\mu + 3r}{2}}\right].
\end{aligned}
$$

Solving the first order condition for the cost minimizing emissions gives us:

$$
q_3^* = \frac{1}{3}\bar{q}_3 - \frac{2}{2}(\mu - r)
$$

Substituting this into the value function gives the optimized value function:

$$
V_3^*(\bar{q}_3, \theta_3) = 3e^{\theta_3 - q_3^*} = 3e^{\theta_3 - \frac{1}{3}\bar{q}_3 + (\mu - r)}. \tag{A-21}
$$

For $j = 4$, we have

$$
\begin{aligned}
V_4(\bar{q}_4, q_4, \theta_4) &= \mathbb{E}_{\theta_4}\left[c(q_4, \theta_4) + e^{-r} V_3^*(\bar{q}_3(\bar{q}_4, q_4), \theta_3) \right] \\
&= \left[e^{\theta_4 - q_4} + e^{-r} \mathbb{E}_{\theta_4}\left[3 e^{\theta_3 - \frac{\bar{q}_4 - q_4 - 3\mu + 3r}{3}} \right] \right] \\
&= e^{\theta_4}\left[e^{-q_4} + 3 e^{-\frac{\bar{q}_4 - q_4 - 6\mu + 6r}{3}} \right].
\end{aligned}
$$

Solving the first order condition for the cost minimizing emissions gives us:

$$
\begin{aligned}
q_4^* &= \frac{3}{4}\frac{\bar{q}_4 - 6\mu + 6r}{3} \\
&= \frac{1}{4}\bar{q}_4 - \frac{3}{2}(\mu - r).
\end{aligned}
\tag{A-22}
$$

Substituting this into the value function gives the optimized value function:

$$
V_4^*(\bar{q}_4, \theta_4) = 4 e^{\theta_4 - q_4^*} = 4 e^{\theta_4 - \frac{1}{4}\bar{q}_4 + \frac{3}{2}(\mu - r)}
\tag{A-23}
$$

Continuing the substitution, we obtain the general form of the optimal dynamic policy is:

$$
q_j^* = \frac{1}{j}\bar{q}_j - \frac{1}{2}(j - 1)(\mu - r).
\tag{A-24}
$$

and in calendar time i, it is:

$$
q_i^* = \frac{1}{N - i + 1}\bar{q}_i - \frac{1}{2}(N - i)(\mu - r).
\tag{A-25}
$$

The general form of the optimized value function is:

$$
V_j^*(\bar{q}_j, \theta_j) = j e^{\theta_j - q_j^*} = j e^{\theta_j - \frac{1}{j}\bar{q}_j + \frac{1}{2}(j - 1)(\mu - r)}
\tag{A-26}
$$

that in calendar time i is:

$$
V_i^*(\bar{q}_i, \theta_i) = i e^{\theta_i - q_i^*} = i e^{\theta_i - \frac{1}{N - i + 1}\bar{q}_i + \frac{1}{2}(N - i)(\mu - r)}.
\tag{A-27}
$$

31

References

Hepburn, C. (2005). Regulating by prices, quantities or both: A Review of Instrument Choice. *Oxford Review of Economic Policy*, 22(2):226–247.

Hoel, M. and Karp, L. (2002). Taxes versus Quotas for a stock pollutant. *Resource and Energy Economics*, 24:367–384.

Karp, L. and Zhang, J. (2005). Regulation of Stock Externalities with Correlated Abatement Costs. *Environmental & Resource Economics*, 32:273–299.

Newell, R. G. and Pizer, A. W. (2003). Regulating stock externalities under uncertainty. *Journal of Environmental Economics and Management*, 45:416–432.

Nordhaus, W. D. (1994). *Managing the Global Commons*. MIT Press, Cambridge.

Pizer, W. A. (1999). The Optimal Choice of Policy in the Presence of Uncertainty. *Resource and Energy Economics*, 21:255–287.

Stavins, R. (1996). Correlated uncertainty and policy instrument choice. *Journal of Environmental Economics and Management*, 30:218–232.

Stern, N. (2006). *The Economics of Climate Change. The Stern Review*. Cambridge University Press, Cambridge - UK.

Weitzman, M. L. (1974). Prices vs. Quantities. *Review of Economic Studies*, 41(4):683–691.