



Stocks and shocks: a clarification in the debate over price vs. quantity controls for greenhouse gases

John E. Parsons and Luca Taschini

March 2011

**Centre for Climate Change Economics and Policy
Working Paper No. 54**

**Grantham Research Institute on Climate Change and
the Environment
Working Paper No. 43**

The Centre for Climate Change Economics and Policy (CCCEP) was established by the University of Leeds and the London School of Economics and Political Science in 2008 to advance public and private action on climate change through innovative, rigorous research. The Centre is funded by the UK Economic and Social Research Council and has five inter-linked research programmes:

1. Developing climate science and economics
2. Climate change governance for a new global deal
3. Adaptation to climate change and human development
4. Governments, markets and climate change mitigation
5. The Munich Re Programme - Evaluating the economics of climate risks and opportunities in the insurance sector

More information about the Centre for Climate Change Economics and Policy can be found at: <http://www.cccep.ac.uk>.

The Grantham Research Institute on Climate Change and the Environment was established by the London School of Economics and Political Science in 2008 to bring together international expertise on economics, finance, geography, the environment, international development and political economy to create a world-leading centre for policy-relevant research and training in climate change and the environment. The Institute is funded by the Grantham Foundation for the Protection of the Environment, and has five research programmes:

1. Use of climate science in decision-making
2. Mitigation of climate change (including the roles of carbon markets and low-carbon technologies)
3. Impacts of, and adaptation to, climate change, and its effects on development
4. Governance of climate change
5. Management of forests and ecosystems

More information about the Grantham Research Institute on Climate Change and the Environment can be found at: <http://www.lse.ac.uk/grantham>.

This working paper is intended to stimulate discussion within the research community and among users of research, and its content may have been submitted for publication in academic journals. It has been reviewed by at least one internal referee before publication. The views expressed in this paper represent those of the author (s) and do not necessarily represent those of the host institutions or funders.

Stocks & Shocks: A Clarification in the Debate Over Price vs. Quantity Controls for Greenhouse Gases*

John E. Parsons^{a†} Luca Taschini^{b‡}

^a*Massachusetts Institute of Technology*

^b*London School of Economics*

March 12, 2011

Abstract

We construct two simple examples that help to clarify the role of a key assumption in the analysis of price or quantity controls of greenhouse gases in the presence of uncertain costs. Traditionally much has been made of the fact that greenhouse gases are a stock pollutant, and that therefore the marginal benefit curve must be relatively flat. This fact is said to establish the preference of a price control over a quantity control. The stock pollutant argument is considered dispositive, so that the preference for price controls is categorical. We show that this argument can only be true if the uncertainty about cost is a special form: all shocks are transitory. We show that in the case of permanent shocks, the traditional comparison of marginal benefits vs. marginal costs is mis-measured. The choice between quantity and price controls becomes ambiguous again and depends upon a more difficult measurement of marginal costs and benefits. The simplicity of the examples and the solutions is a major element of the contribution here. The examples are readily accessible and the comparison of results under the alternative assumptions of transitory and permanent shocks is stark.

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

JEL Classifications: H23, Q28, Q50, Q58.

*Part of Parsons' research was supported by the MIT Center for Energy and Environmental Policy Research and the MIT Joint Program on the Science and Policy of Global Change and by a grant from the Doris Duke Charitable Foundation. Part of Taschini's research was supported by the Grantham Research Institute on Climate Change and by the Centre for Climate Change Economics and Policy, which is funded by the UK Economic and Social Research Council (ESRC) and Munich Re. This draft has benefited from helpful comments from Gib Metcalf, John Reilly, Jake Jacoby and Mort Webster and participants in the MIT Emissions Prediction and Policy Analysis seminar.

[†]Address: Sloan School of Management, MIT, E19-411, 77 Massachusetts Ave., Cambridge, MA 02139 USA. E-mail: jparsons@mit.edu

[‡]Address: The Grantham Research Institute on Climate Change and the Environment, London School of Economics, UK. E-mail: l.taschini1@lse.ac.uk

1 Introduction

Two classic alternatives for regulating 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. In particular, a number of authors make much of our limited knowledge about the impacts of greenhouse gases and the costs of controlling them, and argue that this uncertainty, combined with a special feature of the greenhouse gas problem, leads categorically to the preference of price over quantity controls. In making this argument, these authors lean on the result of Weitzman (1974) that uncertainty about the underlying parameters on benefits and costs undermines the simple duality between price and quantity controls. Quantity control is preferred when the marginal benefits from control are sharply sloping as compared against the marginal costs, while price control is the preferred device when marginal benefits are relatively flat and marginal costs are sharply sloping. These authors then point to the fact that greenhouse gases are stock pollutants. For a stock pollutant, they argue, the marginal benefit function is flat, while the marginal cost function slopes sharply. Hence the clear superiority of a carbon tax over a cap & trade program. See, for example, Nordhaus (1994, Ch. 8, fn. 4), Hoel and Karp (2002), Newell and Pizer (2003), and Newell and Pizer (2006). This claim is widely repeated in policy papers analyzing the economics of controlling greenhouse gases, for example, the Stern Report (2006).

This paper challenges that logic. The original stock pollutant story oversimplifies the problem. In imagining an extension of the Weitzman model from a single period to a multiple one, a strong implicit assumption is made 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. We explicate the key role of the assumption about the type of uncertainty, thereby restoring the empirical burden imposed by the original Weitzman result. The relevant issues remain the shape of the marginal benefit and marginal cost functions. We show that the dynamic issues created by a stock pollutant only complicate the empirical issues at hand in defining and measuring the relevant marginal benefits and marginal costs. The mere fact that greenhouse gases are a stock pollutant is in no way dispositive in favor of price controls.

Other authors have explicitly addressed the dynamic structure of the uncertainty, notably Hoel and Karp (2002), Newell and Pizer (2003), and Karp and Zhang (2005). However, we believe this paper helps to clarify the key issue at hand as well as the limits of some of their results. Our contribution is to develop two pairs of very simple examples that draw out the issue in sharp resolution and clearly expose what is wrong in the oversimplified rhetoric of the stock pollutant argument. Because the examples are readily accessible, they help clarify elements of the dynamic problem that have remained obscure in the existing literature.

In the next section of the paper, we present a mathematically simple model of dynamic abatement cost uncertainty and which can be easily solved for the cost minimizing emissions policy to achieve a given cap on aggregate emissions. The model can be easily 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. In the case of temporary uncertainty, we show that the cost minimizing emissions policy has a high instantaneous volatility in emissions. Cost minimizing emissions are approximately white noise on a linear path, mimicking the stochastic features of the cost parameter. In contrast, along the optimal emissions path, the instantaneous volatility in the shadow price of a unit of emission is small. Therefore the cost minimizing policy can be approximately implemented by a price control, but not by a quantity flow control. In the case of permanent uncertainty, the cost minimizing emissions policy is entirely deterministic i.e., it is unaffected by the evolution of the cost parameter. The shadow price of abatement is volatile, however, mimicking the stochastic features of the cost parameter. The cost minimizing emissions policy can be implemented by a quantity flow constraint, but not by a price control.

In the third section, we address the problem of balancing the cost of abatement against the benefits of abatement. We do this using two stylized examples –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 green-

house 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 $\{i, i = 1, 2, \dots, N\}$. 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 ,

$$c_i(q_i, \theta_i) = e^{\theta_i - q_i}. \quad (1)$$

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

$$\frac{\partial c_i(q_i, \theta_i)}{\partial q_i} = -e^{\theta_i - q_i}. \quad (2)$$

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. To gain the full advantage of the functional form, we allow emissions to be negative at a given instant in time, but the intuition behind the main results obviously extends to the more realistic case in which these are bounded below at zero.

We evaluate two contrasting specifications of cost uncertainty. 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:

$$\theta_i = \theta_0 + i\nu + \sigma\epsilon_i, \quad (3)$$

where θ_0 is the starting cost parameter, ν is the constant expected growth rate in the mean cost parameter, and $\epsilon_i, i = 1, 2, \dots, N$, 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 mean within the period. This defines the temporary character of the shocks.

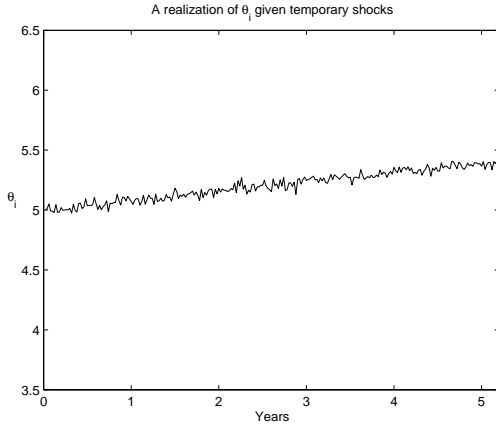
The second specification of the cost parameter θ_i is:

$$\theta_i = \theta_{i-1} + \mu + \sigma\epsilon_i. \quad (4)$$

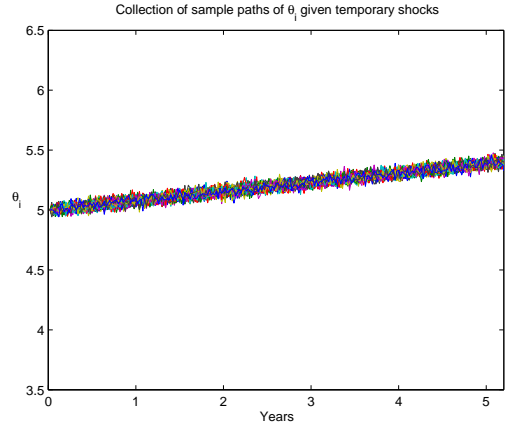
where μ is the constant expected growth in the cost parameter. This process is a random walk with trend. It is often said that the random walk has “infinite memory”. It is in this sense that we say the shocks have a permanent impact on the cost parameter.

To grasp the difference between the two cases, it may help to observe a simulation of the process in each case. For this simulation we generate a set of sample paths for the random errors, $\epsilon_i, i = 1, \dots, N$. We use this one set of random errors to generate a set of sample paths for the cost parameter that follows Equation (3), and to generate a set of sample paths for the cost parameter that follows Equation (4). 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 volatilitys impact along a sample path is different, and the simulation helps one to visualize this difference. Figure 1(a) shows a single sample path of the cost parameter when θ_i follows Equation (3). Figure 1(b) shows the set of sample paths. Because each shocks 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 2(a) shows the

corresponding single sample path of the cost parameter when θ_i follows Equation (4). Figure 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 1(b) and Figure 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

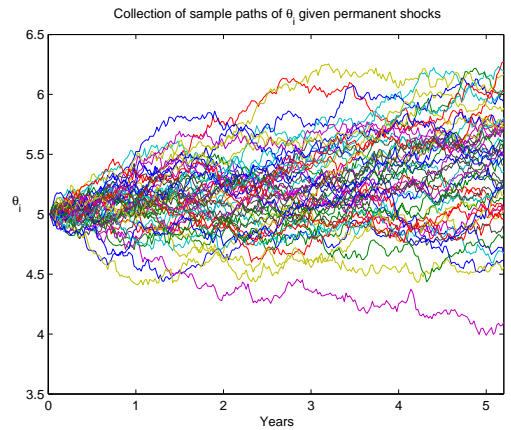


(b) Collection of sample paths given purely temporary shocks

Figure 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 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.

We assume a fixed aggregate emissions constraint, \bar{q} , over the N periods so that

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

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 (3), 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 (4), 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 programming. 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 $i = 1, \dots, N - 1$. 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 $i = 1, \dots, N$, as V_i . In the final period when $i = N$, it is simply the total cost of remaining emissions to abate:

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

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:

$$q_N^*(\bar{q}_N, \theta_N) = \bar{q}_N. \quad (5)$$

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:

$$\begin{aligned} V_N^*(\bar{q}_N, \theta_N) &\equiv V_N\left(\bar{q}_N, q_N^*(\bar{q}_N, \theta_N), \theta_N\right) \\ &= c\left(q_N^*(\bar{q}_N, \theta_N), \theta_N\right) = c(\bar{q}_N, \theta_N). \end{aligned} \quad (6)$$

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

$$p_N^*(\bar{q}_N, \theta_N) \equiv -\frac{\partial c(q_N^*(\bar{q}_N, \theta_N), \theta_N)}{\partial q_N} = c(q_N^*(\bar{q}_N, \theta_N), \theta_N). \quad (7)$$

where $p_i^*(\bar{q}_i, \theta_i)$ represents the 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:

$$V_i(\bar{q}_i, q_i, \theta_i) \equiv c(q_i(\bar{q}_i, \theta_i), \theta_i) + \mathbb{E}_{\theta_i} \left[V_{i+1}^*(\bar{q}_{i+1}(\bar{q}_i, q_i), \theta_{i+1}) \right] \quad \text{for } 1 \leq i < N.$$

The allowed emissions remaining in the subsequent period is, of course, a function of the emissions chosen in the current period. As mentioned above, this corresponds to $\bar{q}_{i+1} = \bar{q}_i - q_i$ for $i = 1, \dots, N-1$. The expectation is taken with respect to the uncertainty about the cost parameter in the subsequent period, θ_{i+1} , given the current value of the cost parameter, θ_i . The cost minimizing emissions level $q_i^*(\bar{q}_i, \theta_i)$ solves

$$\min_{q_i} V_i(\bar{q}_i, q_i, \theta_i). \quad (8)$$

The optimized value function is:

$$V_i^*(\bar{q}_i, \theta_i) \equiv V_i(\bar{q}_i, q_i^*(\bar{q}_i, \theta_i), \theta_i). \quad (9)$$

The marginal cost of emissions is:

$$p_i^*(\bar{q}_i, \theta_i) \equiv -\frac{\partial c(q_i^*(\bar{q}_i, \theta_i), \theta_i)}{\partial q_i} = c(q_i^*(\bar{q}_i, \theta_i), \theta_i). \quad (10)$$

The sequence of emissions functions, $q_i^*(\bar{q}_i, \theta_i), i = 1, \dots, N$, form the cost minimizing dynamic emissions policy. The sequence of price functions, $p_i^*(\bar{q}_i, \theta_i), i = 1, \dots, N$, form the price corresponding to the cost minimizing dynamic emissions policy.

We now turn to examining the solution to this programming problem under different circumstances. We begin by solving the certainty case, since this provides useful intuition for the uncertainty cases. We then solve two uncertainty cases, one when shocks are fully temporary and θ_i follows Equation (3), and one when shocks are fully permanent and θ_i follows Equation (4).

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

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

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

$$q_i^* = \frac{1}{N-i+1}\bar{q}_i - \frac{1}{2}(N-i)(\nu-r), \quad (11)$$

and

$$V_i^*(\bar{q}_i, \theta_i) = ie^{\theta_i - q_i^*} = ie^{\theta_i - \frac{1}{N-i+1}\bar{q}_i + \frac{1}{2}(N-i)(\nu-r)}. \quad (12)$$

The log marginal cost of emissions is:

$$\ln(p_i^*(\bar{q}_i, \theta_i)) = \theta_i - q_i^*(\bar{q}_i, \theta_i) = \theta_i - \frac{1}{N-i+1}\bar{q}_i + \frac{1}{2}(N-i)(\nu-r). \quad (13)$$

The expressions in Equations (11)–(13) 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 of the remaining number of periods. Therefore, it is not immediately clear from Equation (11) 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, $i = 1, 2, \dots, N$, 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 :

$$q_i^*(\bar{q}, N, \nu, r) = \left[\frac{1}{N}\bar{q} - \frac{N}{2}(\nu - r) \right] + i(\nu - r). \quad (14)$$

The marginal cost can also be expressed in terms of the periods, $i = 1, 2, \dots, N$. Transforming the price –the negative of the marginal cost– in log terms for the ease of exposition, p_i^* can be expressed as:

$$\begin{aligned} \ln \left(p_i^*(\bar{q}, N, \nu, r) \right) &= \ln \left(c \left(q_j^*(\bar{q}, N, \nu, r), \theta_i \right) \right) = \theta_i - q_i^* \\ &= \theta_0 + i\nu - \frac{1}{N}\bar{q} + \frac{N}{2}(\nu - r) - i\nu + ir \\ &= \theta_0 - \frac{1}{N}\bar{q} + \frac{N}{2}(\nu - r) + ir. \end{aligned} \quad (15)$$

To understand the optimal emissions policy in Equation (11) 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, $\frac{1}{N}\bar{q}$. The marginal cost of emissions rises, 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 (14) and shown in Figure 3(a), so that the marginal cost of abatement grows at the discount rate, r , as seen in Equation (15) and reported by Figure 3(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.

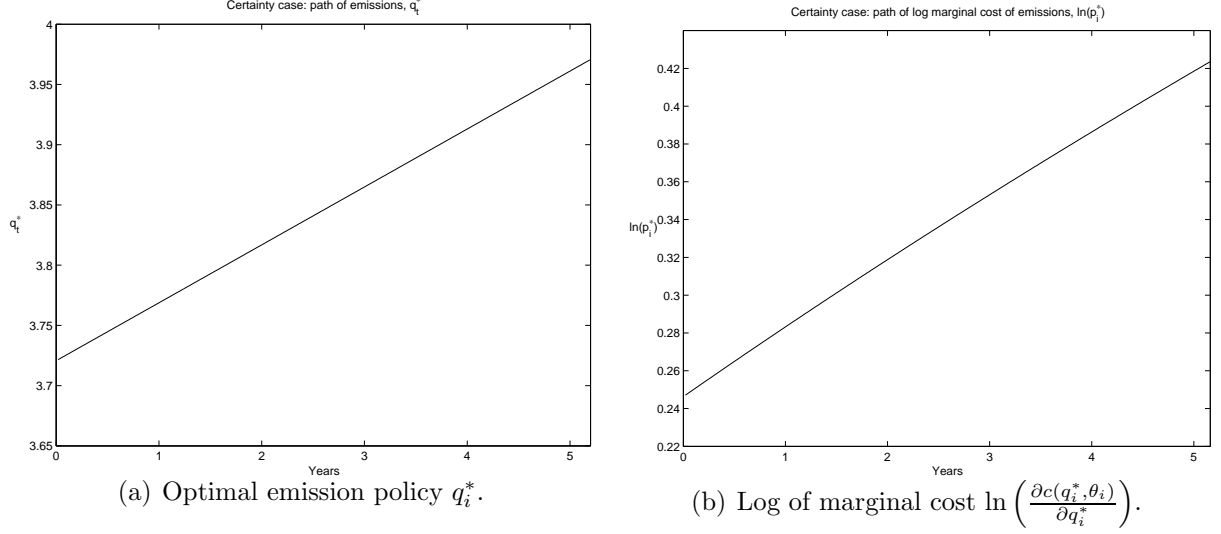


Figure 3: 5-year evolution of the cost minimizing dynamic emission policy, q_i^* , (left) and the log of marginal cost of emissions, $\ln\left(\frac{\partial c(q_i^*, \theta_i)}{\partial q_i^*}\right)$, (right) for the certainty case. In this example $\bar{q} = 1,000$, $\theta_0 = 5$, $\nu = 0.1$, $\sigma = 0$, $r = 0.05$, and the time step corresponds to a week.

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 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.

¹Of 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.

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

$$q_i^* = \frac{1}{N-i+1} \bar{q}_i - A_i \sigma^2 - \frac{1}{2} (N-i)(\nu - r) + \frac{N-i}{N-i+1} \sigma \epsilon_i \quad (16)$$

where

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

The general form of the optimized value function is:

$$V_i^*(\bar{q}_i) = i e^{\theta_i - q_i^*} = i e^{\theta_i - \frac{1}{N-i+1} \bar{q}_i + A_i \sigma^2 + \frac{1}{2} (N-i)(\nu - r) - \frac{N-i}{N-i+1} \sigma \epsilon_i}. \quad (17)$$

The log price is:

$$\begin{aligned} \ln(p_i^*) &= \theta_i - \frac{1}{N-i+1} \bar{q}_i + A_i \sigma^2 + \frac{1}{2} (N-i)(\nu - r) - \frac{N-i}{N-i+1} \sigma \epsilon_i \\ &= \theta_0 + i\nu + \sigma \epsilon_i - \frac{1}{N-i+1} \bar{q}_i + A_i \sigma^2 + \frac{1}{2} (N-i)(\nu - r) - \frac{N-i}{N-i+1} \sigma \epsilon_i \\ &= \theta_0 + i\nu - \frac{1}{N-i+1} \bar{q}_i + A_i \sigma^2 + \frac{1}{2} (N-i)(\nu - r) - \frac{1}{N-i+1} \sigma \epsilon_i \end{aligned} \quad (18)$$

The optimal emissions policy in Equation (16) 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 volatility 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 volatility 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 volatility in the cost parameter is absorbed in adjustment to the current level of emissions. This is

because of the 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 begins to reflect a portion of the volatility of the cost parameter.

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

$$Var_{i-1}(q_i^*) = \frac{N-i}{N-i+1}\sigma, \quad \text{and} \quad Var_{i-1}(\ln(p_i^*)) = \frac{1}{N-i+1}\sigma. \quad (19)$$

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

$$Var_0(q_i^*) = \sqrt{\sum_{h=1}^{i-1} \left(\frac{1}{N-h+1}\right)^2 + \left(\frac{N-i}{N-i+1}\right)^2} \sigma. \quad (20)$$

and

$$Var_0(\ln(p_i^*)) = \sqrt{\sum_{h=1}^i \left(\frac{1}{N-h+1}\right)^2} \sigma. \quad (21)$$

Figure 4(a) shows a pair of one-standard deviation confidence bounds around the expected path of the optimal quantity of emissions through time. Figure 4(b) shows a pair of one-standard deviation confidence bounds around the expected path of the log of marginal cost through time.

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

$$q_i^* = \frac{1}{N-i+1}\bar{q}_i - \frac{1}{2}(N-i)(\mu-r), \quad (22)$$

and

$$V_i^*(\bar{q}_i, \theta_i) = ie^{\theta_i - q_i^*} = ie^{\theta_i - \frac{1}{N-i+1}\bar{q}_i + \frac{1}{2}(N-i)(\mu-r)}. \quad (23)$$

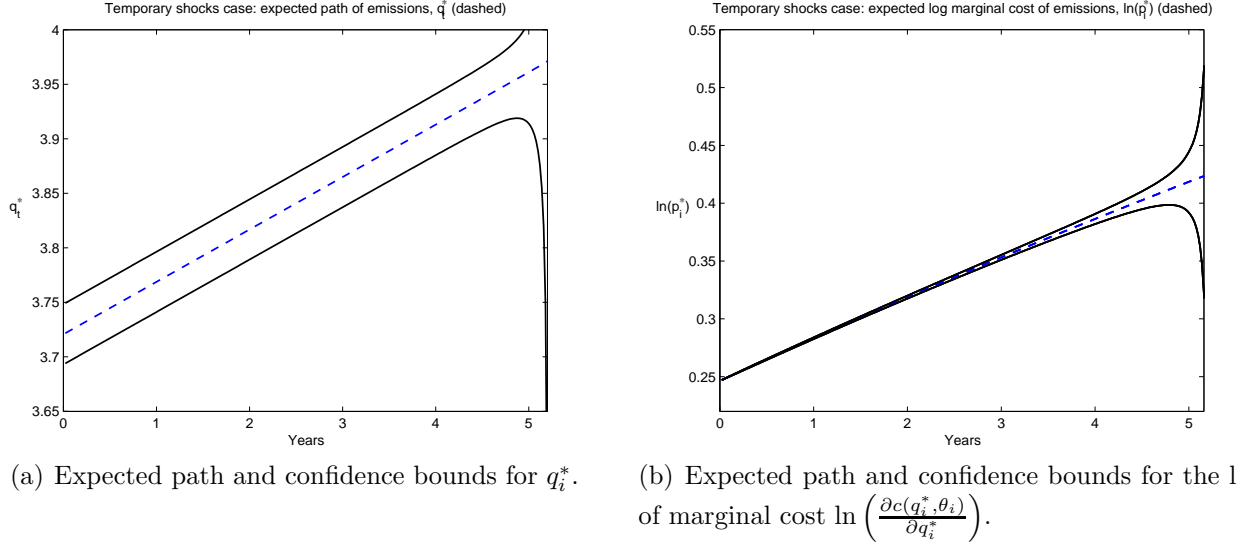


Figure 4: One-standard deviation confidence bounds around the expected path of the optimal quantity of emissions, q_i^* , (left) and around the expected path of the log of marginal cost, $\ln\left(\frac{\partial c(q_i^*, \theta_i)}{\partial q_i^*}\right)$, (right). In this example $\bar{q} = 1,000$, $\theta_0 = 5$, $\nu = 0.1$, $\sigma = 0.2$, $r = 0.05$, 5-year horizon, and the time step corresponds to a week.

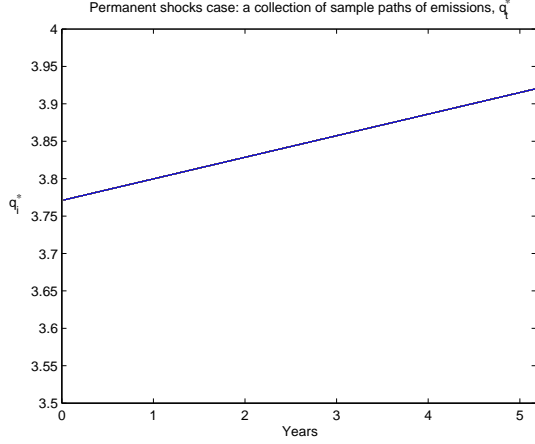
The log price is:

$$\begin{aligned} \ln(p_i^*) &= \theta_i - \frac{1}{N-i+1}\bar{q}_i + \frac{1}{2}(N-i)(\mu-r) \\ &= \theta_{i-1} + \mu + \sigma\epsilon_i - \frac{1}{N-i+1}\bar{q}_i + \frac{1}{2}(N-i)(\mu-r) \end{aligned} \quad (24)$$

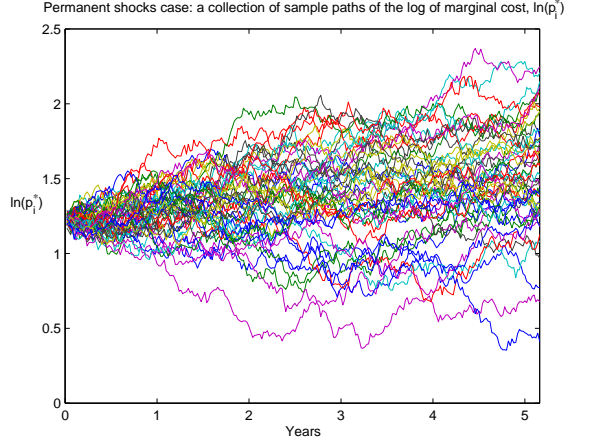
The optimal emissions policy in Equation (22) and represented in Figure 5(a) is identical to the certainty case. Emissions in each period are a proportional fraction of the remaining available allowances as determined by the remaining number of periods over which those allowances must be shared, and adjusted for an allowance for growth in emissions to match the growth rate in the cost parameter. 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 (24).

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

$$Var_{i-1}(q_i^*) = 0, \quad \text{and} \quad Var_{i-1}(\ln(p_i^*)) = \sigma. \quad (25)$$



(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 5: 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.

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

$$Var_0(q_i^*) = 0, \quad \text{and} \quad Var_0(\ln(p_i^*)) = \sqrt{i}\sigma. \quad (26)$$

Figure 5(a) shows the deterministic emissions policy in the face of permanent uncertainty in the cost parameter. Figure 5(b) shows the log of the marginal price.

These two cases provide powerful 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.

This result is a very powerful demonstration of a weakness in the claim that because greenhouse gases are a stock pollutant therefore the optimal policy must be a price control. This only corresponds with the temporary uncertainty case. In the case of permanent uncertainty, a price control will clearly not be optimal since it is price that ought to absorb

all of the shocks to cost, and a price control regime would not accommodate this.

Weitzman's original paper was about a policy maker that was uninformed about cost while the producer was informed. Although we do not formally model a policy maker that is uninformed about cost, we think that our result on the characteristics of the cost minimizing emissions policy is instructive for any model of a policy maker's optimal action given asymmetry between the policy maker and the producer. Given our result that quantity is more variable when uncertainty is temporary, and price is more variable when uncertainty is permanent, regardless of the specific model of the policy maker and the producer, it is less likely that a quantity control will get close to the cost minimizing dynamic emissions path when uncertainty is temporary, and less likely that a price control will get close the cost minimizing path when uncertainty is permanent.

We conclude this section with a brief proviso. Oftentimes economists speak casually of a cap- and-trade system as being a quantity control, in contrast to a carbon tax, which is a price control. In reality, a cap-and-trade system that allows banking and borrowing of allowances across periods can mimic the benefits of a price control regime because agents will reallocate allowances between periods so as to equate marginal cost across periods. In particular, the solution we have derived for the dynamic emissions policy is an optimal allocation of the allowances across the periods included under the cap. This optimal allocation would be implemented by a cap-and-trade system with banking and borrowing. 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. The cap-and-trade with banking and borrowing implements the dynamically efficient allocation of allowances, regardless of the sort of uncertainty it is faced with.

3 A Discrete Time Pair of Examples

In the previous section we solved for the cost minimizing dynamic 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 weighing 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 is extended 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 difficult. In order to simplify things, we construct two extremely stylized examples of temporary and permanent uncertainty within a simple model of costs and benefits. We assume N discrete periods, with no discounting. Emissions in each period are $q_i \geq 0$, with $i = 1, \dots, N$. 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:

$$B(Q_N) = -\frac{b}{2}Q_N^2, \quad (27)$$

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,

$$C(q_i, \theta_i) = \frac{c}{2}(q_i - \bar{q})^2 - \theta_i(q_i - \bar{q}), \quad (28)$$

where $c > 0$ is a fixed parameter, \bar{q} is a reference level of emissions and θ_i is a non-negative random variable whose realization is indicated by $\tilde{\theta}_i$.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 dynamic problem in which the regulator establishes either a quantity or a price constraint at the start of each period, then the uncertain parameter for that period

²The variable θ_i takes truly random variables only at time $i = 1$. For $i = 2, \dots, N$, θ_i is either a given value (θ) or equal to the value taken at time $i = 1$, (θ_1).

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 a new quantity or price constraint must be set.

We use the framework just described to construct two 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. $\tilde{\theta}_1$. 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 periods $i = 2, \dots, N$ 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 $i = 2, \dots, N$. In the second case, the cost parameters for periods $i = 2, \dots, N$ 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 $i = 2, \dots, N$.

We solve the model by backward programming. In both cases, whatever uncertainty existed has been resolved at the conclusion of the first period. Therefore, the optimal level of emissions in every future period can be calculated and readily enforced by the regulator. 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 $i = 2, \dots, N$. 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:

$$V(q_1^*, \tilde{\theta}_1) = \max_{q(q_1^*, \tilde{\theta}_1)} B\left(q_1^* + (N-1)q\right) - (N-1)C(q, \theta). \quad (29)$$

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

$$\max_{q_1(\cdot)} \mathbb{E}_{\tilde{\theta}_1} \left[V(q_1(\cdot), \tilde{\theta}_1) - C(q_1(\cdot), \tilde{\theta}_1) \right]. \quad (30)$$

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, the regulator must either (i) set q_1 , or (ii) fix a price, p_1 , 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. Before solving our problem we first present in Figure 6 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 $i = 2, \dots, N$. The marginal benefit function graphed is:

$$\frac{\partial B(Q_N)}{\partial q_1} = -b(q_1 + (N - 1)q), \quad (31)$$

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

$$\frac{\partial C(q_1, \theta_1)}{\partial q_1} = c(q_1 - \bar{q}) - \theta_1. \quad (32)$$

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 6. 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

³In 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 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 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. A stark illustration is provided in a figure from Pizer (2002) reproduced here as Figure 7. 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 regulate 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 global warming problem.

The problem is that the marginal benefit function written in Equation (31) and shown

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.

in the figure is not the correct marginal benefit function for the first period optimization shown in Equation (30). The correct marginal benefit function is:

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

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:

$$\begin{aligned} \frac{\partial V(q_1, \theta_1)}{\partial q_1} = & -b\left(q_1 + (N-1)q^*(q_1, \theta_1)\right) \left(1 + (N-1)\frac{\partial q^*}{\partial q_1}\right) \\ & -(N-1)c(q^*(q_1, \theta_1) - \bar{q})\frac{\partial q^*}{\partial q_1} + (N-1)\theta\frac{\partial q^*}{\partial q_1}. \end{aligned} \quad (33)$$

The first order condition on q^* implies:

$$-b\left(q_1 + (N-1)q^*(q_1, \theta_1)\right)(N-1)\frac{\partial q^*}{\partial q_1} - (N-1)c(q^*(q_1, \theta_1) - \bar{q})\frac{\partial q^*}{\partial q_1} + (N-1)\theta\frac{\partial q^*}{\partial q_1} = 0, \quad (34)$$

so that by substituting (34) into (33) we have:

$$\frac{\partial V(q_1, \theta_1)}{\partial q_1} = -b\left(q_1 + (N-1)q^*(q_1, \theta_1)\right). \quad (35)$$

The evaluation of Equation (35) 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 (35) reduces to Equation (31) so that Figure 1 is approximately correct.

However, for the second case Figure 6 is entirely inappropriate. In the second case the realization $\tilde{\theta}_1$ affects the cost functions in years $i = 2, \dots, N$, 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 6. 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! This is why Figure 7, taken from Pizer (2002) is misleading if uncertainty takes the form captured in the second case. In the second case we have:

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

which is inconsistent with Figure 6. Figure 8 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.

In the second case, the preference for quantity or price controls depends upon the relative steepness of the marginal benefit function against the marginal *aggregate cost* function. Equation (28) is a per period cost function. The aggregate cost function, $D(Q_N)$, is the result of allocating total emissions efficiently across years:

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

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

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

Whether price controls or quantity controls are preferred is once again an empirical question. For cases that lie in between the extremes of our two cases, it is a difficult

empirical question since the complicated structure of the uncertainty through time must be factored into the measurement of marginal benefit and marginal cost. The contribution of this pair of examples is that they show two extreme alternatives in a form that can be easily solved and that clarifies the essential issue.

4 Discussion

As we mentioned in the introduction, Hoel and Karp (2002), Newell and Pizer (2003), and Karp and Zhang (2005) developed models designed to address the dynamic structure of the uncertainty, and all papers appear to conclude that price controls are better. What is different in our model?

Hoel and Karp (2002) are very explicit in assuming zero correlation between cost shocks through time which is equivalent to assuming all shocks are temporary. Our results only differ when we abandon that assumption and present cases in which cost shocks have a permanent component. This helps to highlight the true source of the stock pollutant argument, i.e., the assumption that cost shocks are temporary.

Newell and Pizer (2003) allow correlated shocks, but solve the open loop problem in which there is no feedback adjustment to the optimal regulation based on information obtained in earlier periods. Their optimal policy is designed at the outset and is constant across all realizations of the uncertainty. In Section 3, where we explicitly weigh costs and benefits, we model the closed loop problem in which the policy at $i = 2$ can be set contingent on the realization of the uncertainty at $i = 1$. Our model still incorporates the restricted strategy space for the policy maker who at $i = 0$ does not know the uncertain parameter for $i = 1$. We assume the policy maker becomes informed after the conclusion of the first period and so can update the regulations for the following periods. In Section 2, where information flows regularly through time and we solve the dynamic cost minimizing emissions policy, we cannot speak of any delay in information to the policy maker. However, our results are stronger since we fully characterize the nature of the cost minimizing emissions policy and therefore the information needed to implement it. 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 with respect to minimizing the cost of achieving a fixed level of emissions. Of course, 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 but not by means of a quantity instrument.

Karp and Zhang (2005) is the closest to our problem in that they allow correlation between cost shocks through time and solve the closed loop problem in which the policy maker learns something about the realization of earlier cost shocks and can adjust the regulations in response. Indeed, they show that the higher is the correlation, the more likely is the optimal policy to be a quantity control.⁴ Our examples provide a simpler or different entry into this result.

Despite their theoretical conclusion that a quantity control can be better, depending upon the parameters, Karp and Zhang (2005) identify the parameters that they believe best match the problem and find that price controls are better. Nothing in this paper rules out this result. Indeed, the objective as stated at the outset is to reestablish the problem as an empirical one. Whether this empirical argument in favor of price control is true depends upon the faith one has in the parameter estimates for the cost and benefit functions and the model of uncertainty.

This would be a good place, however, to call attention to the fact that many comparisons between price and quantity controls, including Karp and Zhang (2005), do not allow banking and borrowing of the quantity limits across time. The actual quantity controls employed in emissions regulations generally allow banking of allowances across time and often allow some borrowing or, what is equivalent, have been designed with an up front loading of allocated allowances so that the optimal decision is to bank. In our solution to the optimal dynamic emissions policy, we have implicitly allowed unlimited banking and borrowing, since we simply ask how should the total emissions be allocated through time. In the case of temporary uncertainty, we show that the quantity of emissions in each period absorbs most of the volatility in the cost parameter, while the price absorbs only a little. While some may take this to be an argument in favor of a price control, this requires

⁴They consider several variations on their model, including various assumptions about trading of quotas and the information learned by the policy maker. Our examples are closest in spirit to theirs with the assumption of tradeable quotas where the policy maker is learning about the cost shocks through time.

jumping to the conclusion by ignoring the possibility of a quantity control implemented not period-by-period, but across the full horizon. If a quantity limit is set across all N periods, a trading of allowances with unlimited banking and borrowing implements the cost minimizing dynamic emissions policy. Therefore, banking and borrowing provide a substitute mechanism for the beneficial flexibility of price controls in the face of temporary shocks to cost uncertainty. While the exercise of examining a strict, period-by-period quantity control as compared against a strict, period-by-period price control are instructive about the important forces at work, making a policy recommendation based on fitting estimated parameters to these two limited alternatives seems distracting. In practice, quantity controls involve some element of banking and borrowing, and the empirical question is whether the amount of flexibility created by the banking and borrowing provisions is sufficient for the amount of temporary cost uncertainty. The exercise in Karp and Zhang (2005) is not informative on this question.

It is worth pausing to ask whether the types of uncertainty at hand in the greenhouse gas problem are best modeled as temporary or permanent or something in between. Hoel and Karp (2002) state categorically, but without explanation, that “In our view, serially correlated shocks are not central to the issue of stock regulation with asymmetric information.” Certainly, some elements of cost uncertainty are likely to be very temporary. One possible example is brief weather episodes such as a hotter than expected summer forcing greater dispatch of coal fired power plants, or a dry year that limits the use of hydro-power. Another example would be a labor dispute such as the strikes which lowered some country emissions during the first, trial phase of the European Union Emissions Trading System. However, we think that there are many other elements of cost uncertainty that arguably have a large permanent element to them. An example of such an uncertainty is the rate of economic growth. At least some authors model this as a random walk like our permanent uncertainty case. Any cost uncertainty based on the development of new technologies is likely to fall into this category. For example, while many policy makers presume that carbon sequestration will become viable, significant uncertainty still exists surrounding this new technology that is untested at scale. Should the technology become proven and the low cost established, it is largely proven for good or at least will not have to be proven again in the same fashion, year after year.

Indeed, what little work has been done on the uncertainties related to the greenhouse gas problem seems to suggest underlying models close to the permanent shocks model rather than the temporary shocks. For example, Webster et al. (2002) and Pizer (2002)

show confidence bounds for baseline emissions which appear to grow through time in a fashion inconsistent with a model of exclusively temporary shocks. Clearly this is an issue worthy of more careful discussion and modeling.

A Appendix

We solve the optimal pollution control problem using backward programming. For the sake of exposure, we count periods backwards from the endpoint using the index j to denote periods from the endpoint, with $j = N, \dots, 2, 1$. Within this backward programming exercise we understand that the subscripting on all variables denotes the period from the endpoint. 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.

Aa 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 the backward programming, 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

$$\begin{aligned} V_2(\bar{q}_2, q_2, \theta_2) &= \mathbb{E}_{\theta_2} \left[c(q_2, \theta_2) + e^{-r} V_1^*(\bar{q}_1(\bar{q}_2, q_2), \theta_1) \right] \\ &= \left[e^{\theta_2 - q_2} + e^{-r} \left[e^{\theta_1 - (\bar{q}_2 - q_2)} \right] \right] \\ &= \left[e^{\theta_2 - q_2} + e^{-r} \left[e^{\theta_2 + \nu - (\bar{q}_2 - q_2)} \right] \right] \\ &= e^{\theta_2} \left[e^{-q_2} + e^{-(r-\nu)} e^{-(\bar{q}_2 - q_2)} \right]. \end{aligned} \tag{36}$$

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

$$q_2^* = \frac{1}{2}\bar{q}_2 - \frac{1}{2}(\nu - r). \quad (37)$$

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

$$V_2^*(\bar{q}_2, \theta_2) = 2e^{\theta_2 - q_2^*} = 2e^{\theta_2 - \frac{1}{2}\bar{q}_2 + \frac{1}{2}(\nu - r)}. \quad (38)$$

For $j = 3$, we have

$$\begin{aligned} V_3(\bar{q}_3, q_3, \theta_3) &= \mathbb{E}_{\theta_3} \left[c(q_3, \theta_3) + e^{-r} V_2^*(\bar{q}_2(\bar{q}_3, q_3), \theta_2) \right] \\ &= \left[e^{\theta_3 - q_3} + e^{-r} \left[2e^{\theta_2 - \frac{\bar{q}_3 - q_3 - \nu + r}{2}} \right] \right] \\ &= \left[e^{\theta_3 - q_3} + 2e^{-r} \left[e^{\theta_3 + \nu - \frac{\bar{q}_3 - q_3 - \nu + r}{2}} \right] \right] \\ &= e^{\theta_3} \left[e^{-q_3} + 2e^{-\frac{\bar{q}_3 - q_3 - 3\nu + 3r}{2}} \right]. \end{aligned} \quad (39)$$

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

$$\begin{aligned} q_3^* &= \frac{2}{3} \frac{\bar{q}_3 - 3\nu + 3r}{2} \\ &= \frac{1}{3}\bar{q}_3 - \frac{2}{2}(\nu - r) \end{aligned} \quad (40)$$

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

$$V_3^*(\bar{q}_3, \theta_3) = 3e^{\theta_3 - q_3^*} = 3e^{\theta_3 - \frac{1}{3}\bar{q}_3 + (\nu - r)}. \quad (41)$$

For $j = 4$, we have

$$\begin{aligned}
V_4(\bar{q}_4, q_4, \theta_4) &= \mathbb{E}_{\theta_4} \left[c(q_4, \theta_4) + e^{-r} V_3^*(\bar{q}_3(\bar{q}_4, q_4), \theta_3) \right] \\
&= \left[e^{\theta_4 - q_4} + e^{-r} \left[3e^{\theta_3 - \frac{\bar{q}_4 - q_4 - 3\nu + 3r}{3}} \right] \right] \\
&= \left[e^{\theta_4 - q_4} + 3e^{-r} \left[e^{\theta_4 + \nu - \frac{\bar{q}_4 - q_4 - 3\nu + 3r}{3}} \right] \right] \\
&= e^{\theta_4} \left[e^{-q_4} + 3e^{-\frac{\bar{q}_4 - q_4 - 6\nu + 6r}{3}} \right].
\end{aligned} \tag{42}$$

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

$$\begin{aligned}
q_4^* &= \frac{3\bar{q}_4 - 6\nu + 6r}{4} \\
&= \frac{1}{4}\bar{q}_4 - \frac{3}{2}(\nu - r).
\end{aligned} \tag{43}$$

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

$$V_4^*(\bar{q}_4, \theta_4) = 4e^{\theta_4 - q_4^*} = 4e^{\theta_4 - \frac{1}{4}\bar{q}_4 + \frac{3}{2}(\nu - r)} \tag{44}$$

The general form of the optimal dynamic policy is:

$$q_j^* = \frac{1}{j}\bar{q}_j - \frac{1}{2}(j-1)(\nu - r). \tag{45}$$

and in calendar period i , it is:

$$q_i^* = \frac{1}{N-i+1}\bar{q}_i - \frac{1}{2}(N-i)(\nu - r). \tag{46}$$

The general form of the optimized value function is:

$$V_j^*(\bar{q}_j, \theta_j) = je^{\theta_j - q_j^*} = je^{\theta_j - \frac{1}{j}\bar{q}_j + \frac{1}{2}(j-1)(\nu - r)} \tag{47}$$

and in calendar period i , it is:

$$V_i^*(\bar{q}_i, \theta_i) = ie^{\theta_i - q_i^*} = ie^{\theta_i - \frac{1}{N-i+1}\bar{q}_i + \frac{1}{2}(N-i)(\nu-r)} \quad (48)$$

Ab 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 the backward programming, 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

$$\begin{aligned} V_2(\bar{q}_2, q_2, \theta_2) &= \mathbb{E}_{\theta_2} \left[c(q_2, \theta_2) + e^{-r} V_1^*(\bar{q}_1(\bar{q}_2, q_2), \theta_1) \right] \\ &= \left[e^{\theta_2 - q_2} + e^{-r} \mathbb{E}_{\theta_2} \left[e^{\theta_1 - (\bar{q}_2 - q_2)} \right] \right] \\ &= \left[e^{\Theta_2 + \sigma\epsilon_2 - q_2} + e^{-r} \mathbb{E}_{\theta_2} \left[e^{\Theta_1 + \sigma\epsilon_1 - (\bar{q}_2 - q_2)} \right] \right] \\ &= \left[e^{\Theta_2 + \sigma\epsilon_2 - q_2} + e^{-r} e^{-(\bar{q}_2 - q_2)} e^{\Theta_2 + \nu} \mathbb{E} \left[e^{\sigma\epsilon_1} \right] \right] \\ &= e^{\Theta_2} \left[e^{\sigma\epsilon_2 - q_2} + e^{-(r - \nu - \frac{\sigma^2}{2})} e^{-(\bar{q}_2 - q_2)} \right] \end{aligned} \quad (49)$$

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

$$q_2^* = \frac{1}{2}\bar{q}_2 - \frac{1}{2}(\nu - r) - \frac{1}{4}\sigma^2 + \frac{1}{2}\sigma\epsilon_2 \quad (50)$$

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

$$V_2^*(\bar{q}_2, \theta_2) = 2e^{\theta_2 - q_2^*} = 2e^{\theta_2 - \frac{1}{2}\bar{q}_2 + \frac{1}{2}(\nu - r) + \frac{\sigma^2}{4} - \frac{1}{2}\sigma\epsilon_2} \quad (51)$$

For $j = 3$, we have

$$\begin{aligned}
V_3(\bar{q}_3, q_3, \theta_3) &= \mathbb{E}_{\theta_3} \left[c(q_3, \theta_3) + e^{-r} V_2^*(\bar{q}_2(\bar{q}_3, q_3), \theta_2) \right] \\
&= \left[e^{\theta_3 - q_3} + e^{-r} \mathbb{E}_{\theta_3} \left[2e^{\theta_2 - \frac{\bar{q}_3 - q_3 - \nu + r - \sigma^2/2 + \sigma\epsilon_2}{2}} \right] \right] \\
&= \left[e^{\Theta_3 + \sigma\epsilon_3 - q_3} + e^{-r} \mathbb{E}_{\theta_3} \left[2e^{\Theta_2 + \sigma\epsilon_2 - \frac{\bar{q}_3 - q_3 - \nu + r - \sigma^2/2 + \sigma\epsilon_2}{2}} \right] \right] \\
&= \left[e^{\Theta_3 + \sigma\epsilon_3 - q_3} + 2e^{-r} e^{-\frac{\bar{q}_3 - q_3 - \nu + r - \sigma^2/2}{2}} e^{\Theta_3 + \nu} \mathbb{E} \left[e^{\sigma\epsilon_2} e^{-\frac{\sigma\epsilon_2}{2}} \right] \right] \\
&= \left[e^{\Theta_3 + \sigma\epsilon_3 - q_3} + 2e^{-r} e^{-\frac{\bar{q}_3 - q_3 - \nu + r - \sigma^2/2}{2}} e^{\Theta_3 + \nu} e^{\frac{\sigma^2}{8}} \right] \\
&= e^{\Theta_3} \left[e^{\sigma\epsilon_3 - q_3} + 2e^{-\frac{\bar{q}_3 - q_3 - 3\nu + 3r - 3\sigma^2/4}{2}} \right] \tag{52}
\end{aligned}$$

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

$$\begin{aligned}
q_3^* &= \frac{2\bar{q}_3 - 3\nu + 3r - 3\sigma^2/4 + 2\sigma\epsilon_3}{3} \\
&= \frac{1}{3}\bar{q}_3 - \frac{2}{3}(\nu - r) - \frac{1}{4}\sigma^2 + \frac{2}{3}\sigma\epsilon_3 \tag{53}
\end{aligned}$$

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

$$V_3^*(\bar{q}_3, \theta_3) = 3e^{\theta_3 - q_3^*} = 3e^{\theta_3 - \frac{1}{3}\bar{q}_3 + (\nu - r) + \frac{1}{4}\sigma^2 - \frac{2}{3}\sigma\epsilon_3} \tag{54}$$

For $j = 4$, we have

$$\begin{aligned}
V_4(\bar{q}_4, q_4, \theta_4) &= \mathbb{E}_{\theta_4} \left[c(q_4, \theta_4) + e^{-r} V_3^*(\bar{q}_3(\bar{q}_4, q_4), \theta_3) \right] \\
&= \left[e^{\theta_4 - q_4} + e^{-r} \mathbb{E}_{\theta_4} \left[3e^{\theta_3 - \frac{\bar{q}_4 - q_4 - 3\nu + 3r - 3\sigma^2/4 - 2\sigma\epsilon_3}{3}} \right] \right] \\
&= \left[e^{\Theta_4 + \sigma\epsilon_4 - q_4} + e^{-r} \mathbb{E}_{\theta_4} \left[3e^{\Theta_3 + \sigma\epsilon_3 - \frac{\bar{q}_4 - q_4 - 3\nu + 3r - 3\sigma^2/4 - 2\sigma\epsilon_3}{3}} \right] \right] \\
&= \left[e^{\Theta_4 + \sigma\epsilon_4 - q_4} + 3e^{-r} e^{-\frac{\bar{q}_4 - q_4 - 3\nu + 3r - 3\sigma^2/4}{3}} e^{\Theta_4 + \nu} \mathbb{E} \left[e^{\sigma\epsilon_3} e^{-\frac{2}{3}\sigma\epsilon_3} \right] \right] \\
&= \left[e^{\Theta_4 + \sigma\epsilon_4 - q_4} + 3e^{-r} e^{-\frac{\bar{q}_4 - q_4 - 3\nu + 3r - 3\sigma^2/4}{3}} e^{\Theta_4 + \nu} e^{\frac{\sigma^2}{18}} \right] \\
&= e^{\Theta_4} \left[e^{\sigma\epsilon_4 - q_4} + 3e^{-\frac{\bar{q}_4 - q_4 - 6\nu + 6r - 11\sigma^2/12}{3}} \right] \tag{55}
\end{aligned}$$

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

$$\begin{aligned} q_4^* &= \frac{3\bar{q}_4 - 6\nu + 6r - 11\sigma^2/12 + 3\sigma\epsilon_4}{3} \\ &= \frac{1}{4}\bar{q}_4 - \frac{3}{2}(\nu - r) - \frac{11}{48}\sigma^2 + \frac{3}{4}\sigma\epsilon_4 \end{aligned} \quad (56)$$

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

$$V_4^*(\bar{q}_4, \theta_4) = 4e^{\theta_4 - q_4^*} = 4e^{\theta_4 - \frac{1}{4}\bar{q}_4 + \frac{3}{2}(\nu - r) + \frac{11}{48}\sigma^2 - \frac{3}{4}\sigma\epsilon_4} \quad (57)$$

The general form of the optimal dynamic policy is:

$$q_j^* = \frac{1}{j}\bar{q}_j - A_j\sigma^2 - \frac{1}{2}(j-1)(\nu - r) + \frac{j-1}{j}\sigma\epsilon_j \quad (58)$$

where

$$A_j = \frac{j-1}{j} \left(A_{j-1} + \frac{1}{2(j-1)^2} \right) \text{ for } j = 2, \dots, N, \text{ and } A_1 = 0. \quad (59)$$

Rewriting in calendar period i , we have:

$$q_i^* = \frac{1}{N-i+1}\bar{q}_i - A_i\sigma^2 - \frac{1}{2}(N-i)(\nu - r) + \frac{N-i}{N-i+1}\sigma\epsilon_j \quad (60)$$

where

$$A_i = \frac{N-i}{N-i+1} \left(A_{i+1} + \frac{1}{2(N-i)^2} \right) \text{ for } i = 1, \dots, N-1, \text{ and } A_N = 0. \quad (61)$$

The general form of the optimized value function is:

$$V_j^*(\bar{q}_j) = j e^{\theta_j - q_j^*} = j e^{\theta_j - \frac{1}{j}\bar{q}_j + A_j\sigma^2 + \frac{1}{2}(j-1)(\nu - r) - \frac{j-1}{j}\sigma\epsilon_j} \quad (62)$$

and in calendar time:

$$V_i^*(\bar{q}_i) = ie^{\theta_i - q_i^*} = ie^{\theta_i - \frac{1}{N-i+1}\bar{q}_i + A_i\sigma^2 + \frac{1}{2}(N-i)(\nu-r) - \frac{N-i}{N-i+1}\sigma\epsilon_i}. \quad (63)$$

Ac 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 the backward programming, 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

$$\begin{aligned} V_2(\bar{q}_2, q_2, \theta_2) &= \mathbb{E}_{\theta_2} \left[c(q_2, \theta_2) + e^{-r} V_1^*(\bar{q}_1(\bar{q}_2, q_2), \theta_1) \right] \\ &= \left[e^{\theta_2 - q_2} + e^{-r} \mathbb{E}_{\theta_2} \left[e^{\theta_1 - (\bar{q}_2 - q_2)} \right] \right] \\ &= \left[e^{\theta_2 - q_2} + e^{-r} \mathbb{E}_{\theta_2} \left[e^{\theta_2 + (\mu - \frac{1}{2}\sigma^2) + \sigma\epsilon_1 - (\bar{q}_2 - q_2)} \right] \right] \\ &= \left[e^{\theta_2 - q_2} + e^{-r} e^{-(\bar{q}_2 - q_2)} e^{\theta_2 + \mu - \frac{1}{2}\sigma^2} \mathbb{E}_{\theta_2} \left[e^{\sigma\epsilon_1} \right] \right] \\ &= e^{\theta_2} \left[e^{-q_2} + e^{-(r - \mu + \frac{1}{2}\sigma^2)} e^{-(\bar{q}_2 - q_2)} e^{\frac{1}{2}\sigma^2} \right]. \end{aligned} \quad (64)$$

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

$$q_2^* = \frac{1}{2}\bar{q}_2 - \frac{1}{2}(\mu - r). \quad (65)$$

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

$$V_2^*(\bar{q}_2, \theta_2) = 2e^{\theta_2 - q_2^*} = 2e^{\theta_2 - \frac{1}{2}\bar{q}_2 + \frac{1}{2}(\mu - r)}. \quad (66)$$

For $j = 3$, we have

$$\begin{aligned}
V_3(\bar{q}_3, q_3, \theta_3) &= \mathbb{E}_{\theta_3} \left[c(q_3, \theta_3) + e^{-r} V_2^*(\bar{q}_2(\bar{q}_3, q_3), \theta_2) \right] \\
&= \left[e^{\theta_3 - q_3} + e^{-r} \mathbb{E}_{\theta_3} \left[2e^{\theta_2 - \frac{\bar{q}_3 - q_3 - \mu + r}{2}} \right] \right] \\
&= \left[e^{\theta_3 - q_3} + 2e^{-r} \mathbb{E}_{\theta_3} \left[e^{\theta_3 + (\mu - \frac{1}{2}\sigma^2) + \sigma\epsilon_2 - \frac{\bar{q}_3 - q_3 - \mu + r}{2}} \right] \right] \\
&= \left[e^{\theta_3 - q_3} + 2e^{-r} e^{-\frac{\bar{q}_3 - q_3 - \mu + r}{2}} e^{\theta_3 + \mu - \frac{1}{2}\sigma^2} \mathbb{E}_{\theta_3} \left[e^{\sigma\epsilon_2} \right] \right] \\
&= e^{\theta_3} \left[e^{-q_3} + 2e^{-\frac{\bar{q}_3 - q_3 - 3\mu + 3r}{2}} \right].
\end{aligned} \tag{67}$$

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

$$\begin{aligned}
q_3^* &= \frac{2}{3} \frac{\bar{q}_3 - 3\mu + 3r}{2} \\
&= \frac{1}{3} \bar{q}_3 - \frac{2}{2}(\mu - r)
\end{aligned} \tag{68}$$

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

$$V_3^*(\bar{q}_3, \theta_3) = 3e^{\theta_3 - q_3^*} = 3e^{\theta_3 - \frac{1}{3}\bar{q}_3 + (\mu - r)}. \tag{69}$$

For $j = 4$, we have

$$\begin{aligned}
V_4(\bar{q}_4, q_4, \theta_4) &= \mathbb{E}_{\theta_4} \left[c(q_4, \theta_4) + e^{-r} V_3^*(\bar{q}_3(\bar{q}_4, q_4), \theta_3) \right] \\
&= \left[e^{\theta_4 - q_4} + e^{-r} \mathbb{E}_{\theta_4} \left[3e^{\theta_3 - \frac{\bar{q}_4 - q_4 - 3\mu + 3r}{3}} \right] \right] \\
&= \left[e^{\theta_4 - q_4} + 3e^{-r} \mathbb{E}_{\theta_4} \left[e^{\theta_4 + (\mu - \frac{1}{2}\sigma^2) + \sigma\epsilon_3 - \frac{\bar{q}_4 - q_4 - 3\mu + 3r}{3}} \right] \right] \\
&= \left[e^{\theta_4 - q_4} + 3e^{-r} e^{-\frac{\bar{q}_4 - q_4 - 3\mu + 3r}{3}} e^{\theta_4 + \mu - \frac{1}{2}\sigma^2} \mathbb{E}_{\theta_4} \left[e^{\sigma\epsilon_3} \right] \right] \\
&= e^{\theta_4} \left[e^{-q_4} + 3e^{-\frac{\bar{q}_4 - q_4 - 6\mu + 6r}{3}} \right].
\end{aligned} \tag{70}$$

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

$$\begin{aligned} q_4^* &= \frac{3\bar{q}_4 - 6\mu + 6r}{4} \\ &= \frac{1}{4}\bar{q}_4 - \frac{3}{2}(\mu - r). \end{aligned} \quad (71)$$

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

$$V_4^*(\bar{q}_4, \theta_4) = 4e^{\theta_4 - q_4^*} = 4e^{\theta_4 - \frac{1}{4}\bar{q}_4 + \frac{3}{2}(\mu - r)} \quad (72)$$

The general form of the optimal dynamic policy is:

$$q_j^* = \frac{1}{j}\bar{q}_j - \frac{1}{2}(j-1)(\mu - r). \quad (73)$$

and in calendar time i , it is:

$$q_i^* = \frac{1}{N-i+1}\bar{q}_i - \frac{1}{2}(N-i)(\mu - r). \quad (74)$$

The general form of the optimized value function is:

$$V_j^*(\bar{q}_j, \theta_j) = je^{\theta_j - q_j^*} = je^{\theta_j - \frac{1}{j}\bar{q}_j + \frac{1}{2}(j-1)(\mu - r)} \quad (75)$$

that in calendar time i is:

$$V_i^*(\bar{q}_i, \theta_i) = ie^{\theta_i - q_i^*} = ie^{\theta_i - \frac{1}{N-i+1}\bar{q}_i + \frac{1}{2}(N-i)(\mu - r)}. \quad (76)$$

References

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.
- Newell, R. G. and Pizer, W. A. (2006). Indexed regulation. RFF Discussion Paper 06-32., Resources for the Future.
- Nordhaus, W. D. (1994). *Managing the Global Commons*. MIT Press, Cambridge.
- Pizer, W. A. (2002). Combining Price And Quantity Controls To Mitigate Global Climate Change. *Journal of Public Economics*, 85:409–434.
- 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.
- Webster, M., Babiker, M., Mayer, M., Reilly, J. M., Harnisch, J., Hyman, R., and Sarofim, M. C. (2002). Uncertainty In Emissions Projections For Climate Models. *Atmospheric Environment*, 36:3659–3670.
- Weitzman, M. L. (1974). Prices vs. Quantities. *Review of Economic Studies*, 41(4):683–691.

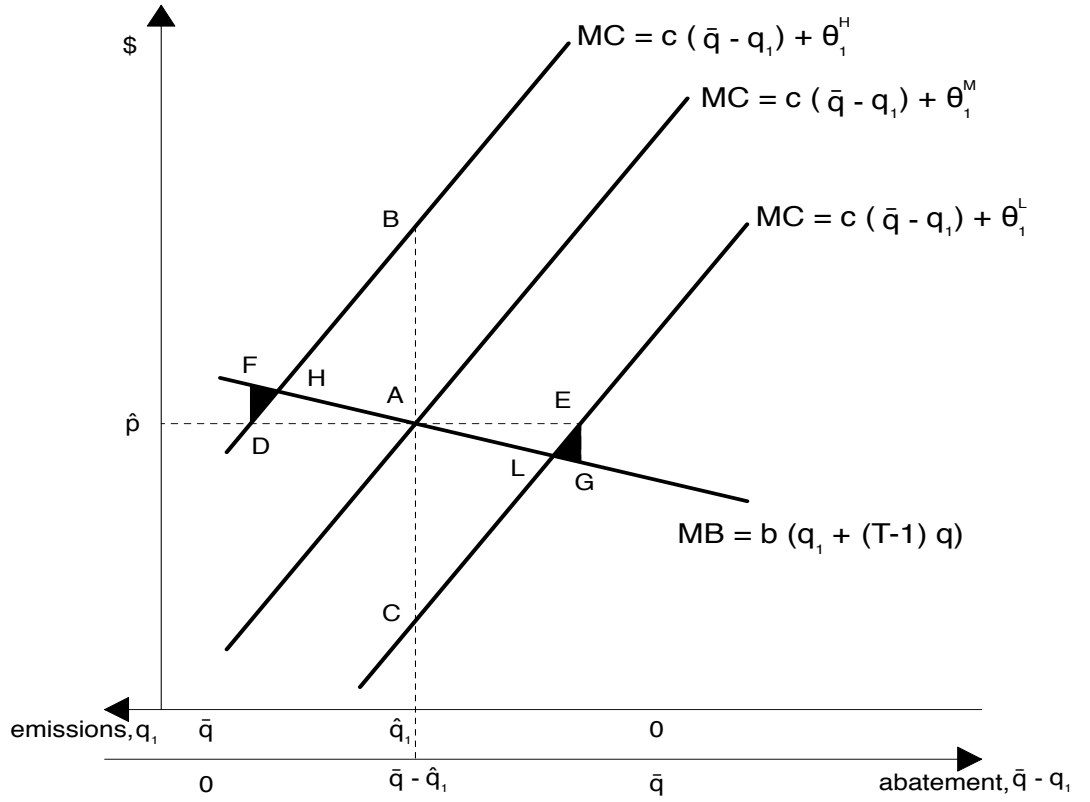


Figure 6: **Marginal costs and marginal benefits of abatement.** It is most common to show the marginal benefits and marginal costs as a function of abatement, which is the difference between emissions and a benchmark. This is equivalent to charting the negative of the marginal benefit and marginal cost functions from Equations (31) and (32), 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.

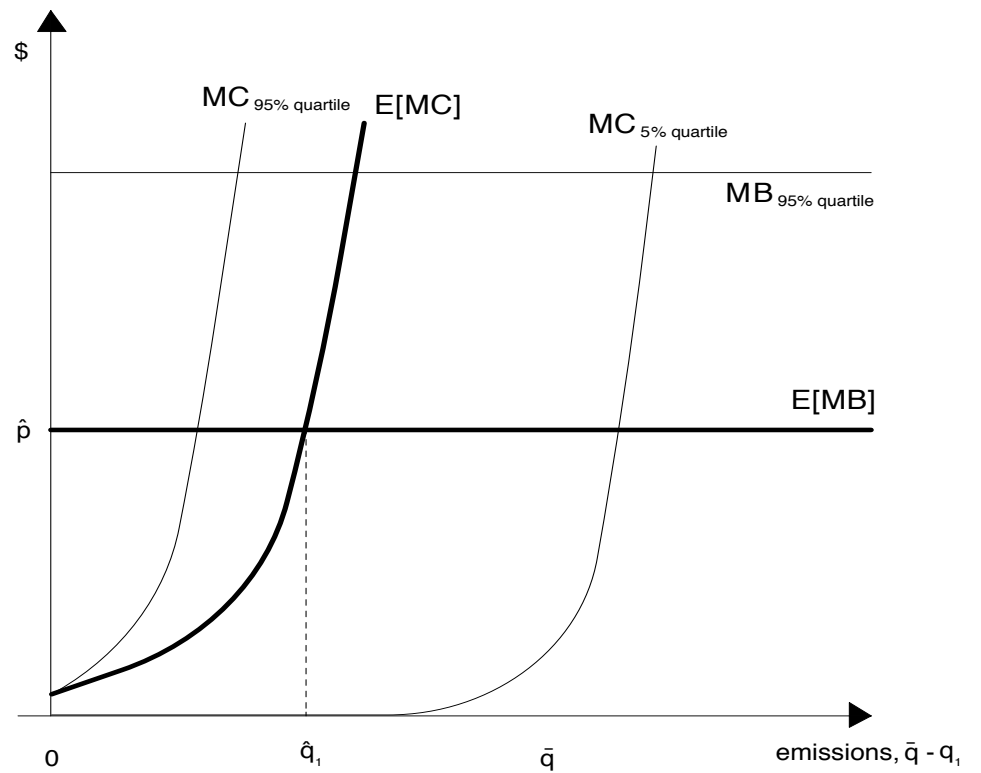


Figure 7: **Estimated comparison of annual marginal costs vs. aggregate marginal benefits.** Replication of the Figure in Pizer (2002) that represents the estimation of marginal costs and benefits in 2010. The 5% quantile of marginal benefits overlaps the x-axis.

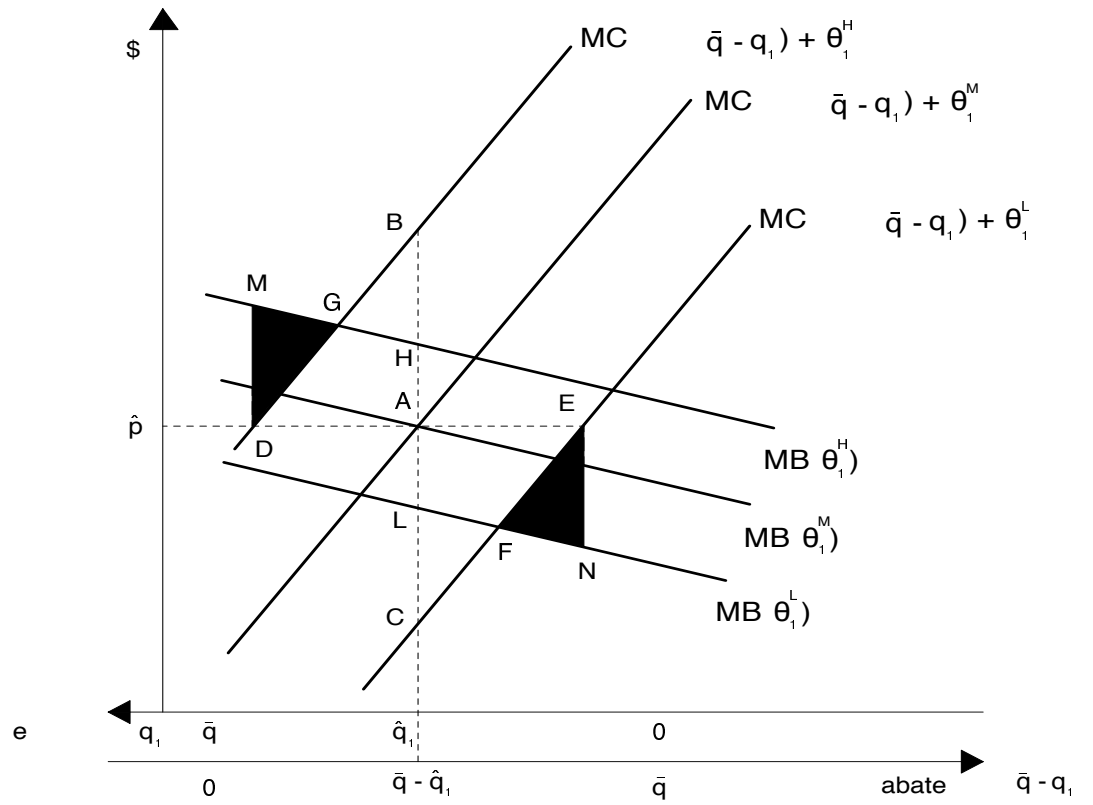


Figure 8: Marginal costs and marginal benefits for the second case. MC and MB represents marginal costs and marginal benefits, respectively.