Available online at www.sciencedirect.com

ACADEMIC PRESS

Journal of Environmental Economics and Management 45 (2003) 416–432

JOURNAL OF ENVIRONMENTAL ECONOMICS AND MANAGEMENT

http://www.elsevier.com/locate/jeem

# Regulating stock externalities under uncertainty[☆]

Richard G. Newell* and William A. Pizer

*Resources for the Future, 1616 P Street, NW, Washington, DC 20036, USA*

Received 26 February 2001; revised 11 December 2001

## Abstract

Using a simple analytical model incorporating benefits of a stock, costs of adjusting the stock, and uncertainty in costs, we uncover several important principles governing the choice of price-based policies (e.g., taxes) relative to quantity-based policies (e.g., tradable permits) for controlling stock externalities. As in Weitzman (Rev. Econom. Stud. 41(4) (1974) 477), the relative slopes of the marginal benefits and costs of controlling the externality continue to be critical determinants of the efficiency of prices relative to quantities, with flatter marginal benefits and steeper marginal costs favoring prices. But some important adjustments for dynamic effects are necessary, including correlation of cost shocks across time, discounting, stock decay, and the rate of benefits growth. Applied to the problem of greenhouse gases and climate change, we find that a price-based instrument generates several times the expected net benefits of a quantity instrument.

© 2003 Elsevier Science (USA). All rights reserved.

*Keywords:* Stock; Externality; Policy; Regulation; Uncertainty; Prices; Quantities; Taxes; Tradable permits

## 1. Introduction

The threat of global climate change is one of the most important and challenging problems facing the world community. Without a concerted international effort to reduce emissions of greenhouse gases, the world could face climatic changes with profound impacts on the global population and economy. At the same time, these reductions entail equally profound changes in the world's reliance on fossil fuels, the primary source of these emissions.

---

[☆] The paper has benefited from helpful comments from Larry Karp, Preston McAfee, Steve Polasky, Jim Sanchirico, Martin Weitzman, anonymous referees, seminar participants at Harvard, RFF, Stanford, and the University of California, support under NSF Grant #9711607, and research assistance from Brian Nadreau.

*Corresponding author.

*E-mail address:* newell@rff.org (R.G. Newell).

0095-0696/03/$ - see front matter © 2003 Elsevier Science (USA). All rights reserved.
doi:10.1016/S0095-0696(02)00016-5

There is another element to this challenge: Global climate change is a stock externality. The consequences depend not on emissions in a single year, but on the accumulated stock of emissions over time. In fact, many of today's most pressing policy concerns—from hazardous waste to research, development, and educational attainment—are characterized by stock-based externalities.

When addressing of these problems, policymakers face the inevitable task of not just setting a target for policy, but also of choosing a policy instrument for attaining that target. The Kyoto Protocol and Framework Convention on Climate Change, for example, consider restrictions on the annual flow of greenhouse gases, with the ultimate goal of limiting their atmospheric stock. In such policy settings, this paper offers guidance on the best instrument for regulating stocks when the costs of regulation are uncertain. The primary focus is negative externalities—pollutants such as carbon dioxide and other greenhouse gases, hazardous waste, pesticides in groundwater, and ozone depleting substances.[1] They are produced continually as a byproduct of economic activity, but—unlike other pollutants such as airborne particulate matter or volatile organic compounds—their harmful consequences are a function of a much larger stock accumulating in the environment, rather than an annual flow. The common element among these pollutants is that they *persist* for a long period of time.

Regulating such stocks involves considerable uncertainty. The benefits and costs of control are often known only approximately, while increases to the stock in any year can persist far into the future, introducing additional uncertainty about valuation in those future periods. Even as firms learn about costs and respond to a policy, the ability of regulators to ascertain and use this information is limited. Policy adjustments introduce incentives to behave strategically and realistic policy choices often involve only simple and infrequently adjustable controls.

Economic analyses of climate change policy have thus far developed the clearest intuition about instrument choice in a stock setting, but this has not generally been a focus of that research. Nordhaus [24] observes that damages in his climate policy simulations are essentially a linear function of emissions (i.e., marginal benefits are flat) and, based on an appeal to Weitzman [33], that this implies a preference for price instruments. He also suggests that this preference might extend to stock pollutants more generally, a point hinted at by Weitzman [33] when he notes that:

> ...in situations with more rather than fewer independent units producing outputs which substitute for each other in yielding benefits, there is a correspondingly greater relative advantage to the price mode of control. (p. 488)

While Weitzman describes physically distinct production, the stock issue describes inter-temporally distinct production—but the intuition is the same.[2]

---

[1] Examples of positive stock externalities include policies designed to protect wildlife habitats and preserve species, as well as provision of durable public goods, such as highways and national defense technologies.

[2] A literal use of Weitzman's approach in Section 5 of his 1974 article is limited to comparing prices in all periods versus quantities in all periods, as he does not consider the possibility that some "units" experience price controls while others quantity controls. Although unusual in the static cross-sectional setting that he goes on to elaborate, it is a natural possibility in our dynamic problem where instrument choice might switch over time. This is particularly true for stock externalities in which the legacy of an $n$-period price or quantity policy can persist far beyond the $n$-period policy duration, requiring explicit or implicit assumptions about what happens after $n$-periods (e.g., whether quantity or price policies continue).

Pizer [26], on the other hand, directly investigates the price versus quantity question for climate change using Monte Carlo simulations. Like Nordhaus [24] and Kolstad [16], he observes linear damages, but is further able to demonstrate that price policies indeed lead to much better welfare outcomes. Finally, McKibbin and Wilcoxen [21] argue that the absence of an obvious benefit to stabilizing either the flow or stock weighs heavily in favor of a price mechanism. None of these authors, however, clearly examine the link between the flatness of marginal benefits and issues of policy choice, or explain the conditions under which their observed results and conjectures will continue to hold. That is our goal: to develop a simple analytical model of policy choice for stock externality regulation and to apply the resulting framework to the important issue of greenhouse gas policy.

### 1.1. *Policy choice under uncertainty*

In a deterministic world, it is widely recognized that regulation based on either prices (e.g., taxes or subsidies) or quantities (e.g., tradable permits) can yield any desired level of output, including the economically efficient outcome. Although state-contingent policies could, in principle, be designed to maintain this proposition even under conditions of uncertainty,[3] such policies would be of little if any practical use. Recognizing this, Weitzman [33] initiated a discussion about the relative efficiency of alternative regulatory instruments in a distinctly different world characterized by uncertainty, asymmetry of information, second-best policy alternatives, and costly policy adjustments. Under these more realistic conditions, Weitzman found that there are fundamental differences in the consequences of price versus quantity regulation. He described the resulting divergence in efficiency as a function of basic and intuitive economic variables, including the slopes of marginal benefits and costs and the degree of cost uncertainty.

When uncertainty exists about costs, and policies must be fixed before the uncertainty is resolved, priced policies will lead to distinctly different outcomes than quantity policies. Pollution taxes, for example, encourage firms to reduce emissions until the marginal cost of reductions equals the tax. The tax leads to a range of possible emission levels depending on how uncertainty is resolved—but will fix marginal cost at the tax level. Conversely, a tradable permit system will fix the level of emissions, with the permit price determined by the marginal cost of meeting the emission constraint. The permit mechanism will therefore lead to a range of possible marginal costs depending on how uncertainty is resolved, but will lead to a fixed level of aggregate emissions. Different expected net benefits will therefore be associated with these alternate policies.

Weitzman's remarkable insight was that, on economic efficiency grounds, a flat expected marginal benefit function (relative to marginal costs) favors prices, while a steep benefit function favors quantities. Intuitively, flat marginal benefits imply a constant benefit per unit, suggesting that a tax could best correct the externality. In contrast, steep marginal benefits imply a dangerous threshold that should be avoided—a threshold that is efficiently enforced by a quantity control. Despite its substantial insight, however, the Weitzman result remains primarily a static story.

---

[3] Laffont [18] provides a careful description of the information structure assumed in policy choice problems formulated in the Weitzman tradition. Iterative policy processes with truth-revealing procedures could, in principle, be designed to yield first-best outcomes even in the face of uncertainty and information asymmetries [5,17]. Benford [3] extends this line of research to stock externalities, resulting in a "dynamic Kwerel scheme".

Although his discussion involving many production units could be applied to the case where those production units are located in different time periods, additional structure is necessary both to understand the importance of dynamic parameters and to consider the application to climate change.[4]

### 1.2. Stock regulation under uncertainty

Research on optimal policy choice under uncertainty has generally dealt with situations where both benefits and costs are a function only of current output, as in Weitzman's original formulation. Those that have explored the question in a dynamic context have generally done so in elaborate models where the general theory of stock externality regulation was not the object.[5] In contrast, we use an otherwise simplified model to explore how uncertainty influences the choice of policies to regulate a periodic flow when benefits are a function of the accumulated stock. As in [33], we find that the relative slopes of the marginal benefits and costs of controlling the externality continue to be critical determinants of the efficiency of prices relative to quantities, with flatter marginal benefits and steeper marginal costs favoring prices. But important adjustments for dynamic effects are necessary, including correlation of cost shocks across time, discounting, stock decay, and the rate of benefits growth. Applied to the problem of greenhouse gas policy, we find a price advantage of many times the welfare gain associated with quantity-based controls.

Our approach differs from related research by Hoel and Karp [7,8] in several ways. Most importantly, we model uncertainty in costs as being correlated across time, while Hoel and Karp assume independence of cost shocks. As described further below, we believe that not only is an assumption of correlated shocks more reasonable, but assuming otherwise ignores an important reason for possibly favoring quantity controls. In addition, given the long timeframes that are relevant for many stock externalities, we allow for benefits, costs, and baseline emissions to change over time. Finally, Hoel and Karp [8] assume uncertainty enters in a multiplicative, rather than additive fashion, and Hoel and Karp [7] focus in part on the frequency of adjustment of firms and policy to new information—an issue we do not address. Nonetheless, their main result is generally consistent with our own—that prices dominate quantities for climate policy.

In Section 2, we present our model and describe the optimal quantity and price policies. Section 3 derives an expression for the relative advantage of price versus quantity policies and Section 4 applies the result to climate policy. We conclude in Section 5.

---

[4] As we describe further below, our approach allows us to identify the incremental consequences of switching from price to quantity controls in year $t$ rather than $t-1$. This in turn allows us to add these consequences over varying durations of the price policy (e.g., the next year, the next ten years, or the next 40 years), an issue we consider important.

[5] Most economic research on stock pollutants has focused on optimal control in a deterministic setting [6,10,12,27,31]. This tends to remove the distinction among alternative instruments and sidesteps the issue of choosing an efficient instrument. Plourde and Yeung [28] extend these models to incorporate stochastic elements, but they only consider uncertainty in the amount of pollution generated and the stock decay rate, that is, on the benefit side. Without cost uncertainty, price and quantity controls remain "equivalent". Research on efficient fisheries management has addressed the issue in a context where benefits and costs are a function of the stock of fish, via the harvesting function [1,2,14,15]. However, parsimonious, intuitive results regarding efficient policy for controlling other types of stocks are obscured in these analyses by the particulars of optimal fishery modeling. A recent exception is Weitzman [35], which finds that in the presence of recurrent ecological uncertainty in fish stocks, landing fees unambiguously dominate harvest quotas. See Stavins [32] for a review of the literature on efficient policy instrument choice under uncertainty.

## 2. A model of stock externality regulation under uncertainty

In developing a model to explore issues of policy instrument choice for stock externality regulation, we were guided by three goals: (i) to keep the model parsimonious, (ii) to follow previous convention where possible, and (iii) to nonetheless include elements that are essential for the application to greenhouse gas policy. For these reasons, we chose to specify quadratic costs and benefits in the manner of Weitzman [33],[6] and to focus on basic price and quantity policies. The focus on these policies is appealing not only for its simplicity, but also due to the infeasibility and costliness of complex policies that entail continual readjustment or large amounts of information. It is possible of course to combine price and quantity mechanisms to form superior hybrid policies [29,34], but these are rarely, if ever, seen in practice. We also do not model so-called banking and/or borrowing which, in a dynamic context, can make quantity policies more flexible and efficient.[7] Without loss of generality, we use a discrete time framework, which is both more consistent with realistic policy design and avoids the use of stochastic calculus.

This approach allows us to compare and contrast our results for stock externalities with Weitzman's well-known results that are applicable to a single-period flow externality. To examine stock externalities, we make several modifications to his original framework, including omission of certain complicating features that turn out to be irrelevant for the final results.[8] First, in our model benefits are a function of the stock and costs are a function of the flow, whereas in Weitzman's model both costs and benefits were a function of the flow. Next, because changes in the stock level persist across time, it is necessary to set the model in a multi-period context. This dynamic context has several key features: stock depreciation, time discounting, cost correlation, and growth in baseline conditions, benefits, and costs.

With benefits and costs occurring at different points in time, intertemporal prices, or discount factors, are required. The persistence of stocks and the possibility of stock decay introduces a depreciation rate. In addition, just as the results in the static analysis can depend on the correlation of benefit and cost uncertainty within the single period [32,33], in a multi-period setting results will also hinge on correlation among costs in different periods. This correlation represents persistence in

---

[6] Weitzman [33] described quadratic costs and benefits as local approximations around the optimal quantity control which only needed to prevail over the range of likely disturbances in order for the results to be valid. In our model, since both the stock and flow change over time, we focus on globally quadratic costs and benefits.

[7] Banking is addressed in a deterministic fashion by Kling and Rubin [13] and in the presence of cost uncertainty by Yates and Cronshaw [39], Williams [37], and Yates [38]. Absent permanent cost shocks, full banking and borrowing across all periods would make a quantity control behave much like a price control—that is, quantities rather than marginal costs would fluctuate. This is analogous to how the marginal utility of consumption fluctuates only slightly in response to transient income shocks under the permanent income hypothesis. Yates and Cronshaw [39] find that allowing intertemporal permit trading indeed raises net benefits when marginal benefits are less steep than marginal costs.

[8] First, we omit uncertainty in the base level of costs and benefits (i.e., the constant terms in the cost and benefit functions) because Weitzman shows that it has no effect on the relative advantage of alternative policies. We also omit benefit uncertainty because it does not affect the behavior of firms or individuals in response to a policy and therefore cannot affect the relative advantage of alternative policies unless it is correlated with costs. Weitzman [33] and Stavins [32] have explored the potential importance of such benefit-cost correlation in detail. We suspect that including correlation in costs and benefits in our model would have implications similar to those found in Weitzman's original analysis, namely that positive correlation in benefits and costs will tend to favor quantity controls; the exact form of the consequences remains an issue for further research.

shocks to technology and baseline emissions, and would typically be positive—especially over annual intervals where business cycles and other macroeconomic shocks clearly persist. This feature is also essential in order for the discrete approach to match results in continuous time since the cost shocks will be increasingly and positively correlated over smaller and smaller intervals.[9] Finally, in order to properly apply the model to longer-term problems such as the regulation of greenhouse gas emissions, we need to allow for growth or decline in benefits, costs, and uncontrolled output due to changes over time in important factors such as income, population, and technology.[10]

### 2.1. Benefits, costs, and stock accumulation

Based on these choices, we specify benefits as

$$B_t(S_t) = -\frac{b_t}{2}(S_t - \bar{S}_t)^2, \tag{1}$$

where $t$ indexes time, $S_t$ is the stock of the regulated good, $\bar{S}_t$ is the benefit-maximizing stock, and $b_t$ is the per-period slope of marginal benefits—all of which can change over time. We assume benefits are concave ($b_t \geqslant 0$), so that benefits are maximized at $\bar{S}_t$ and there are diminishing returns, regardless of whether we are seeking to increase the stock (for a good) or decrease it (for a bad). While changes in $b_t$ over time allow for benefit growth, changes in $\bar{S}_t$ could be associated with adjustments in the desired level of a good due to economic growth. In the case of a negative externality like pollution, we assume the benefit maximizing level is zero, after normalizing if there is some positive natural level of the stock (as in the case of greenhouse gases such as carbon dioxide). To eliminate unnecessary complexity, we also omit the constant term from the quadratic form since it has no effect on the results; the linear term vanishes when the form is written in terms of deviations from the maximizing level. We apply similar simplifications below with costs.

Costs each period are given by

$$C_t(q_t, \theta_t) = \theta_t(q_t - \bar{q}_t) + \frac{c_t}{2}(q_t - \bar{q}_t)^2, \tag{2}$$

where $q_t$ is the quantity of the regulated good (or bad) produced, $\bar{q}_t$ is the cost-minimizing output of the good in the absence of regulation, $c_t$ is the slope of marginal costs, and $\theta_t$ is a shock to the marginal cost function.[11] Potential changes in $c_t$ and $\bar{q}_t$ allow for cost reductions and growth in

---

[9] This contrasts sharply with a result of Hoel and Karp [7], which assumes independent cost shocks and finds an unambiguous preference for prices as the time interval tends to zero.

[10] Uncontrolled output (of a negative externality) is likely to increase over time due to economic growth. Benefits also seem likely to increase due to growth in population and per capita income. Technological change may, however, mitigate growth in baseline output and benefits. In the climate change context, for example, energy efficiency improvements can counter the effect of economic growth on greenhouse gas emissions and adaptation can reduce the damages resulting from climate change. Technological change will also tend to reduce per unit production (abatement) costs, although increased scarcity in important inputs could counter these cost reductions.

[11] Despite its intuitive and analytical attraction, there is admittedly an asymmetry in our description of costs and benefits since one could imagine stock effects on the *cost* side, representing either knowledge or capital. While this remains an interesting area for further research, our intuition is that such an effect would favor price controls. Namely, if costs depend on both the regulated flow as well as a secondary (capital and/or knowledge) stock that is fixed in the short run, it will become increasingly expensive to make large, positive changes to the regulated flow in the short run. This would show up in our model as convexity in the marginal cost schedule and, as suggested by Yohe [40], this will tend to favor price controls.

uncontrolled output. The cost shock $\theta_t$ has an autoregressive form $\theta_t = \rho\theta_{t-1} + \varepsilon_t$, with correlation $\rho$ across time, and error $\varepsilon_t$ with zero mean and variance $\sigma_0^2$.[12] Reduction of a negative externality, for example, is therefore represented by $q_t < \bar{q}_t$. We assume costs are convex ($c_t > 0$) so that costs are minimized at $\bar{q}_t - \theta_t/c_t$ (ignoring the potential benefits). Any deviation from this rate, whether to increase output in the case of a good or decrease it in the case of a bad, leads to increasing costs at an increasing rate.

We represent the dynamic nature of the stock by an accumulation equation,

$$S_t = (1-\delta)S_{t-1} + q_t, \tag{3}$$

which specifies that the stock decays at rate $0 \leqslant \delta \leqslant 1$ in addition to the contribution of $q_t$. The decay rate can take on values representing cases ranging from a "pure stock externality" that persists forever ($\delta = 0$) to a "flow externality" ($\delta = 1$) that replicates the traditionally analyzed case.

### 2.2. Optimal quantity and price policies

Based on the benefits and costs defined in the previous section, we can write the discounted net benefits at time $t$ as

$$NB_t = \sum_{s=t}^{\infty} \frac{B_s(S_s) - C_s(q_s, \theta_s)}{(1+r)^{s-t}}, \tag{4}$$

where $r$ is the discount rate used to value costs and benefits in different periods. We can now determine quantity and price policies that will maximize these net benefits. We then derive the relative advantage of price versus quantity policies as the difference in their expected net benefits, and comment on how relative policy performance is influenced by the shape of the cost and benefit functions, discount and decay rates, cost correlation, and the duration of the policies being compared.

We determine the optimal quantity policy at time $t$ by maximizing the expected net benefits of stock control (Eq. (4)) with respect to $q_t$, subject to the stock accumulation relationship given in Eq. (3) (see [23]). The resulting first-order conditions yield

$$E[C_t'(q_t^*, \theta_t)] = E\left[\sum_{s=t}^{\infty} B_s'(S_s^*) \frac{(1-\delta)^{s-t}}{(1+r)^{s-t}}\right]. \tag{5}$$

That is, the expected sum of discounted and depreciated marginal benefits associated with the output quantity in a particular period should equal the expected marginal cost.[13]

Now consider price-based mechanisms. Since costs and benefits depend on $q_t$ and $S_t$, the first step in calculating the optimal price policy is to determine the quantity that would result from a particular price policy. This *response function* $q_t(P_t, \theta_t)$ depends on both the price $P_t$ set by the

---

[12] The cost variance at time $t$, $\mathrm{var}(\theta_t)$, is therefore given by $\mathrm{var}(\theta_t) = \sigma_t^2 = \sigma_0^2(1-\rho^{2t})/(1-\rho^2)$, assuming that $\theta_{t<1} = 0$.

[13] As noted earlier, in order to simplify the model and presentation we do not consider the possibility that future policies adjust to past ones. For that reason, policies today have consequences far into the future as demonstrated in Eq. (5). Discounting and stock decay diminishes the importance of this assumption.

regulator and the cost shock $\theta_t$. The price policy will take the form of a subsidy in the case of a positive externality and a tax in the case of a negative externality. Assuming firms respond to the price policy $P_t$ by equating it to the marginal cost of output, $\theta_t + c_t(q_t - \bar{q}_t)$, the response function will be given by

$$q_t(P_t, \theta_t) = \bar{q}_t + \frac{P_t - \theta_t}{c_t}. \tag{6}$$

The optimal price policy at time $t$ maximizes expected net benefits with respect to $P_t$. The optimal response function is (see [23]):

$$q_t(P_t^*, \theta_t) = q_t^* - \frac{\theta_t}{c_t}, \tag{7}$$

where $-\theta_t/c_t$ is the quantity deviation resulting from the cost shock in period $t$. In other words, the optimal price policy delivers the optimal quantity level "on average", but delivers more or less than the quantity policy as marginal costs fall and rise, respectively, thereby taking advantage of low costs and avoiding high costs.

As a consequence of deviations in the quantities induced by the price policy, the stock level associated with the price policy will also be a function of the realized cost shocks:

$$S_t = S_t^* - \sum_{s=0}^{t} (1-\delta)^s \frac{\theta_{t-s}}{c_{t-s}}, \tag{8}$$

where $S_t^*$ is the stock associated with the optimal quantity policy. Since differences in the costs and benefits of alternative policies depend only on differences in the flows and stocks associated with those policies, Eqs. (7) and (8) set the stage for deriving the relative advantage of prices versus quantities.

## 3. The relative advantage of price over quantity policies

Our basic metric for policy comparisons builds on Weitzman's method of computing the "comparative advantage" of price instruments over quantity instruments, by deriving and comparing the expected benefits and costs of these alternative policies.[14] Our approach reveals an interesting feature of the dynamic problem compared to a simple static model. That is, in order to compute the relative advantage from choosing prices rather than quantities in period $t$, we need to specify the policies in all other periods. In the simplest sense, this is necessary because shocks in the quantity produced under a price policy interact over time in the calculation of benefits—an interaction that depends on the instrument choice in other periods. In a more subtle sense, a careful accounting is required so that as we consider the choice of instrument over time, these interactions are counted—exactly once.

---

[14] The discussions of the cost and benefit differences hinge only on the price policy having the same expected quantity outcome as the quantity policy (and similarly, the quantity policy having the same expected price outcome as the price policy). Under those conditions, the expressions for the cost and benefit differences are equally applicable to both optimal and sub-optimal policy comparisons.

In particular, we compute the difference $\Delta_t$ in the expected net benefits of using a price rather than quantity policy in period $t$ as follows:

$$\Delta_t = E[NB_{t,\text{price}}] - E[NB_{t,\text{quantity}}], \tag{9}$$

where the "price" subscript represents the use of a price policy in period $t$, while the "quantity" subscript represents the use of quantity policy in period $t$. When measuring the net benefits for *both* the price and quantity policies in period $t$, it is assumed that a price policy is in place for all periods prior to $t$, and a quantity policy is in place for all periods after $t$. This accounting has the effect of measuring the *incremental* effect of continuing to use prices for one more period, rather than switching to quantities in period $t$. Thus, $\Delta_t > 0$ indicates that using the optimal price-based policy rather than the optimal quantity-based policy in period $t$ has higher net benefits, while $\Delta_t < 0$ indicates the reverse, assuming the use of price controls before period $t$ and quantity controls afterwards.

We use this ordering based on our sense that optimal climate policies may naturally sequence from price to quantity instruments over time, as we discuss further below. Alternative assumptions about the pattern of price and quantity use in other periods influences the magnitude of the correlation term (see $\Omega_{\rho,t}$ in Eq. (10)) but leave our qualitative results unchanged. Note that when we sum $\Delta_t$ over an infinite horizon, we arrive at an expression comparing the use of prices in all periods to the use of quantities in all periods.

Our basic result, derived in [23], is[15]

$$\Delta_t = \frac{\sigma_t^2}{2c_t^2}(c_t - b_t \Omega_\delta \Omega_{\rho,t}), \tag{10}$$

where the persistence factor $\Omega_\delta$ captures future benefit losses, and is given by

$$\Omega_\delta = \frac{1+r}{1+r-(1+g_b)(1-\delta)^2},$$

assuming that the benefit parameter $b_t$ grows at the constant rate $g_b$, whereas $\Omega_{\rho,t}$ accounts for correlated costs:

$$\Omega_{\rho,t} = 1 + \frac{c_t^2}{\sigma_t^2} 2(1-\delta) cov\left(-\frac{\theta_t}{c_t}, S_{t-1}\right).$$

From (8), we note that a larger correlation $\rho$ will raise $\Omega_{\rho,t}$, a point we consider in more detail in [23], where we assume that the cost parameter $c_t$ grows at the constant rate $g_c$.[16]

Prices are therefore preferred ($\Delta_t > 0$) if $c_t > b_t \Omega_\delta \Omega_{\rho,t}$, quantities are preferred ($\Delta_t < 0$) if $c_t < b_t \Omega_\delta \Omega_{\rho,t}$, and there is indifference between prices and quantities ($\Delta_t = 0$) if $c_t = b_t \Omega_\delta \Omega_{\rho,t}$. Our first observation is readily evident.

---

[15] This result can be viewed as an extension of Weitzman's [33] consideration of prices versus quantities with many production units and, in particular, his Eq. (26). In Weitzman's notation, we have considered only the cost savings from one unit along with the benefit consequences arising from correlations with all *previous* units (note that our "units" are ordered). This allows us to add up the consequences over different price policy durations without double counting, as noted previously.

[16] If the slope of marginal costs can change over time, one might inquire whether there should also be a trend to the cost variance. When we consider multiple-period policies in the next section, increasing variance would tend to raise the importance of future periods. As we soon discuss, this would make it more likely that quantity controls are preferred.

> More steeply sloped marginal costs tend to favor price controls, while more steeply sloped marginal benefits tend to favor quantity controls for regulating stock externalities.[17]

This observation reaffirms Weitzman's original result—the relative slopes of marginal costs and benefits continue to be fundamental to policy choice in the dynamic context of a stock externality. Quantities tend to be preferred in cases where strong nonlinearities or thresholds lead to steep marginal benefits. Less curvature in benefits tends to favor prices.

Now compare the above expression to Weitzman's original expression $\Delta = \frac{\sigma^2}{2c^2}(c-b)$. First, note that the sign of Weitzman's $\Delta$, which indicates the policy preference, depends *only* on the relative magnitude of $c$ and $b$ (i.e., the relative slopes). In contrast, the sign of Eq. (10) depends on the magnitude of $c_t$ relative to $b_t$ *multiplied by the persistence and correlation terms* $\Omega_\delta$ and $\Omega_{\rho,t}$. This is necessary because the production costs in a given period occur only in that period while the associated benefits persist into the indefinite future. In the absence of persistence—the special case of a flow externality—these two terms drop out and our expression reduces to Weitzman's formula. We therefore make two additional observations:

> Lower stock decay rates, lower discount rates, and greater rates of benefits growth tend to favor quantity controls for regulating stock externalities.

and

> Greater correlation in costs across time tends to favor quantity controls.

Under a price policy, slower stock decay causes price-induced deviations in the stock level to persist longer, thereby increasing the variability of the stock. This leads to lower expected benefits because benefits are a concave function of the stock level. Lower discount rates and greater rates of benefits growth will give these future losses greater weight. The persistence term $\Omega_\delta$ captures these effects. When there is positive correlation in the cost shocks ($\rho > 0$), the expected loss of benefits rises, as indicated by the correlation term $\Omega_{\rho,t}$ multiplying $b_t$. When costs are positively correlated across time, deviations in the stock arising under a price mechanism tend to accumulate rather than canceling out. This exacerbates variation in the stock level and further lowers expected benefits under a price policy. Greater correlation and slower depreciation will increase the covariance term, thereby increasing the magnitude of $\Omega_{\rho,t}$ (as shown in [23]).

Depending on the rates of time discounting and benefit growth, these two factors have the potential to greatly increase the relative importance of marginal benefits. In particular, regardless of how $b_t$ and $c_t$ compare, if we care enough about the future (e.g., $r$ near zero and/or large $g_b$) it is always possible that $c_t < b_t\Omega_\delta\Omega_{\rho,t}$, implying a preference for quantity controls. In the climate change policy debate, for example, this is one possible explanation for the persistent emphasis on quantity controls by some stakeholders.

It is apparent that the relative advantage expression (Eq. (10)) can change over time due to changes in $c_t$ and $b_t$, and due to accumulated correlation among cost shocks. As a consequence, it

---

[17] Note that our use of the words "tend to favor" and "make it more likely" is meant in a comparative static sense. It is not necessarily the case that low discount rates, low decay rates, highly correlated costs, or steep marginal benefits imply that quantities *are* preferred—or even that movements in those directions will lower the value of $\Delta_t$—but rather that movements in the stated direction will *widen the range* of conditions in which quantities are preferred.

will not always be optimal to use the same instrument in every period—it may be optimal to use price policies in some periods and quantity policies in others. In particular, it is reasonable to assume that costs decline over time ($g_c < 0$) due to technological improvements and benefits rise over time due to growth in the economy and population subject to the externality ($g_b > 0$) (see footnote 10). This implies an eventual preference for quantity controls in the future based on Eq. (10). As marginal costs fall, the cost savings under price policies become less important. Meanwhile, with marginal benefits rising, the stock certainty assured by quantity policies becomes more important.

Now recall that $\Delta_t$ represents the advantage of choosing price controls relative to choosing quantity controls in a *single* period $t$ (but counting the consequences in future periods of this policy choice in period $t$). In order to consider the relative advantage of fixing these policy instruments over a longer timeframe $T$—as in our application to climate change policy below—we simply compute the present value $\Delta^T$, where

$$\Delta^T = \sum_{t=1}^{T} (1+r)^{-t}\Delta_t, \tag{11}$$

where, as noted earlier, this discounted sum measures the cumulative effect of using prices rather than quantities prior to and including period $T$, assuming quantity controls are used in all later periods.

With policies fixed over longer timeframes and $\Delta_t$ changing over time, the choice of policy instrument may depend on how long the policy remains in place. For example, the choice between price and quantity controls may differ when considering a price policy for a 5-year versus a 40-year duration due to changes in benefits and costs. The likely direction of these effects—which we have noted seem to favor quantity controls in the future—implies that shorter-term policies could favor price controls while longer-term policies would favor quantities.

Although our focus at this point is on quadratic benefits that rise geometrically over time, we can use the model to qualitatively understand the potential effects of more dramatic consequences in the future due, for example, to thresholds beyond which stock consequences greatly increase. In the case of climate change, the possibility of melting polar ice caps followed by significant sea level rise is an example of such catastrophic effects. This situation can be captured by a large increase in the slope of marginal benefits in the future. Ceteris paribus, such an increase will tilt the balance toward quantity controls by raising the benefit loss and lowering $\Delta_t$. However, this effect will be diminished if it occurs in the future. Benefit losses in some future period from using prices in the current period depends not only on the increased severity of damages in the future, but also on the amount of time that will pass before the threshold is reached.

This distance in time matters for two reasons: decay and discounting.[18] Greater decay of price-induced shocks to the stock will lessen the effect of current policy on future benefits, and greater discounting will lessen the value of these future benefits. Therefore, the further we are from the threshold, the greater is the mitigating effect of stock decay and discounting. Just as our plausible growth assumptions ($g_c<0$ and $g_b>0$) suggest the use of a near-term price policy coupled with a quantity control in the future, the presence of stock thresholds may suggest a

---

[18] It also depends on the ability to correct fluctuations from a price policy at some point in the future—an issue we do not address in the current analysis.

similar strategy.[19] That is, provided the thresholds remain significantly beyond the initial planned duration of the price policy. In contrast, a more imminent threshold would tend to demand quantity controls now.

## 4. Application to climate change policy

We now apply the above modeling results to the case of regulating the stock of carbon dioxide in the atmosphere in order to mitigate the externality of global climate change. Table 1 presents the benchmark values we used for this application, based on currently available information.[20] In addition to the parameters discussed so far, we introduce $g_q$ to represent a constant growth rate in $\bar{q}_t$.

Using Eqs. (10) and (11) it is straightforward to compute $\Delta^T$ for various policy durations. With a single-period duration, for example,

$$\Delta^1 = \frac{\sigma_1^2}{2(c_0(1+g_c))^2}\left(c_0(1+g_c) - \frac{b_0(1+g_b)(1+r)}{1+r-(1-\delta)^2(1+g_b)}\right) = \$520 \text{ million.}$$

That is, a price instrument used to regulate carbon dioxide emissions generates an expected \$520 million gain relative to a quantity instrument in a single year (but counting the consequences in all future years). To get a better sense of the relative gain from using a price policy, we can also find the welfare gain associated with a quantity-based policy and compute the percentage difference. To do this, we first numerically compute the optimal deterministic control path. We then compare the net benefits of this optimized quantity policy over a particular horizon versus the no-policy alternative. In the first period, we find a discounted marginal benefit of 9 \$/ton (which is comparable to the estimates in [24]) and total net discounted benefits of \$225 million. The price policy therefore generates over three times the expected net benefits of a comparable quantity policy in the first period.

Table 2 also shows results for alternative policy durations of 5, 10, 20, and 40 years, where the instrument is fixed for the indicated period but consequences are computed over an infinite horizon. The results indicate that with longer policy durations, the price advantage rises for climate change policy. With a 40-year duration, price controls generate \$35 billion in higher expected benefits compared to quantity controls, although this is now "only" 2 times the net benefits of the quantity policy. We find that prices generate up to nearly 5 times the

---

[19] Our model ignores the possibility that we could react to price-induced shocks after they occur and before reaching the threshold, further enhancing the case for price controls in the short term.

[20] See [23] for much greater detail on data sources and our basis for the range of sensitivity analyses we consider. All monetary values are denominated in 1998 US dollars, if necessary adjusted using the price index for gross domestic product [4]. $S_0$ is based on data from [11,22]; $\bar{q}_0$is based on data from [9,20,24]; $g_q$ is based on IPCC scenario IS92a [9]; $b_0$ is based on [25] as reported in [30]; $c_0$ is based on results from 10 models that participated in the Energy Modeling Forum's EMF 16, as reported in [36]; $g_b$ is the annualized growth rate in GWP from [19] for the central IPCC scenario IS92a; $\sigma_0$ is also based on modeling results from [36] and a first-order autoregressive maximum-likelihood model of global carbon emissions using data from [20]; $\delta$ is from [24]; and $r$ is the central value from [26].

Table 1
Information for analysis of climate change policy

| Parameter | Annualized value |
|---|---|
| Decay rate of stock ($\delta$) | 0.83% |
| Discount rate ($r$) | 5.0% |
| Marginal benefit slope ($b_0$) | $8.7 \times 10^{-13}$ \$/ton$^2$ |
| Marginal cost slope ($c_0$) | $1.6 \times 10^{-7}$ \$/ton$^2$ |
| Cost uncertainty ($\sigma_0$) | 13 \$/ton |
| Cost correlation ($\rho$) | 0.80 |
| Benefit growth rate ($g_b$) | 2.5% |
| Cost growth rate ($g_c$+$g_q$) | –1.0% |
| Baseline emissions growth rate ($g_q$) | 1.5% |
| Initial stock ($S_0$) | $1.7 \times 10^{11}$ tons |
| Initial emissions ($\bar{q}_0$) | $5.0 \times 10^{9}$ tons |

*Note*: $ refers to 1998 US dollars and tons refers to metric tons of carbon. See footnote 20 for detail on data sources for parameter values.

Table 2
Relative advantage of prices over quantities for optimal climate change policy

| Price policy | Expected price advantage | | Benefits required for indifference | |
|---|---|---|---|---|
| duration (years) | \$ billions | $NB_{\text{price}}/NB_{\text{quantity}}$ | \$/ton$^2$ | Relative to benchmark |
| 1 | 0.52 | 3.3 | $6.1 \times 10^{-9}$ | 7300 |
| 5 | 4.6 | 4.5 | $1.7 \times 10^{-9}$ | 2000 |
| 10 | 11 | 4.8 | $9.8 \times 10^{-10}$ | 1200 |
| 20 | 21 | 3.7 | $5.7 \times 10^{-10}$ | 680 |
| 40 | 35 | 2.2 | $3.0 \times 10^{-10}$ | 360 |

*Note*: $ refers to 1998 US dollars.

expected welfare gains of quantities, with consequences on the order of many billions of dollars per year.[21]

Given that there is a wide range of opinion surrounding the benefits of climate change mitigation, we can also ask how large the slope of marginal benefits would have to be in order for us to be indifferent between prices and quantities. We use the notation $(b_0/c_0)_{\text{crit}}$ to denote the ratio of marginal benefit and cost slopes at which we are indifferent between prices and quantities, conditional on the remaining parameters. For the given parameters, this "critical relative slope" defines a relative benefit level above which quantities are preferred and below which prices are preferred. In the case of a policy with a single-period duration, the critical relative slope is obtained by rearranging Eq. (10) when $\Delta_t = 0$, obtaining $(b_0/c_0)_{\text{crit}} = 1/(\Omega_\delta \Omega_{\rho,t})$. The benefit

---

[21] One might be confused by our qualitative statement that longer durations of the price policy favor quantitative controls and the quantitative result showing that longer durations raise the gain to price controls. The explanation is the difference between the expression determining the sign of $\Delta_t$ and the expression determining its *magnitude*, a point we discuss further below.

slope itself is obtained by multiplying $(b_0/c_0)_{\text{crit}}$ by the estimate of $c_0$ in Table 1. As shown in Table 2, we find that there would have to be at least a 7300-fold increase in the slope of marginal benefits (relative to the benchmark value of $b_0$ in Table 1) in order for prices and quantities to generate the same expected net benefits over a single year.

With a longer duration $T$, we can find $\Delta^T = 0$ from Eq. (11), obtaining the functional relationship

$$(b_0/c_0)_{\text{crit}} = (b_0/c_0)_{\text{crit}}(T, \rho, g_b, g_c, r, \delta), \tag{12}$$

where the parameter $\sigma_0^2$ vanishes since it only scales $\Delta^T$. These values must be computed numerically. From Table 2, we can see that the benefit condition for indifference is relaxed over longer policy durations, though it is still quite high. A 360-fold increase in the slope of marginal benefits relative to costs is required for indifference between price and quantity controls even over a 40-year duration. Thus, despite the considerable uncertainty surrounding the consequences of global climate change, the advantage of prices over quantities remains unless the true benefits of carbon mitigation are many orders of magnitude greater than our best estimate.

We now focus on a 10-year policy duration and further analyze the sensitivity of the relative advantage to variation in key parameters. We vary the parameter values one variable at a time, holding the other parameters at the base values given in Table 1.[22] Table 3 displays the range of values used in the sensitivity analysis, with scenario 3 being the base case given in Table 1. The range of values considered in the sensitivity analysis is very broad, and more than covers the plausible range of values for these key parameters. Table 4 provides the results of this sensitivity analysis in terms of billions of dollars that the price policy delivers in excess of the quantity policy. Table 5, on the other hand, provides the sensitivity results as a ratio of the net benefits of the price policy to the quantity policy. Positive values in Table 4 therefore indicate a preference for the price policy, while in Table 5 values greater than 1 indicate a price preference.

The main message from Tables 4 and 5 is that a price policy continues to be preferred to a quantity policy for climate change mitigation over a very wide range of values for key parameters—in fact, all of those we consider. Regarding specific results, we must point out that the scenarios are ordered so that higher numbered scenarios are "more likely" to be associated with a preference for price policies. That is, according to the comparative static results given in Section 3 regarding the *sign* of the relative advantage expression, which depends only on the portion of the expression within the parentheses of Eq. (10). However, the results presented in Tables 4 and 5 are for the *magnitude* of the full relative advantage expression. As a consequence, the results for sensitivity analysis of variables that also multiply outside the parentheses, that is $c_0$, $g_c$, and $\rho$, are ordered in a direction that might appear counterintuitive (a similar explanation covers the earlier results concerning policy duration). This points to the fact that within a Weitzman prices versus quantities framework, it is necessary to distinguish between the effect of alternate parameter values on the "likelihood" of a preference for price or quantity controls and the effect of those values on the magnitude of $\Delta_t$. Finally, it is important to keep in mind that these

---

[22] A more complete, although much more complex analysis, could optimize the instrument choice over these parameters, while treating them as uncertain as is done with marginal costs.

Table 3
Parameter values for sensitivity analysis

| Scenarios | $\delta$ (%) | $r$ (%) | $b_0$ | $c_0$ | $\rho$ | $g_b$ (%) | $g_c$ (%) |
|---|---|---|---|---|---|---|---|
| 1 | 0.01 | 1 | $8.7\times10^{-11}$ | $1.6\times10^{-9}$ | 0.99 | 3.5 | −3.5 |
| 2 | 0.10 | 3 | $8.7\times10^{-12}$ | $1.6\times10^{-8}$ | 0.90 | 3.0 | −3.0 |
| 3 (base) | 0.83 | 5 | $8.7\times10^{-13}$ | $1.6\times10^{-7}$ | 0.80 | 2.5 | −2.5 |
| 4 | 10 | 7 | $8.7\times10^{-14}$ | $1.6\times10^{-6}$ | 0.40 | 0.0 | 0.0 |
| 5 | 50 | 9 | $8.7\times10^{-15}$ | $1.6\times10^{-5}$ | 0.00 | −1.0 | 1.0 |

Table 4
Sensitivity analysis of the relative advantage of prices ($ billions)

| Scenarios | $\delta$ | $r$ | $b_0$ | $c_0$ | $\rho$ | $g_b$ | $g_c$ |
|---|---|---|---|---|---|---|---|
| 1 | 10 | 13 | 9.6 | 956 | 23 | 10 | 11 |
| 2 | 10 | 12 | 10 | 104 | 15 | 10 | 11 |
| 3 (base) | 11 | 11 | 11 | 11 | 11 | 11 | 11 |
| 4 | 11 | 9.4 | 11 | 1.1 | 5.4 | 11 | 9.1 |
| 5 | 11 | 8.5 | 11 | 0.10 | 4.7 | 11 | 8.6 |

Table 5
Sensitivity analysis of the relative advantage of prices ($NB_{\text{price}}/NB_{\text{quantity}}$)

| Scenarios | $\delta$ | $r$ | $b_0$ | $c_0$ | $\rho$ | $g_b$ | $g_c$ |
|---|---|---|---|---|---|---|---|
| 1 | 2.1 | 1.0 | 1.0 | 9.2 | 9.0 | 2.6 | 5.0 |
| 2 | 2.3 | 1.5 | 1.1 | 6.7 | 6.1 | 3.4 | 4.8 |
| 3 (base) | 4.8 | 4.8 | 4.8 | 4.8 | 4.8 | 4.8 | 4.8 |
| 4 | 584 | 13 | 287 | 3.9 | 2.9 | 21 | 3.8 |
| 5 | 45,282 | 27 | 23,692 | 3.5 | 2.6 | 36 | 3.6 |

results are conditional on the assumptions that costs and benefits are well-behaved, smooth functions, an assumption violated by a dramatic, near-term damage threshold.[23]

## 5. Conclusions

Our results extend to the case of stock externalities seminal work by Weitzman distinguishing between otherwise equivalent price and quantity controls when uncertainty exists about control costs. His original conclusion, that price mechanisms are more efficient when marginal benefits are relatively flat and quantity mechanisms are more efficient when benefits are relatively steep, carries over to the case of stock externalities considered in this paper. Flatter benefits continue to

---

[23] Of course, in the case of a dramatic, near-term damage threshold, there are other problems besides the choice of policy instrument.

favor price controls, but the story is complicated in several ways. It is no longer a simple relative slopes argument. The slope of the cost curve must instead be compared to an adjusted measure of marginal benefit, which takes into account growth, discounting, depreciation, and correlation of cost shocks. In addition to the obvious application to stock pollutants, these results could usefully be applied to issues of species preservation, land-use policy, education, research, highways, and national defense as areas where policymakers wish to regulate a stock-like externality.

Regarding climate change, these results have important implications for current policy discussions, including the Kyoto Protocol, the European Commission's recent proposal for an European Community greenhouse gas emissions trading program, along with other proposals for domestic and international carbon/greenhouse gas trading systems. The application of our relative advantage expression finds that price-based instruments for carbon reduction—such as a carbon tax—are likely to generate several times the expected welfare gains of quantity-based instruments, such as tradable carbon permits. Yet, programs like the Kyoto Protocol require binding, quantity-based reductions. At a minimum, this suggests that there could be great value in incorporating price elements into quantity-based policies, such as a cap on emission permit prices.

## References

[1] E. Anderson, Taxes vs. quotas for regulating fisheries under uncertainty: a hybrid discrete-time continuous-time model, Mar. Res. Econom. 3 (1986) 183–207.

[2] R.A. Androkovich, K.R. Stollery, Tax versus quota regulation: a stochastic model of the fishery, Amer. J. Agric. Econom. 73 (2) (1991) 300–308.

[3] F.A. Benford, On the dynamics of the regulation of pollution: incentive compatible regulation of a persistent pollutant, J. Environ. Econom. Management 36 (1998) 1–25.

[4] Council of Economic Advisors, 1999 Economic Report of the President, Council of Economic Advisors, Washington, DC, 1999.

[5] P. Dasgupta, P. Hammond, E. Maskin, On imperfect competition and pollution control, Rev. Econom. Stud. 47 (1980) 857–981.

[6] I. Falk, R. Mendelsohn, The economics of controlling stock pollutants: an efficient strategy for greenhouse gases, J. Environ. Econom. Management 25 (1993) 76–88.

[7] M. Hoel, L. Karp, Taxes versus quotas for a stock pollutant, Unpublished Manuscript, University of Oslo, Oslo, 1998.

[8] M. Hoel, L. Karp, Taxes and quotas for a stock pollutant with multiplicative uncertainty, J. Pub. Econom. 82 (1) (2001) 91–114.

[9] A. Kattenberg, F. Giorgi, H. Grassl, G.A. Meehl, J. Mitchell, R.J. Stouffer, T. Tokioka, A.J. Weaver, T.M. Wigley, Climate models—projections of future climate, in: J.T. Houghton, L.G. Meiro Filho, B.A. Callendar, N. Harris, A. Kattenberg, K. Maskell (Eds.), Climate Change 1995: The Science of Climate Change, Cambridge University Press, Cambridge, UK, 1996, pp. 285–358.

[10] E. Keeler, M. Spence, R. Zeckhauser, The Optimal Control of Pollution, J. Econom. Theory 4 (1971) 19–34.

[11] C.D. Keeling, T.P. Whorf, Atmospheric carbon dioxide record from mauna loa, Trends: A Compendium of Data on Global Change, Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, U.S. Department of Energy, Oak Ridge, TN, 1999.

[12] Y. Kitabatake, Optimal exploitation and enhancement of environmental resources, J. Environ. Econom. Management 16 (1989) 224–241.

[13] C. Kling, J. Rubin, Bankable permits for the control of environmental pollution, J. Pub. Econom. 64 (1) (1997) 101–115.

[14] E.F. Koenig, Controlling stock externalities in a common property fishery subject to uncertainty, J. Environ. Econom. Management. 11 (1984) 124–138.

[15] E.F. Koenig, Fisheries regulation under uncertainty: a dynamic analysis, Mar. Res. Econom. 1 (1984) 193–208.

[16] C.D. Kolstad, Learning and stock effects in environmental regulation: the case of greenhouse gas emissions, J. Environ. Econom. Management 31 (1) (1996) 1–18.

[17] E. Kwerel, To tell the truth: imperfect information and optimal pollution control, Rev. Econom. Stud. 44 (3) (1977) 595–601.

[18] J.J. Laffont, More on prices vs. quantities, Rev. Econom. Stud. 44 (1977) 177–182.

[19] J. Leggett, W.J. Pepper, R.J. Swart, Emission Scenarios for the IPCC, in: J.T. Houghton, B.A. Callander, S.K. Varney (Eds.), Climate Change 1992: The Supplementary Report to the IPCC Scientific Assessment, Cambridge University Press, Cambridge, UK, 1992, pp. 75–95.

[20] G. Marland, T.A. Boden, R.J. Andres, A.L. Brenkert, C. Johnston, Global, regional, and national $CO_2$ emissions, Trends: A Compendium of Data on Global Change, Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, U.S. Department of Energy, Oak Ridge, TN, 1999.

[21] W.J. McKibbin, P.J. Wilcoxen, A better way to slow global climate change, Policy Brief #17, The Brookings Institution, Washington, DC, 1997.

[22] A. Neftel, H. Friedli, E. Moor, H. Lötscher, H. Oeschger, U. Siegenthaler, B. Stauffer, Historical carbon dioxide record from the siple station ice core Trends: A Compendium of Data on Global Change, Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, U.S. Department of Energy, Oak Ridge TN, 1999.

[23] R.G. Newell, W.A. Pizer, Regulating stock externalities under uncertainty, RFF Discussion Paper 99-10 (Revised), Resources for the Future, Washington, DC, 1999.

[24] W.D. Nordhaus, Managing the Global Commons, MIT Press, Cambridge, 1994.

[25] W.D. Nordhaus, Expert Opinion on Climate Change, Amer. Sci. 82 (1994) 45–51.

[26] W.A. Pizer, Combining price and quantity controls to mitigate global climate change, J. Pub. Econom. 85 (3) (2002) 409–434.

[27] C. Plourde, A model of waste accumulation and disposal, Canad. J. Econom. 5 (1) (1972) 119–125.

[28] C. Plourde, D. Yeung, A model of industrial pollution in a stochastic environment, J. Environ. Econom. Management 16 (2) (1989) 97–105.

[29] M.J. Roberts, M. Spence, Effluent charges and licenses under uncertainty, J. Pub. Econom. 5 (3,4) (1976) 193–208.

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

[31] V.L. Smith, Dynamics of waste accumulation: disposal versus recycling, Quart. J. Econom. 86 (1972) 600–616.

[32] R.N. Stavins, Correlated uncertainty and policy instrument choice, J. Environ. Econom. Management 30 (1996) 218–232.

[33] M.L. Weitzman, Prices vs. quantities, Rev. Econom. Stud. 41 (4) (1974) 477–491.

[34] M.L. Weitzman, Optimal rewards for economic regulation, Am. Econom. Rev. 68 (1978) 683–691.

[35] M.L. Weitzman, Landing fees vs. harvest quotas with uncertain fish stocks, J. Environ. Econom. Management 43 (2) (2002) 325–338.

[36] J.P. Weyant, J. Hill, Introduction and overview, in: J.P. Weyant (Ed.), The Costs of the Kyoto Protocol: A Multi-Model Evaluation, Special Issue of the Energy Journal, International Association for Energy Economics, Cleveland, OH, 1999.

[37] R.C. Williams, Should emissions permit programs allow trading and banking? Working Paper, University of Texas at Austin, 2001.

[38] A.J. Yates, Decentralization in pollution permit markets, J. Pub. Econom. Theory 4 (4) (2002) 641–660.

[39] A.J. Yates, M.B. Cronshaw, Pollution permit markets with intertemporal trading and asymmetric information, J. Environ. Econom. Management 42 (1) (2001) 104–118.

[40] G.W. Yohe, Towards a general comparison of price controls and quantity controls under uncertainty, Rev. Econom. Stud. 45 (1978) 229–238.