



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

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

Help ensure our sustainability.

Give to AgEcon Search

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

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

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

Regulating Stock Externalities Under Uncertainty

Richard G. Newell and William A. Pizer

Original version: June 1998

Revised: December 1998, Revised: May 2000

Discussion Paper 99-10-REV



Resources for the Future
1616 P Street, NW
Washington, D.C. 20036
Telephone: 202-328-5000
Fax: 202-939-3460
Internet: <http://www.rff.org>

© 2000 Resources for the Future. All rights reserved. No portion of this paper may be reproduced without permission of the authors.

Discussion papers are research materials circulated by their authors for purposes of information and discussion. They have not undergone formal peer review or the editorial treatment accorded RFF books and other publications.

Regulating Stock Externalities Under Uncertainty

Richard G. Newell and William A. Pizer

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. 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. As in Weitzman (1974), 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.

We also demonstrate an important link between instrument choice and policy stringency, based on the observation that both of these elements of policy design depend on the same underlying information, namely the marginal benefit and cost slopes. This result is especially useful when there is disagreement about benefits, since it can be used to restrict the set of efficient policies even in such cases. For negative externalities, we find that less stringent controls and price instruments are both associated with flatter marginal benefits, while aggressive controls and quantity instruments are associated with steeper marginal benefits. Intuitively, damages (marginal benefits) can only be steep enough to recommend quantity controls in the near term if they are steep enough to recommend substantial reductions. Furthermore, as long as the existing stock is large relative to the annual flow, marginal benefits either (i) will appear very flat over range of emissions in a single year, or (ii) will be sufficiently high to warrant near complete abatement. This generic characteristic of what it means to be a stock externality therefore weighs heavily in favor of price instruments for their control, so long as the optimal control falls short of stabilization at the current stock level.

In the case of greenhouse gas policy, we show that any benefit slope consistent with quantity controls would imply an optimal emission level of *zero*. Under more general conditions we further demonstrate that price instruments are preferred in cases where optimal abatement rates are below about 40 percent, unless the initial stock level is small (less than twenty times the flow rate) or benefits and costs diverge quickly (by more than 10 percent annually). This result has important implications for the Kyoto Protocol, which mandates reductions among industrialized countries equal to roughly 5 percent of forecast global emissions. Regardless of one's beliefs about various climate change parameters, these relatively low aggregate abatement levels are inconsistent with quantity-based emission limits.

Key words: Stock, Externality, Regulation, Policy, Uncertainty, Price, Quantity, Tax, Tradable Permit, Pollution, Climate Change, Greenhouse Gas, Instrument Choice

JEL Classification Numbers: Q28, D81, C61

Contents

Abstract	ii
Contents	iii
I. Introduction	1
A. Policy Choice Under Uncertainty	2
B. Stock Regulation Under Uncertainty	3
II. A Model of Stock Externality Regulation Under Uncertainty	4
A. Benefits, Costs, and Stock Accumulation	6
B. Optimal Quantity and Price Policies	7
III. The Relative Advantage of Price Over Quantity Policies	9
A. Cost Savings from Price Policy	9
B. Benefit Losses from Price Policy	10
C. The Relative Advantage of Prices	12
D. Instrument Choice Over Time	14
E. Application to Climate Change Policy	15
IV. Instrument Choice When Benefits Are Not Directly Known	17
A. Establishing the Link Between Policy Instrument and Stringency	18
B. Calculating Critical Output Levels Above which Prices are Preferred	20
C. Application to Climate Change Policy	21
D. Sensitivity Analysis and the General Case of a Negative Stock Externality	21
V. Conclusions	24
Appendix	26
A. Details on Correlation Factor	26
B. Data for Climate Policy Application and Sensitivity Analysis	27
C. Establishing the Link Between Policy Instrument and Stringency	31
D. Signing the Derivatives of $(b_0/c_0)_{\text{crit}}$ and q_1^*/\bar{q}_1	32
References	37

Regulating Stock Externalities Under Uncertainty

Richard G. Newell and William A. Pizer*

I. Introduction

The threat of global climate change is one of the most important and challenging problems currently facing the human race. Without a concerted international effort to reduce manmade emissions of heat-trapping greenhouse gases in the coming decades, the world could face climatic changes with potentially profound impacts on the global population and economy. Global climate change is a stock externality: the consequences depend not on the emissions of greenhouse gases in a single year, but on the accumulated stock of emissions over many decades. 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 contemplating regulation of these problems, policymakers face the inevitable and important 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 linked to global warming, with the ultimate goal of limiting the atmospheric stock of greenhouse gases. 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. 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. Examples of positive 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.

Regulating such stocks involves considerable uncertainty. The magnitude of the associated benefits and costs of control are often known only approximately, while increases to the stock in any single year can persist far into the future, introducing additional uncertainty about conditions and valuation in those future periods. Even as regulated firms learn about costs and respond to an announced

* Newell and Pizer are Fellows at Resources for the Future (RFF). The paper has benefited from helpful comments from Larry Karp, Preston McAfee, 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.

price or quantity 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.¹

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 (1994a) observes that damages in his climate change policy simulations are essentially a linear function of emissions and, based on an appeal to Weitzman (1974), that this implies a preference for price instruments. He also suggests that this preference might extend to stock pollutants more generally. Pizer (1997), on the other hand, directly investigates the price versus quantity question for climate change using Monte Carlo simulations. Like Nordhaus and Kolstad (1996), he observes linear damages, but is further able to demonstrate that price policies indeed lead to much better welfare outcomes. Finally, McKibbin and Wilcoxen (1997) 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 examines 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.

A. 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., tradeable 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, such policies would be of little if any practical use. Recognizing this, Weitzman (1974) 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 fairly basic and intuitive economic variables, including the slopes of marginal benefits and costs and the degree of cost uncertainty.

¹ Laffont (1977) 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 (Kwerel 1977; Dasgupta, Hammond, and Maskin 1980).

When uncertainty exists about costs, and policies must be fixed before the uncertainty is resolved, priced-based policies will lead to distinctly different outcomes than quantity-based policies. Pollution taxes, for example, operate by encouraging firms to reduce emissions as long as the marginal cost of reducing an additional unit of pollution is less than the level of the pollution tax. The tax mechanism will lead to a range of possible emission levels across different cost outcomes—but will fix marginal cost at the level of the tax. Conversely, a quantity-based policy such as a tradable permit system will fix the total volume of emissions, with the equilibrium permit price determined by the marginal cost of the last reduction necessary to meet the emission constraint. The permit mechanism will therefore lead to a range of possible permit prices (and marginal costs) across different cost outcomes—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, relatively flat marginal benefits imply a constant benefit per unit, suggesting that a tax could perfectly correct the externality. In contrast, steep marginal benefits imply a dangerous threshold that should be avoided at all costs—a threshold that is efficiently enforced by a quantity control. Despite its substantial insight, however, the Weitzman result remains a static story. Additional modeling is necessary as we turn to the specific question of controlling stock externalities, where consequences occur in a dynamic setting.

B. 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.² 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. In contrast, we use an otherwise simplified model to explore how uncertainty could influence the choice of policies to regulate a periodic flow when benefits are a function of the accumulated stock. Our results add several important

² Most economic research on stock pollutants has focused on optimal control in a deterministic setting (Keeler, Spence, and Zeckhauser 1971; Plourde 1972; Smith 1972; Kitabatake 1989; Falk and Mendlesohn 1993). This removes the distinction among alternative instruments and sidesteps the issue of choosing an efficient instrument. Plourde and Yeung (1989) 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 (Koenig 1984a, 1984b; Anderson 1986; Androkovitch and Stollery 1991). However, parsimonious, intuitive results regarding efficient policy for controlling other types of stocks are obscured in the analyses by the particulars of optimal fishery modeling. In isolation from our own work, Hoel and Karp (1998) have recently written about regulating a stock pollutant under uncertainty, with a focus on the frequency of adjustment of firms and policy to new information. Otherwise, see Stavins (1996) for a review of the literature on efficient policy instrument choice under uncertainty.

new dimensions to the existing view of the problem, with potential relevance to a wide variety of problems where policy interventions might be justified based on the existence of positive or negative stock externalities.

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 prices relative to quantities for controlling stock externalities. As in Weitzman (1974), 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 gas reductions to mitigate the effects of climate change, we find a price advantage of almost five times the welfare gain associated with quantity-based controls.

But we can say more, especially when one considers the disagreement surrounding the benefits of climate change mitigation. Holding other parameters constant, we show that flatter benefits not only favor price policies, but also lead to lower abatement levels. We then characterize exactly how low the optimal abatement level must be before it reveals underlying benefits that recommend price policies. In our climate example, we find that optimal abatement levels even marginally below 100 percent are sufficient to recommend price policies. More importantly, we demonstrate optimal emission reductions of less than about 40 percent are sufficient to reveal a general preference for price policies, unless the stock is particularly small or there is a rapid divergence between benefits and costs over time. Returning to the question of optimal climate change policy, we argue that the 5 percent aggregate reductions mandated by the Kyoto Protocol are inconsistent with quantity-based regulation regardless of parameterization.

In section 2, we present our model of stock externality regulation under uncertainty and describe the optimal quantity and price policies. Section 3 derives an expression for the difference in the expected net benefits of price versus quantity policies and applies the result to the problem of global climate change. Section 4 develops the connection between the optimal control path and instrument choice and again considers the climate change example. Details of the data used in the climate change example as well as detailed analytical derivations are left to the appendix.

II. 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 (1974),³ 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 (Roberts and Spence 1976, Weitzman 1978), but these are rarely, if ever, seen in practice. Without loss of generality, we also 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.⁴ 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 (Weitzman 1974, Stavins 1996), in a multi-period setting results will also hinge on correlation among costs in different periods. This correlation represents persistence in 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.

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.

³ Weitzman (1974) 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. The results can be generalized to a path of quadratic cost-benefit approximations.

⁴ 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 (1974) and Stavins (1996) 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.

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.

A. Benefits, Costs, and Stock Accumulation

Based on these choices, we specify benefits as

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

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 ($b_t \geq 0$), 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 , 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 atmospheric carbon). 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

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

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 θ_t is a shock to the marginal cost function.⁵ Potential changes in c_t and \bar{q}_t allow for cost reductions and growth in uncontrolled output.

⁵ 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

The cost shock θ_t has an autoregressive form $\theta_t = \rho\theta_{t-1} + \varepsilon_t$, with correlation ρ across time, and error ε_t with zero mean and variance σ_0^2 . The cost variance at time t , $\text{var}(\theta_t)$, is therefore given by

$$\text{var}(\theta_t) = \sigma_t^2 = \frac{\sigma_0^2(1-\rho^{2t})}{(1-\rho^2)},$$

assuming that $\theta_{t-1} = 0$. 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 (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,

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

which specifies that the stock decays at rate $0 \leq \delta \leq 1$ in addition to the contribution of q_t . The depreciation 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.

B. Optimal Quantity and Price Policies

We first determine the conditions under which a quantity policy for period t will maximize expected net benefits NB , where

$$(4) \quad NB_t = \sum_{s=t}^{\infty} \frac{B_s(S_s)}{(1+r)^s} - C_t(q_t, \theta_t).$$

Next we consider the optimal price policy. 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, and cost correlation.

We determine the optimal quantity policy at time t by maximizing the expected net benefits of stock control (Equation (4)) with respect to q_t , subject to the stock accumulation relationship given in Equation (3). The Lagrangian form of this optimization problem is

$$(5) \quad \max_{q_t} E \left[\sum_{s=t}^{\infty} \frac{B_s(S_s) - \lambda_s (S_s - (1-\delta)S_{s-1} - q_s)}{(1+r)^{s-t}} - C_t(q_t, \theta_t) \right],$$

where $E[\cdot]$ is the expectations operator. This yields two first order conditions in addition to the stock constraint:

$$(6) \quad E[-C'_t(q_t^*, \theta_t)] + \lambda_t = -c_t (q_t^* - \bar{q}_t) + \lambda_t = 0$$

Yohe (1978), this will tend to favor price controls.

$$(7) \quad E[B'_s(S_s^*)] - \lambda_s + \frac{1-\delta}{1+r} \lambda_{s+1} = -b_s (S_s^* - \bar{S}_s) - \lambda_s + \frac{1-\delta}{1+r} \lambda_{s+1} = 0.$$

where Equation (7) holds for all $s \geq t$, and q_t^* and S_t^* are the optimal quantity control and resulting stock level, respectively. Note that Equation (7) can be recursively substituted and then combined with Equation (6) to yield

$$(8) \quad E[C'_t(q_t^*, \theta_t)] = E\left[\sum_{s=t}^{\infty} B'_s(S_s^*) \frac{(1-\delta)^{s-t}}{(1+r)^{s-t}}\right].$$

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.

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 regulator and the cost shock θ_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

$$(9) \quad q_t(P_t, \theta_t) = \bar{q}_t + \frac{P_t - \theta_t}{c_t}.$$

To compute the optimal price policy P_t^* , we employ the Lagrangian given by Equation (5) with two changes: (i) the choice variable is now P_t , rather than q_t ; and (ii) we substitute the response function (Equation (9)) for q_t . The resulting first order conditions are

$$\begin{aligned} E\left[(-C'_t(q_t, \theta_t) + \lambda_t)q'_t(P_t^*)\right] &= (-c_t(E[q_t] - \bar{q}_t) + \lambda_t)(1/c_t) = 0 \\ E[B'_s(S_s)] - \lambda_s + \frac{1-\delta}{1+r} \lambda_{s+1} &= -b_s (E[S_s] - \bar{S}_s) - \lambda_s + \frac{1-\delta}{1+r} \lambda_{s+1} = 0. \end{aligned}$$

Since $q'_t(P_t) = 1/c_t$ is a constant (see Equation (9)), this term will drop out of the first condition and yield the same first order conditions for the optimal price policy as for the optimal quantity policy, in expectation. Thus, the optimal price policy P_t^* must result in the optimal quantity policy q_t^* on average (that is, $E[q_t(P_t^*, \theta_t)] = q_t^*$), and thus follows $P_t^* = c_t(q_t^* - \bar{q}_t)$. The optimal response function is therefore

$$(10) \quad q_t(P_t^*, \theta_t) = q_t^* - \frac{\theta_t}{c_t},$$

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:

$$(11) \quad S_t = S_t^* - \sum_{s=0}^t (1-\delta)^s \frac{\theta_{t-s}}{c_{t-s}},$$

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, Equations (10) and (11) set the stage for deriving the relative advantage of prices versus quantities.

III. 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. This approach involves deriving and comparing the expected benefits and costs of alternative policies.⁶ In particular, we compute the difference Δ_t in the expected net benefits of using a price rather than quantity policy in period t as follows:

$$(12) \quad \Delta_t = E[NB_{t, \text{price}} - NB_{t, \text{quantity}}] = E[(B_{t, \text{price}} - B_{t, \text{quantity}}) - (C_{t, \text{price}} - C_{t, \text{quantity}})].$$

Thus, $\Delta_t > 0$ indicates that the optimal price-based policy performs better than the optimal quantity-based policy in period t , while $\Delta_t < 0$ indicates the reverse. Both derivation and intuition are easier if we first derive the difference in the expected costs of prices versus quantities, then derive the difference in expected benefits, and finally find Δ_t by taking the difference of these two differences (as shown in Equation (12)).

A. Cost Savings from Price Policy

The expected difference in costs each period is given by

⁶ 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.

$$\begin{aligned} E[C_{t, \text{price}} - C_{t, \text{quantity}}] &= E\left[C_t\left(q_t^* - \frac{\theta_t}{c_t}, \theta_t\right) - C_t\left(q_t^*, \theta_t\right)\right] \\ &= E\left[\theta_t\left(q_t^* - \frac{\theta_t}{c_t} - \bar{q}_t\right) + \frac{c_t}{2}\left(q_t^* - \frac{\theta_t}{c_t} - \bar{q}_t\right)^2 - \left(\theta_t\left(q_t^* - \bar{q}_t\right) + \frac{c_t}{2}\left(q_t^* - \bar{q}_t\right)^2\right)\right], \end{aligned}$$

where $-\theta_t/c_t$ is again the quantity deviation resulting from the use of a price policy in period t . Many terms cancel or drop out upon taking expectations, yielding

$$(13) \quad E[C_{t, \text{price}} - C_{t, \text{quantity}}] = E\left[-\frac{\theta_t^2}{2c_t}\right] = -\frac{c_t}{2} \text{var}\left(-\frac{\theta_t}{c_t}\right) = -\frac{\sigma_t^2}{2c_t}.$$

The cost difference expression is negative, indicating that the price policy leads to lower expected costs than the quantity policy. The price policy relaxes the quantity when costs are high and increases it when costs are low, resulting in lower average costs. The expected savings will depend on the magnitude of cost uncertainty and the slope of marginal costs. The expression also reveals the interesting result that the difference in costs depends only on the deviations of the price from the quantity policy, and not the level of the optimal quantity policy itself.

B. Benefit Losses from Price Policy

The expected difference in benefits also depends on the quantity deviations resulting from the price policy, but in a somewhat more complicated fashion since benefits are a function of the stock. In particular, use of a price policy in period t affects not only the stock in period t , but in every future period $t+u$. To properly account for these effects, we introduce notation representing the benefits in period $t+u$ of alternative policies in period t , $B_{t, \text{policy}}^{t+u}$, where “policy” is either price or quantity. The entire stream of benefits accruing to policy in period t , $B_{t, \text{policy}}$, will be the discounted sum of these $B_{t, \text{policy}}^{t+u}$ terms. We therefore proceed by first specifying the difference in the benefits of period t policies in future periods $t+u$, then taking expectations of this difference, and finally summing up the effects in all periods to find the total benefit difference for policy in period t , $E[B_{t, \text{price}} - B_{t, \text{quantity}}]$.

Using s to sum over the cost shocks in previous periods, we can write:

$$\begin{aligned} \mathbb{E}\left[B_{t, \text{price}}^{t+u} - B_{t, \text{quantity}}^{t+u}\right] &= \mathbb{E}\left[B_{t+u}\left(S_{t+u}^* - \sum_{s=1}^t (1-\delta)^{t+u-s} \frac{\theta_s}{c_s}\right) - B_{t+u}\left(S_{t+u}^* - \sum_{s=1}^{t-1} (1-\delta)^{t+u-s} \frac{\theta_s}{c_s}\right)\right] \\ &= \mathbb{E}\left[-\frac{b_{t+u}}{2}\left(S_t^* - \sum_{s=1}^t (1-\delta)^{t+u-s} \frac{\theta_s}{c_s} - \bar{S}_t\right)^2\right. \\ &\quad \left.-\left(-\frac{b_{t+u}}{2}\left(S_{t+u}^* - \sum_{s=1}^{t-1} (1-\delta)^{t+u-s} \frac{\theta_s}{c_s} - \bar{S}_t\right)^2\right)\right], \end{aligned}$$

where this difference measures the benefits of a price policy in all periods up to and including t , versus the benefits of a quantity policy in period t coupled with price policies in all periods up to $t-1$. This measures the incremental effect of a prices in the last period t , and thus avoids double counting when we later compute the accumulated effect of price policies in every period up to t .

Taking expectations over our expression for the effect of period t policy in period $t+u$, many terms cancel or drop out, yielding

$$\mathbb{E}\left[B_{t, \text{price}}^{t+u} - B_{t, \text{quantity}}^{t+u}\right] = -\frac{b_{t+u}}{2}(1-\delta)^{2u} \left(\text{var}\left(-\frac{\theta_t}{c_t}\right) + 2(1-\delta) \text{cov}\left(-\frac{\theta_t}{c_t}, S_{t-1}\right) \right).$$

This can be written more simply as

$$(14) \quad \mathbb{E}\left[B_{t, \text{price}}^{t+u} - B_{t, \text{quantity}}^{t+u}\right] = -\frac{\sigma_t^2}{2c_t^2} b_{t+u} (1-\delta)^{2u} \Omega_{\rho, t}$$

by replacing the covariance term by a correlation factor $\Omega_{\rho, t}$, where

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

Based on Equation (11) for S_{t-1} , we see that $\Omega_{\rho, t}$ will be non-negative as long as correlation (i.e., the covariance term) is positive. Further, $\Omega_{\rho, t}$ disappears ($\Omega_{\rho, t} = 1$) if correlation is zero or if decay is immediate ($\delta = 1$). More generally, greater correlation or slower depreciation will increase the covariance term, thereby increasing the magnitude of $\Omega_{\rho, t}$ (as shown in part A of the appendix).

In order to measure the entire stream of expected benefit differences associated with a price policy in period t , we add up the discounted effects in all future periods $t+u$,

$$\begin{aligned} (15) \quad \mathbb{E}\left[B_{t, \text{price}} - B_{t, \text{quantity}}\right] &= \sum_{u=0}^{\infty} (1+r)^{-u} \mathbb{E}\left[B_{t, \text{price}}^{t+u} - B_{t, \text{quantity}}^{t+u}\right] \\ &= -\frac{\sigma_t^2}{2c_t^2} \left(\sum_{u=0}^{\infty} b_{t+u} \frac{(1-\delta)^{2u}}{(1+r)^u} \right) \Omega_{\rho, t} \\ &= -\frac{\sigma_t^2}{2c_t^2} b_t \Omega_{\delta} \Omega_{\rho, t} \end{aligned}$$

where the persistence factor Ω_δ captures future benefit losses; Ω_δ is given by

$$(16) \quad \Omega_\delta = \frac{1+r}{1+r - (1+g_b)(1-\delta)^2},$$

assuming that the benefit parameter b_t grows at the constant rate g_b . The loss generated in any particular period depreciates in the future and, since the loss depends on the squared deviation, the factor $(1-\delta)^2$ appears alongside the discount rate in the expression. Lower decay, lower discount rates, and higher rates of benefit growth will increase the magnitude of Ω_δ . Note that $\Omega_\delta = 1$ and drops out for the case of a flow externality that decays immediately (i.e., $\delta = 1$). If $\delta = 0$ and $g_b = 0$ —corresponding to a pure stock externality with no benefits growth—then $\Omega_\delta = (1+r)/r$, the present value factor for an infinite series.

Equation (15) is unambiguously negative, indicating that the price policy leads to lower benefits, thereby creating a policy tradeoff between saved costs and lost benefits. Variation in q_t leads to lower expected benefits from the price policy because benefits are concave (the quantity policy, in contrast, fixes q_t). The factor Ω_δ augments b_t to take account of the persistence of benefits over time, including adjustment for decay and discounting. 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 in Equation (15). With positive correlation, deviations tend to build on each other, rather than canceling out.

C. The Relative Advantage of Prices

By subtracting the cost savings (Equation (13)) from the benefit losses (Equation (15)), we find the relative advantage of using prices rather than quantities in period t :

$$(17) \quad \Delta_t = \frac{\sigma_t^2}{2c_t^2} (c_t - b_t \Omega_\delta \Omega_{\rho,t}).$$

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.

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.

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 Δ , 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 Equation (17) depends on the magnitude of c_t relative to b_t , *multiplied by the persistence and correlation terms*. 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. As described above, the persistence term Ω_δ captures the effect of benefit growth, stock decay, and discounting on marginal benefits, while the correlation term $\Omega_{p,t}$ adjusts for the presence of cost correlation. Faster benefit growth, lower decay rates, lower discount rates, or greater cost correlation will increase the magnitude of $\Omega_\delta \Omega_{p,t}$, thereby increasing the possibility that quantities are preferred. 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. Depending on the rates of time discounting and benefit growth, this “persistence effect” has 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) it is always possible that $c_t < b_t \Omega_\delta \Omega_{p,t}$, 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. 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 again lowers expected benefits under a price policy.

Note that our use of the words “tend to favor” 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, but rather that movements of the parameters in the stated direction will *widen the range* of conditions in which quantities are preferred.

D. Instrument Choice Over Time

It is apparent that the relative advantage expression (Equation (17)) 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 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 ($g_b > 0$). This implies an eventual future preference for quantity controls based on Equation (17). 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 Δ_t represents the advantage of prices relative to quantities in a *single* 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 Δ^T , where

$$(18) \quad \Delta^T = \sum_{t=1}^T (1+r)^{-t} \Delta_t .$$

With policies fixed over longer timeframes and Δ_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 5-year versus a 40-year time horizon due to changes in benefits and costs. The likely direction of these effects—which we noted tend 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 Equation (15) to 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 b_{t+u} , 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 (from Equation (15)) and lowering Δ_t (from Equation (17)). However, this effect will be diminished if it occurs in the future. Equation (15) indicates that the benefit losses in some future period $t+u$ from using prices in the current period t depends not only on the increased severity of damages in the future (b_{t+u}), but also on the amount of time that will pass before the threshold is reached (u periods).

This distance in time matters for two reasons: decay and discounting. 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 similar strategy. That is, provided the thresholds remain significantly beyond our initial planning horizon T in Equation (18). In contrast, a more imminent threshold would tend to demand quantity controls now.

E. 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. The rationale and data sources for these values are described in detail in part B of the appendix. In addition to the parameters discussed so far, we introduce g_q and g_c to represent constant geometric growth rates for \bar{q}_t and c_t , respectively.

TABLE 1—INFORMATION FOR ANALYSIS OF CLIMATE CHANGE POLICY

Parameter	Annualized value
Decay rate of stock (δ)	0.83%
Discount rate (r)	5.0%
Marginal benefit slope (b_0)	8.7×10^{-13} \$/ton ²
Marginal cost slope (c_0)	1.6×10^{-7} \$/ton ²
Cost uncertainty (σ_0)	13 \$/ton
Cost correlation (ρ)	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)	1.7×10^{11} tons
Initial emissions (\bar{q}_0)	5.0×10^9 tons

Note: \$ refers to 1998 US dollars and tons refers to metric tons of carbon. See part B of the appendix for detail on data sources for parameter values.

Using Equations (17) and (18) it is straightforward to compute Δ^T for various policy horizons. With a single period horizon, 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. 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 by solving Equation (5). 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 Nordhaus (1994a)) and total net discounted benefits of \$225 million. The price policy therefore offers a welfare improvement equal to 2.3 times the value of a comparable quantity policy in the first period. In other words, prices generate over three times the expected net benefit of quantities in the first period.

Table 2 also shows results for alternative horizons of 5, 10, 20, and 40 years. The results indicate that over longer horizons, the price advantage rises for climate change policy. With a forty-year horizon, price controls generate \$35 billion in higher expected benefits compared to quantity controls, although this is now “only” an 120 percent improvement over quantities. In general, we find that prices generate nearly 5 times the expected welfare gains of quantities, with consequences on the order of many billions of dollars per year.

TABLE 2—RELATIVE ADVANTAGE OF PRICES OVER QUANTITIES FOR OPTIMAL CLIMATE CHANGE POLICY

Policy horizon	Expected price advantage		Benefits required for indifference	
	\$ billions	Relative to quantities	\$/ton ²	Relative to benchmark
1 year	0.52	2.3	6.1×10^{-9}	7,300
5 years	4.6	3.5	1.7×10^{-9}	2,000
10 years	11	3.8	9.8×10^{-10}	1,200
20 years	21	2.7	5.7×10^{-10}	680
40 years	35	1.2	3.0×10^{-10}	360

Note: \$ refers to 1998 US dollars.

Given that there is a wide range of options 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 horizon, the critical relative slope is obtained by rearranging Equation (17) when $\Delta_t = 0$,

obtaining $(b_0/c_0)_{\text{crit}} = 1/(\Omega_\delta \Omega_{\rho,t})$. The benefit 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 7,300-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 horizon T , we can find solve $\Delta^T = 0$ from Equation (18), obtaining the functional relationship

$$(19) \quad (b_0/c_0)_{\text{crit}} = (b_0/c_0)_{\text{crit}}(T, \rho, g_b, g_c, r, \delta),$$

where the parameter σ_0^2 vanishes since it only scales Δ^T . These values must be computed numerically. From Table 2, we can see that the benefit condition for indifference is relaxed over longer time horizons, 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 forty-year horizon. 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.

IV. Instrument Choice When Benefits Are Not Directly Known

Thus far we have demonstrated that Weitzman's main finding—that the choice between prices and quantities depends on the relative slopes of marginal costs and benefits—requires some important adjustments for growth, decay, discounting, and cost correlation when applied to the case of stock externalities. We have shown how one can determine both *whether* a price or quantity instrument should be preferred for a particular case and *by how much* that instrument is preferred. That is, Equation (17) gives us both the sign and the magnitude of the relative advantage. Even given our simplified model, however, the information requirements for this result may be excessive for many real-world problems. As suggested above, the slope of the benefit function may be especially hard to pin down given that it represents the value of an externality, which by definition is not correctly revealed in market prices. Nonetheless, in this section we demonstrate that even in the absence of direct benefit measures, we can still say something about efficient instrument choice.

Namely, we establish a link between the efficient choice of policy instrument and the optimal production level of the stock externality, demonstrating that certain combinations of instrument and stringency are inconsistent with optimal behavior regardless of the true benefit level. In the previous section, we described and calculated the relative slope of marginal benefits required for indifference between price and quantity policies for climate mitigation, noting that lower benefits favor prices and higher benefits favor quantities. We now take this one step further to consider the optimal stock path associated with this critical relative slope and show that lower mitigation benefits indicate not only a

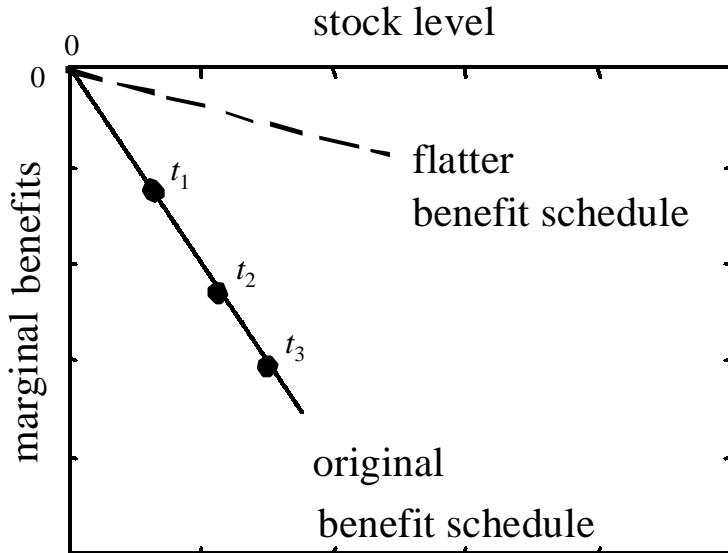
preference for prices, but also a higher optimal stock level in all periods (and lower initial abatement levels). We thereby establish a link between policy instrument and policy stringency. Stock externalities that require only gradual action will favor the use of price-based policies. In contrast, cases that require stringent near-term response will favor the use of quantity policies. Thus, there is typically an inconsistency inherent in stringent price policies or gradual quantity policies for regulating stock externalities. We heuristically outline our supporting argument below, leaving a more detailed proof for the appendix.

Applied to the case of global climate change policy, our parameter estimates (from Table 1) indicate that quantity controls are unjustified without the *complete near-term elimination* of carbon dioxide emissions. Our sensitivity analysis demonstrates that under more generous assumptions favoring quantity controls, price policies continue to be preferred unless mitigation benefits are steep enough to recommend abatement exceeding 40 percent of uncontrolled carbon emissions. These results are appropriate for a generic negative stock externality unless the initial stock level is small or there is a large difference in the growth rates of marginal benefit and marginal cost slopes.

A. Establishing the Link Between Policy Instrument and Stringency

We begin by considering how a change in the benefit parameter b_0 will affect the optimal stock path. Ignoring growth (so $b_t = b_0$), Figure 1 shows how the marginal benefits in a single period, $-b_0(S_t - \bar{S}_t)$, would be associated with several points along the optimal stock path. Note that for the case of a negative externality, movements down the vertical axis indicate negative values of increasing magnitude. In order to satisfy the optimality condition (Equation (8)), the present discounted and decayed value of these marginal benefits must equal marginal costs in period t .

FIGURE 1—FLATTER BENEFITS IMPLY HIGHER OUTPUT (LESS REDUCTION) AND HIGHER STOCKS



Under conditions with flatter marginal benefits (due to smaller b_0), indicated by the dashed line, the magnitude of marginal benefits is lower at each of the indicated stock levels. In order to maintain the optimality condition (Equation (8)), the stock level and the associated marginal benefit level must change. In particular, the stock level will need to rise in each period to partially compensate for the fall in b_0 and return marginal benefits toward their earlier levels. At least initially, this will involve an increase in the output level, which will lower the magnitude of marginal costs and again help rebalance the optimality condition. Flatter marginal benefits thus represent a weakening of the externality, leading to less aggressive action. As we show in part C of the appendix, this behavior is formalized by the following comparative dynamic result:

For a negative stock externality, a lower slope of marginal benefits relative to marginal costs implies a higher stock level in every period and a higher output level in the first period. This establishes a link between optimal policy stringency and instrument choice since a lower slope ratio also makes price policies more desirable. An analogous result holds for positive stock externalities.

This result is useful because it establishes a link between the choice of optimal policy *level* and *instrument*. In the previous section, we computed a “critical slope ratio” associated with indifference between price and quantity controls, below which prices are preferred. We can now use this critical slope ratio to compute an associated level of policy stringency, a “critical output level,” indicating optimal output in the first period based on the critical slope ratio and the remaining parameters. Relative benefit slopes below the critical slope ratio are not only consistent with price controls, but imply output levels

above the critical output level for a negative externality (or lower output levels for a positive externality). In other words, despite a lack of agreement about the benefits associated with a particular externality, an unambiguous consistency must remain between a chosen output level and the instrument used to obtain that output level *if the policy is to be optimal given the remaining parameters*.

B. Calculating Critical Output Levels Above which Prices are Preferred

To compute this critical output level, we use the critical relative slope (Equation (19)), along with the remaining parameters, to solve the dynamic optimization given by Equation (5). Although the general problem does not have an analytic solution, we can write the functional relationship for the optimal first-period output level q_l^* relative to the baseline \bar{q}_l

$$(20) \quad q_l^*/\bar{q}_l = q_l^*/\bar{q}_l \left(S_0/\bar{q}_l, b_0/c_0, g_b, g_c, g_q, r, \delta \right),$$

where we have simplified for the case of a negative externality ($\bar{S}_t = 0$) and expressed benefits and costs, as well as the initial stock and flow, in relative terms. The use of relative measures follows from: (i) our ability to re-scale benefits and costs by the cost parameter c_0 without affecting the solution to Equation (5); and (ii) the recognition that both flow and stock variables in the optimality condition (Equation (8)) and the accumulation equation (3) can be scaled by an arbitrary constant. We scale the stock and flow by $1/\bar{q}_l$, which allows us to focus on the optimal emission level of the negative externality as well as the initial stock level, both *relative* to uncontrolled emissions. The direction of the effect of the remaining parameters on q_l^*/\bar{q}_l are shown above Equation (20) based on derivations given in part D of the appendix. Substituting, Equation (19) into Equation (20), we have

$$(21) \quad \begin{aligned} \left(q_l^*/\bar{q}_l \right)_{\text{crit}} &= q_l^*/\bar{q}_l \left(S_0/\bar{q}_l, (b_0/c_0)_{\text{crit}} (T, \rho, g_b, g_c, r, \delta), g_b, g_c, g_q, r, \delta \right) \\ &= \left(q_l^*/\bar{q}_l \right)_{\text{crit}} \left(S_0/\bar{q}_l, T, \rho, g_b, g_c, g_q, r, \delta \right) \end{aligned}$$

where the signed effects on $(b_0/c_0)_{\text{crit}}$ are given in part D of the appendix.

Equation (21) allows us to compute a critical output level $\left(q_l^*/\bar{q}_l \right)_{\text{crit}}$ for a given set of parameters $\{S_0/\bar{q}_l, T, \rho, g_b, g_c, g_q, r, \delta\}$. Relative slopes below the critical slope ratio indicate both a preference for prices as well as an optimal output level higher than this $\left(q_l^*/\bar{q}_l \right)_{\text{crit}}$. Similarly, relative slope values above the critical slope ratio indicate a preference for quantity controls and an optimal output

level below $(q_1^* / \bar{q}_1)_{\text{crit}}$. In other words, the critical output level defines a threshold for optimal policy stringency consistent with use of quantity controls. In the case of a negative externality, if the optimal level of output is greater than the threshold, it reveals marginal benefits flat enough to generate a positive comparative advantage, implying price controls are preferred.

C. Application to Climate Change Policy

In the case of climate change, we can use the benchmark parameter estimates from Table 1 to compute both $(b_0 / c_0)_{\text{crit}}$ and $(q_1^* / \bar{q}_1)_{\text{crit}}$. We focus on a horizon of 20 years since it is likely that a climate change policy chosen today could be revised within that time if circumstances warrant. Based on Equation (21), and using our benchmark parameter estimates for climate change, we find $(q_1^* / \bar{q}_1)_{\text{crit}} = 0$. That is, unless we consider policies to *completely eliminate* near-term carbon dioxide emissions, our revealed beliefs about costs and benefits are inconsistent with the use of quantity controls. Based on the signed effect of T in Equation (21), a horizon of less than 20 years would only lower the critical output level and strengthen this result. The next section shows that while the quantitative result $(q_1^* / \bar{q}_1)_{\text{crit}} = 0$ is sensitive to certain assumptions, particularly concerning the natural stock level and benefit/cost growth, the qualitative result that less aggressive abatement policies go hand-in-hand with price controls is remarkably robust.

D. Sensitivity Analysis and the General Case of a Negative Stock Externality

A lack of direct knowledge and disagreement about mitigation benefits and the parameter b_0 motivated our focus on the link between the optimal control rate and the choice of instrument. Of course, other parameters describing future growth rates, cost correlation, the discount rate, and even the decay rate may not be precisely known or agreed upon. It turns out that our results about the critical output level can be generalized even when many of these remaining parameters are uncertain. In particular, we can characterize *conservative* critical output levels using a small number of parameters. By conservative, we mean that the critical output level will be as high as possible for a negative externality, shrinking the region where optimal output exceeds the critical level and where prices are revealed to be preferred. We consider this conservative because despite our attempts to shrink this region, price controls continue to be preferred in cases that would typically be identified as “stock” externalities.

As described further in part B of the appendix and summarized in Table 3, we begin this characterization by establishing reasonable value ranges for the remaining parameters. We then either fix each parameter at the most conservative value based on the derivative signs established in Equation (21)

or consider a grid of discrete values for the parameter. We exclude certain combinations of parameters that represent unrealistic cases. The result is 29,196 parameter combinations or “scenarios” for which we first solve for the critical ratio $(b_0/c_0)_{\text{crit}}$ associated with indifference between prices and quantities (as in Table 2), then use this critical ratio to compute an optimal control path over a 200-year horizon.

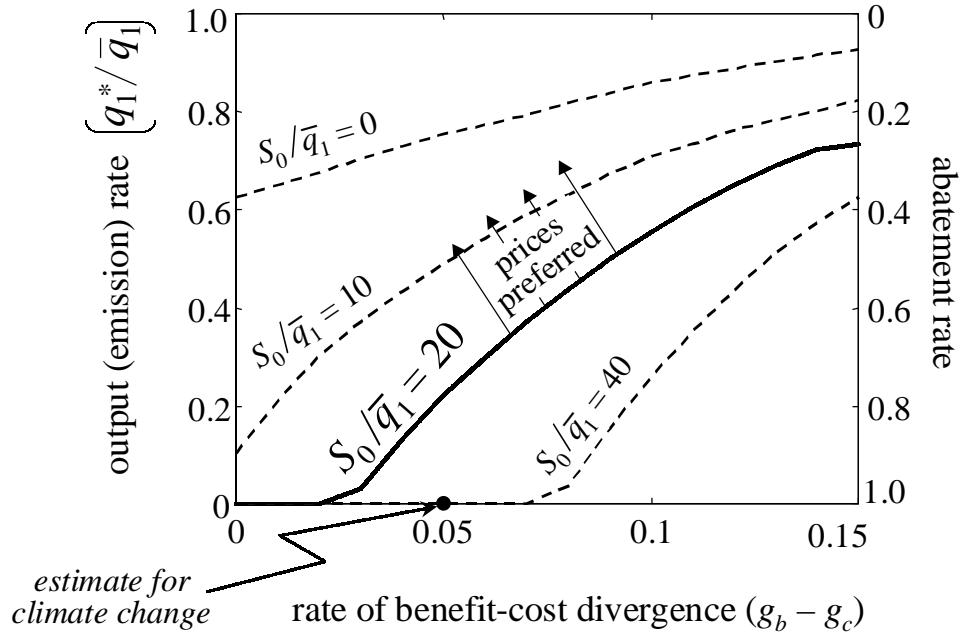
TABLE 3—PARAMETER VALUES FOR COMPUTING CRITICAL OUTPUT LEVELS ABOVE WHICH PRICES ARE PREFERRED

Parameter	Value(s)
Initial stock level (S_0/\bar{q}_1)	0, 10, 20, 40
Time horizon (T)	20 years
Baseline output growth rate (g_q)	0, 1, 2, 3, 4, 5%
Cost growth rate ($g_c + g_q$)	-5, -4, -3, -2, -1, 0%
Benefit growth rate (g_b)	0, 1, 2, 3, 4, 5%
Discount rate (r)	2, 3, 4, 5, 6, 7%
Decay rate of stock (δ)	0-100%
Correlation of cost shocks (ρ)	1

Note: See part B of the appendix for detail on the rationale for parameter value ranges and restrictions imposed on parameter combinations.

Figure 2 shows the resulting critical output levels in the first period, $(q_1^*/\bar{q}_1)_{\text{crit}}$, plotted against the rate of divergence between benefit and cost growth $g_b - g_c$. Downward movements along the vertical axis represent lower output levels and therefore increasingly stringent controls (higher abatement). The figure displays curves representing the upper envelope of the critical output estimates, revealing the value of optimal output levels above which only prices remain efficient for all parameter combinations. Thus, output levels above the curves are consistent only with the use of price policies, while output levels below the curves may or may not be consistent with the use of prices given our conservative approach. We indicate four different loci of critical output levels differentiated by the initial stock/flow ratios S_0/\bar{q}_1 with $S_0/\bar{q}_1 = 20$ highlighted by the use of a solid rather than dashed line. We emphasize the case where $S_0/\bar{q}_1 = 20$ because the idea of a large accumulated stock relative to small annual flow is, it would seem, the most natural characterization of a stock externality—whether it is climate change, hazardous waste, national defense, or knowledge. For example, $S_0/\bar{q}_0 \approx 34$ for atmospheric carbon (see Table 1). The figure also shows the specific point corresponding to the climate change parameters given in Table 1.

FIGURE 2—CRITICAL OUTPUT LEVELS ABOVE WHICH PRICES ARE PREFERRED



Note: For a given stock/flow ratio (S_0/\bar{q}_1), areas above each curve designate control levels that are consistent with the use of prices and are inconsistent with the use of quantities.

Based on the benchmark climate change parameters given in Table 1, the figure illustrates that only policies involving the *complete near-term elimination* of emissions are consistent with the use of quantity instruments. More generally, if the initial stock/flow ratio is 20 or greater and $g_b - g_c$ is below 10 percent, the figure tells us that only policies involving 40 percent or greater rates of abatement are potentially consistent with the use of quantity policies. Since a large accumulated stock is the typical characterization of a stock externality and benefit/cost growth rates are limited in the long run by economic growth, this is not only a robust result for climate change, but a general result for most stock externalities. This result is limited only by cases of high growth or low initial stock levels.

This result clarifies the intuition indirectly suggested by earlier authors such as Nordhaus (1994), Kolstad (1996), Pizer (1997), and McKibbin and Wilcoxen (1998). Intuitively, that work suggests that as long as the existing stock is large relative to the annual flow, marginal benefits will tend to look very flat over the range of annual emissions since the reductions that could be taken in a given year will never be enough to significantly alter the stock. Based on Weitzman (1974), this generic characteristic of what it means to be a stock externality weighs heavily in favor of price instruments for their control. Figure 2 shows that this is generally true for carbon mitigation and other pollution stocks unless: (i) marginal benefits are high enough to warrant 100 percent abatement, suggesting that benefits are indeed steep; or

(ii) the benefit/cost slope ratio is rapidly diverging, so that a relatively low and flat marginal benefit schedule today may become a high and steep schedule in the near future.

Turning to the figure in more detail, there is an intuitive relationship between the critical output level and the two parameters shown in Figure 2, S_0/\bar{q}_1 and $g_b - g_c$. For any parameterization, a low initial stock level S_0/\bar{q}_1 provides some ability to delay action until the stock accumulates to significant levels. When S_0/\bar{q}_1 is large, this flexibility does not exist. Therefore, high initial stock levels imply that optimal policy will move more quickly towards stock stabilization and higher abatement levels.

The rate of divergence between benefits and costs $g_b - g_c$ has a very different effect. As shown in part D of the appendix, we know that both faster benefit growth and slower cost growth will lower the critical benefit slope, tending to favor quantities. Intuitively, increasing values of $g_b - g_c$ distinguish cases where, over twenty years, we are indifferent between prices and quantity controls in *each* year versus cases where we are indifferent *overall*, even though prices are better in the initial years and quantity controls are better towards the end. In the latter case, with higher benefit growth and lower initial benefits, the optimal initial abatement will be lower (and emissions higher) in advance of the future rise in mitigation benefits and fall in abatement costs. Hence, we see a positive relation between benefit-cost divergence on the horizontal axis and optimal initial emissions on the vertical axis.

V. 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 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.

Beyond characterizing the relative advantage of price controls for a stock externality, we further extend the Weitzman results by linking the choice of instrument and policy stringency. This result is especially useful when there is disagreement about benefits, since it can be used to restrict the set of efficient policies even in such cases. In particular, we demonstrate that quantity-based instruments can only be associated with more aggressive controls, that is, high abatement levels for a stock pollutant. For

this kind of negative stock externality, we characterize a critical output level above which the optimally chosen output level reveals beliefs are consistent with price controls and inconsistent with quantity controls. We show that only when the existing stock level is relatively small and the rate of future divergence between benefits and costs is quite large, can quantity controls be consistent with levels of abatement below 40 percent. Otherwise, price controls are preferred. While our choice of parameter values is motivated by the example of climate change and carbon dioxide emission reductions, we believe these results are applicable to a wide range of market failures involving negative stock externalities.

This confirms earlier intuition with two caveats: As long as the existing stock is large relative to the annual flow, it has been argued that marginal benefits will tend to look very flat over the range of annual emissions, since the reductions that could be taken in a given year will never be enough to significantly alter the stock. Based on Weitzman's relative slope argument, this generic characteristic of what it means to be a stock externality weighs heavily in favor of price instruments for their control. Our results demonstrate that this is true unless marginal benefits are high enough to warrant high abatement levels in the immediate future, or if benefits grow rapidly relative to costs.

Regarding climate change, these results have important implications for the recently negotiated Kyoto Protocol. First, application of our relative advantage expression finds that price instruments for carbon reduction generate up to five times the expected welfare gains of quantity instruments, depending on the policy horizon. Yet, the Kyoto Protocol requires binding, quantity-based reductions. Second, the Kyoto Protocol mandates average reductions of 4 percent below 1990 levels for industrialized countries, or about a 10 percent reduction below 2010 projected baseline levels (Energy Information Administration 1998). This represents only a 5½ percent decline in global flows into the atmospheric stock of carbon since Kyoto does not limit emissions from non-industrialized countries—which will soon comprise about half of global emissions. The Kyoto mandates are therefore well below even our conservative estimate of the reductions necessary (40%) to justify the use of quantity policies in efficiency terms. Kyoto thus embodies an economic inconsistency between targets and instruments.

What does this tell us? If we believe the Kyoto Protocol reveals a balance between costs and benefits, it tells us that the Kyoto goals should be treated as “targets” to be achieved through policies that include flexible price-based elements rather than rigid quantity mandates. On the other hand, if we believe that the threat of climate change demands rigid quantity controls, the limited reductions associated with Kyoto are wholly inadequate. One might argue that the Kyoto Protocol should be viewed as a delicate balance among many competing factors, some economic and some geo-political. Perhaps so. But at a minimum, our analysis suggests that the cost of this economic sacrifice is on the order of many billions of dollars and should not be made lightly.

Appendix

A. Details on Correlation Factor

We introduce the correlation factor $\Omega_{\rho,t}$ to capture the potential correlation between current and previous period shocks to the stock level, which will increase benefit losses under a price policy. Here we derive an analytic expression for $\Omega_{\rho,t}$ and establish the sign of various derivatives. Our expression for $\Omega_{\rho,t}$ follows from two earlier results, namely

$$S_t = E[S_t] - \frac{\theta_t}{c_t} - \frac{(1-\delta)\theta_{t-1}}{c_{t-1}} - \dots - \frac{(1-\delta)^{t-1}\theta_1}{c_1}$$

from Equation (11), and

$$E[\theta_t \theta_s] = \sigma_{\min(s,t)}^2 \rho^{|s-t|} = \sigma_0^2 \rho^{|s-t|} (1 + \rho^2 + \dots + \rho^{2 \cdot \min(s,t)-2}) = \sigma_0^2 \rho^{|s-t|} \frac{1 - \rho^{2 \cdot \min(s,t)}}{1 - \rho^2},$$

which follows from the autoregressive evolution of θ_t : $\theta_t = \rho \theta_{t-1} + \varepsilon_t$ and $\theta_{t < 1} = 0$. With these two results, it is straightforward to compute

$$\begin{aligned} \Omega_{\rho,t} &= 1 + \frac{2c_t^2}{\sigma_t^2} \text{cov}\left(-\frac{\theta_t}{c_t}, (1-\delta)S_{t-1}\right) \\ (22) \quad &= 1 + 2 \left(\frac{c_t}{c_{t-1}} \frac{\rho(1-\delta)\sigma_{t-1}^2}{\sigma_t^2} + \dots + \frac{c_t}{c_1} \frac{\rho^{t-1}(1-\delta)^{t-1}\sigma_1^2}{\sigma_t^2} \right) \\ &= 1 + 2 \left(\rho(1-\delta)(1+g_c) \frac{(1-\rho^{2t-2})}{(1-\rho^{2t})} + \dots + \rho^{t-1}(1-\delta)^{t-1}(1+g_c)^{t-1} \frac{(1-\rho^2)}{(1-\rho^{2t})} \right) \end{aligned}.$$

Equation (22) can be summed to create the expression

$$\Omega_{\rho,t} = 1 + \frac{2\rho(1-\delta)(1+g_c)}{1-\rho^{2t}} \left(\frac{1 - (\rho(1-\delta)(1+g_c))^{t-1}}{1 - \rho(1-\delta)(1+g_c)} - \rho^{2t-2} \frac{1 - ((1-\delta)(1+g_c)/\rho)^{t-1}}{1 - (1-\delta)(1+g_c)/\rho} \right).$$

We can establish the sign of several derivatives of $\Omega_{\rho,t}$ by considering an arbitrary term in Equation (22):

$$\rho^s (1-\delta)^s (1+g_c)^s \frac{(1-\rho^{2t-2s})}{(1-\rho^{2t})}.$$

From this term we immediately have $\frac{\partial \Omega_{\rho,t}}{\partial \delta} < 0$ and $\frac{\partial \Omega_{\rho,t}}{\partial g_c} > 0$. Taking the derivative of $\Omega_{\rho,t}$ with respect to ρ and rearranging we have

$$\begin{aligned} \frac{\partial}{\partial \rho} \left[\rho^s \frac{(1-\rho^{2t-2s})}{(1-\rho^{2t})} \right] &= \frac{\rho^{-s-1} (s\rho^{2s} + s\rho^{2t} - s\rho^{4t} - s\rho^{2t+2s} - 2t\rho^{2t} + 2t\rho^{2t+2s})}{(1-\rho^{2t})^2} \\ &= t\rho^{-s-1} \frac{[xy^x + xy - xy^2 - xy^{1+x} - 2y + 2y^{1+x}]}{(1-\rho^{2t})^2} \end{aligned},$$

where $x = s/t$ and $y = \rho^{2t}$ (so $0 < x, y < 1$). Rearranging the bracketed term, we have

$$xy^x + xy - xy^2 - xy^{1+x} - 2y + 2y^{1+x} = y(1+y^{x-1}) \left[x(1-y) - 2 \left(\frac{1-y^x}{1+y^{x-1}} \right) \right] \geq 0.$$

The inequality follows from: (i) the fact that $2(1-y^0)/(1+y^{-1}) = 0$ and $2(1-y^1)/(1+y^0) = 1-y$, so the two terms in brackets are equal when $x=0$ and $x=1$ for any $1 > y > 0$; and (ii) the second term is convex:

$$\frac{\partial^2}{\partial x^2} \left[\frac{1-y^x}{1+y^{x-1}} \right] = \frac{y^{1+2x}(1+y)(1-y^{1-x})(\log y)^2}{(y+y^x)^3} > 0,$$

versus the first term which is linear. Thus, $\frac{\partial \Omega_{\rho,t}}{\partial \rho} > 0$.

Finally, consider the derivative of $\Omega_{\rho,t}$ with respect to time t ,

$$\frac{\partial}{\partial t} \left[\frac{1-\rho^{2t-2s}}{1-\rho^{2t}} \right] = \frac{2\rho^{2t-2s}(1-\rho^{2s})(-\log \rho)}{(1-\rho^{2t})^2} > 0.$$

Since every term in Equation (22) is increasing in t and raising t one period will also add another positive term, $\frac{\partial \Omega_{\rho,t}}{\partial t} > 0$.

B. Data for Climate Policy Application and Sensitivity Analysis

In this section we describe the rationale and sources for the values we used for our application to climate policy and general sensitivity analyses, as presented in Tables 2 and 3. All monetary values are denominated in 1998 US dollars, if necessary adjusted using the price index for gross domestic product (Council of Economic Advisors 1999).

Uncontrolled Output (\bar{q}_0), *Output Growth* (g_q) and *Initial Stock Level* (S_0). The pre-industrial concentration of carbon in the atmosphere is 6.13×10^{11} tons from Neftel et al (1999), converted from 288 ppm to metric tons using a conversion rate of 2.13×10^9 tons/ppm. The

concentration of carbon in 1998 is 367 ppm or 7.81×10^{11} tons from Keeling and Whorf (1999). Since we measure the stock relative to its benefit-maximizing level, our initial stock level ($S_0 = 1.7 \times 10^{11}$ tons) is found by taking the difference between these two values. Initial (1998) global carbon emissions ($\bar{q}_0 = 5.0 \times 10^9$ tons) are based on 6.52×10^9 tons of emissions from fossil-fuel burning, cement manufacture, and gas flaring in 1996 (Marland et al 1999), plus 1.25×10^9 tons of emissions from land use changes (e.g., deforestation) (Kattenberg et al 1996). Total carbon emissions of 7.77×10^9 tons are then adjusted by multiplying by 0.64 to account for short-run decay (Nordhaus 1994a). This yields an initial stock/output ratio (S_0/\bar{q}_0) of 34. For sensitivity analysis, we consider initial stock/output ratios between zero and forty.

The growth rate of baseline carbon emissions ($g_q = 1.5\%$) is based on our estimate of the rate of emissions growth over the 25 years 1972–1996 using the data and model described below in the paragraph on cost uncertainty. It is equal to the average 1.5% rate used in IPCC scenario IS92a over the period 2000–2020 (Kattenberg et al 1996). For sensitivity analysis we consider rates of output growth ranging from 0% to 5%. This range was determined by the assumption that output growth is limited in the long run by economic growth, which seems unlikely to exceed 5%.

Initial Benefits (b_0), Costs (c_0), and Growth (g_b , g_c). The initial value for the marginal benefit slope ($b_0 = 8.7 \times 10^{-13}$ \$/ton²) is based on Nordhaus' (1994b) survey of climate experts as reported in Roughgarden and Schneider (1999). The survey asked respondents to estimate the loss associated with 3°C warming, producing answers that ranged from zero to a 21% loss of gross world product (GWP). We use the median response of a 1.85% loss in our analysis. We associate 3°C warming with a 1.12×10^{12} ton increase over pre-industrial carbon concentrations, based on the assumptions that a doubling of carbon concentrations leads to a 2°C warming and that temperature change is proportional to the change in the log of the carbon stock (Kattenberg et al 1996). In order to convert the fractional loss of GWP to dollars, we multiply by the \$29.5 trillion IMF estimate of 1998 GWP (International Monetary Fund 1998). This gives us \$546 billion in damages from 1.12×10^{12} tons of carbon which, applied to the benefit relation (Equation (1)), generates our estimate of b_0 . This estimate leads to an initial marginal benefit of abatement of roughly \$9/ton (discounted and depreciated), which is comparable to earlier work by Falk and Mendelsohn (1993) and Nordhaus (1994a).

The initial value for the marginal cost slope ($c_0=1.6 \times 10^{-7}$ \$/ton²) is based on results from 10 models that participated in the Energy Modeling Forum's EMF 16, as reported by Weyant and Hill (1999). For each model, we considered a 10% emission reduction. We divided the estimated marginal cost by 0.64 to account for short-run emissions decay, converting from dollars per ton of *emitted* carbon to dollars per ton of carbon *retained* in the atmosphere. Second, we computed a marginal cost slope for each model by dividing the adjusted marginal costs by 5.0×10^8 tons, which is 10% of our estimate of initial retained emissions (\bar{q}_0). We used the mean value of these slopes as our estimate of c_0 . Although this estimate is several times higher than those reported in Falk and Mendelsohn (1993) and Nordhaus (1994a), it reflects more recent modeling results suggesting less elasticity in carbon demand.

The benefit growth rate ($g_b=2.5\%$) is the annualized growth rate in GWP for 2000–2050 from Leggett et al (1992) for the central IPCC scenario IS92a. This value captures benefit growth due to both population and income per capita. For sensitivity analysis we allow g_b to range from 0% to 5%. As with output, this is based on the assumption that benefit growth is limited by economic growth and that 5% is a reasonable upper bound.

The cost growth rate ($g_q+g_c=-1.0\%$) represents our assumed rate of change in the marginal cost associated with a fixed fractional reduction in emissions. (Note that g_c alone indicates the rate of change in marginal costs associated with the first ton of abatement, regardless of the uncontrolled emission level.) There is limited information in the literature upon which to base this parameter; we have chosen it to equal the average rate of change in marginal costs for the 10 EMF models. As a measure of technological change, we allow g_q+g_c to take on values ranging from 0% to -5% in our sensitivity analysis.

Cost uncertainty (σ_0) and Correlation (ρ). Cost uncertainty ($\sigma_0=13$ \$/ton) is the sum of uncertainty in the cost function itself and uncertainty in baseline emissions. Our estimate of uncertainty in the cost function itself (6 \$/ton) is the standard error of the adjusted marginal cost of a 1.1% emission reduction from the 10 EMF models, as described above. Reductions of 1.1% are optimal in the first period for the deterministic optimal control problem in Equation (5) based on the preceding parameter values. Our estimate of uncertainty in baseline emissions is from a first-order autoregressive maximum-likelihood model of global carbon emissions over the 25 year period 1972–1996 using data from Marland et al (1999). We estimated the model

$$\bar{q}_t = \alpha \exp(g_q t) + \mu_t,$$

where \bar{q}_t is the historic uncontrolled emission level, $\mu_t = \rho\mu_{t-1} + \nu_t$ and ν_t is iid. The estimated standard error of ν_t was 1.1×10^8 tons and the estimate of ρ was 0.66. We converted the emissions uncertainty into cost uncertainty by multiplying by 0.64 (to adjust for short-run decay) and then by our estimated slope of marginal costs; the result was 11 \$/ton. We computed total cost uncertainty (13 \$/ton) by summing the variances of its two underlying components.

If we use the same approach to measure cost uncertainty after twenty periods, when optimal reductions are 2.7%, we find that the 10 EMF models generate a standard error of 15 \$/ton (due to divergence of the models at higher reduction levels) while correlation of uncontrolled emissions shocks generates an error of 15 \$/ton. Summing the variances we find a total cost error of 21 \$/ton after 20 years which, beginning with cost shocks of 13 \$/ton in year one, would occur if the overall correlation were 0.80, our choice for ρ . That is, we choose ρ to match the estimated cost errors in year one and year twenty.

Decay (δ) and Discounting (r). The decay rate of carbon dioxide in the atmosphere ($\delta=0.83\%$) is from Nordhaus (1994a). For our sensitivity analyses, we allow δ to vary between 0 and \bar{q}_0/S_0 , since a decay rate of δ can lead to a stock/flow ratio of at most $1/\delta$. We use a discount rate of $r = 5\%$ based on the central value from Pizer (1997). For our sensitivity analysis, we considered a recent review of discount rates used by policymakers (Bazelon and Smetters 1999) and use the range 2% to 7%.

Approach for Sensitivity Analysis. For our sensitivity analysis, we set $\rho=1$ so that the critical output level will be as high as possible for a negative externality, thereby being biased in favor of quantity controls. We establish reasonable value ranges for the remaining parameters, as described above, and for each parameter we consider six evenly spaced values, except S_0/\bar{q}_1 where we consider four exponentially spaced values (i.e., 0, 10, 20, and 40). We then consider all possible combinations of these parameter values, thereby generating an initial set of $4 \times 6 \times 6 \times 6 \times 6 = 31,104$ different parameter combinations or stock externality “scenarios”.

In addition to establishing reasonable ranges for individual parameters, however, it is also important consider the implications of certain parameter combinations. In particular, we restrict our attention to cases where $r + 2\delta - g_b > 0$ for the following reason. We know from Equation

(16) that $\frac{(1-\delta)^2(1+g_b)}{1+r} \geq 1$ (or, approximately, $r + 2\delta - g_b \leq 0$) implies quantities will be

preferred for any non-zero value of b_0 and therefore $(q_1^*/\bar{q}_1)_{\text{crit}} = 1$. Intuitively, the concavity of benefits is growing at a sufficiently fast rate to cause any price-induced variation in the stock level to have infinitely negative consequences—even allowing for discounting and decay.

We exclude these combinations as representing unusual, unrealistic cases both because (i) they imply that the long-run stock should be zero ($\lim_{t \rightarrow \infty} S_t = 0$) and (ii) they suggest a negative rate of pure time preference. First, $r + 2\delta - g_b \leq 0$ implies $r \leq g_b$ and therefore that the present value benefit associated with any non-zero stock level S in the distant future

$$\lim_{t \rightarrow \infty} -b_0 S^2 \left(\frac{1+g_b}{1+r} \right)^t$$

is minus infinity. This means that the optimal long-run stock level is zero—a very unusual case that renders the problem uninteresting. In the case of climate change, for example, there are frequent discussions of stabilizing atmospheric carbon dioxide concentrations at a level of 450 ppm or higher, but rarely trying to return concentrations to their pre-industrialization (and presumably harmless) level of 288 ppm.

Second, the discount rate can be decomposed into components associated with the pure rate of time preference and with productivity growth (Arrow et al 1996):

$$(23) \quad r = (\text{pure time preference}) + \tau (\text{productivity growth})$$

where τ measures the curvature of a constant relative risk aversion utility function; empirically τ is usually greater than one. To the extent that benefit growth is limited by productivity growth, so $g_b < (\text{productivity growth})$, Equation (23) implies $r > g_b$. This is inconsistent with the possibility that $r + 2\delta - g_b \leq 0$. For both of these reasons, we focus our attention on the 29,196 cases where $r + 2\delta - g_b > 0$. Based on our use of evenly-spaced parameter values, this actually implies $r + 2\delta - g_b > 1\%$ which, from the decomposition in Equation (23), implies a pure rate of time preference of at least 1%.

C. Establishing the Link Between Policy Instrument and Stringency

We construct a proof by contradiction of the comparative dynamic result by showing that if the stock were to *not* decrease in response to an increase in the marginal benefits slope (for a negative externality), this would lead to a violation of the conditions for optimality. Therefore our result holds: that steeper benefits imply lower initial output of a negative stock externality

and increased output of a positive stock externality and the reverse for flatter benefits. We begin with an optimal production and stock path, q_t^* and S_t^* , for a particular path of cost and benefit parameters, c_t and b_t , as well as the cost minimizing and benefit maximizing production and stock levels, \bar{q}_t and \bar{S}_t , respectively. We focus on the case of a bad, where $S_t > \bar{S}_t$ and $q_t > \bar{q}_t$; the proof proceeds analogously with the inequalities reversed for a good.

Consider an alternative set of steeper benefit parameters $\hat{b}_t > b_t$. We begin with the supposition (later shown to be false) that along the new optimal path, \hat{q}_t^* and \hat{S}_t^* , there exists a period \hat{t} where the stock path associated with steeper benefits is *not* decreasing relative to the baseline stock path (i.e., $\hat{S}_{\hat{t}-1}^* \leq S_{\hat{t}-1}^*$ and $\hat{S}_{\hat{t}}^* \geq S_{\hat{t}}^*$). From the accumulation equation (3) we know that $\hat{q}_{\hat{t}}^* \geq q_{\hat{t}}^*$. Rearranging the Euler equation (7), we have

$$(24) \quad c_t (q_t - \bar{q}_t) = -b_t (S_t - \bar{S}_t) + \left(\frac{1-\delta}{1+r} \right) c_{t+1} (q_{t+1} - \bar{q}_{t+1}),$$

which implies that $\hat{q}_{\hat{t}+1}^* > q_{\hat{t}+1}^*$. Recursively applying this technique implies $\hat{q}_t^* > q_t^*$ and $\hat{S}_t^* > S_t^*$ for all $t > \hat{t}$.

Now consider Equation (8) (the recursively substituted Euler equation):

$$(25) \quad c_t (q_t - \bar{q}_t) = - \sum_{s=t}^{\infty} b_s (S_s - \bar{S}_s) \frac{(1-\delta)^{s-t}}{(1+r)^{s-t}}.$$

For a bad, where $S_t > \bar{S}_t$, these results ($\hat{b}_t > b_t$, $\hat{S}_t^* > S_t^*$, and $\hat{q}_t^* > q_t^*$ for all $t > \hat{t}$) imply that after period \hat{t} , the right side of Equation (25) unambiguously decreases and the left side increases. Therefore, our original supposition—that there exists a period \hat{t} when the stock path associated with steeper benefits is *not* decreasing relative to the original stock path—is false. Since $\hat{S}_0^* = S_0^*$, this implies that with steeper benefits it must be the case that the optimal stock level falls for all $t > 0$ ($\hat{S}_t^* < S_t^*$) and $\hat{q}_1^* < q_1^*$ for a negative externality.

D. Signing the Derivatives of $(b_0/c_0)_{\text{crit}}$ and q_1^*/\bar{q}_1

We first compute the derivatives of $(b_0/c_0)_{\text{crit}}$ with respect to different arguments, finding the following:

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

where signs above parameters give the direction of effects. We sign these effects by applying the implicit function theorem to $\Delta^T = \Delta(T, \sigma_0^2, c_0, (b_0/c_0)_{\text{crit}}, \rho, g_b, g_c, r, \delta)$. From Equation (17)

we have $\frac{\partial \Delta_t}{\partial b_t} < 0$ and with $b_t = (b_0/c_0) c_0 \exp(g_b t)$, $\frac{\partial \Delta_t}{\partial (b_0/c_0)} < 0$. From Equation (18)

$\frac{\partial \Delta^T}{\partial (b_0/c_0)} < 0$. This provides the first step. In order to sign each derivative x , we determine the

sign of $\frac{\partial \Delta^T}{\partial x}$ and, applying the implicit function theorem, $\frac{\partial (b_0/c_0)}{\partial x} = -\frac{\partial \Delta^T}{\partial x} / \frac{\partial \Delta^T}{\partial (b_0/c_0)}$, we find

that the sign of $\frac{\partial (b_0/c_0)}{\partial x}$ will be the same as the sign of $\frac{\partial \Delta^T}{\partial x}$.

In several situations, we will appeal to the fact that whenever $\Delta^T = 0$ there always exists a period $\hat{t} < T$ where $\Delta_t \geq 0$ for $t \leq \hat{t}$ and $\Delta_t < 0$ for $t > \hat{t}$. This is true because we assume $g_b \geq 0 \geq g_c$ (see data section) and $\Omega_{\rho, t+1} \geq \Omega_{\rho, t}$ (from our discussion of $\Omega_{\rho, t}$). Therefore, $c_t - b_t \Omega_\delta \Omega_{\rho, t} < c_{t-1} - b_{t-1} \Omega_\delta \Omega_{\rho, t-1}$. Based on the existence of \hat{t} , $\sum_{t \leq \hat{t}} f(t) \Delta_t / (1+r)^t < 0$ for any monotonically increasing non-negative function $f(\cdot)$, since it places greater weight on the values of $\Delta_t < 0$ when $t > \hat{t}$.

We now consider the effect of each of the arguments of $(b_0/c_0)_{\text{crit}}$ in turn.

T . Since $c_t - b_t \Omega_\delta \Omega_{\rho, t} < c_{t-1} - b_{t-1} \Omega_\delta \Omega_{\rho, t-1}$ and $\Delta_T < 0$, it follows that $\Delta_{T+1} < 0$.

Therefore, $\frac{\partial \Delta^T}{\partial T} < 0$ and it follows that $\frac{\partial (b_0/c_0)_{\text{crit}}}{\partial T} < 0$.

ρ . We have $\frac{\partial \Delta^T}{\partial \rho} = \sum_{t=1}^T \frac{\Delta_t}{(1+r)^t} \left(\frac{1}{\sigma_t^2} \frac{\partial \sigma_t^2}{\partial \rho} \right) - \sum_{t=1}^T \frac{1}{(1+r)^t} \frac{\sigma_t^2}{c_t^2} b_t \Omega_\delta \frac{\partial \Omega_{\rho, t}}{\partial \rho}$. We can compute.

$$\begin{aligned} \frac{1}{\sigma_t^2} \frac{\partial \sigma_t^2}{\partial \rho} &= \frac{\sigma_0^2}{\sigma_t^2} \frac{\partial}{\partial \rho} [1 + \rho^2 + \dots + \rho^{2t-2}] \\ &= \frac{2\rho + 4\rho^3 + \dots + (2t-2)\rho^{2t-3}}{1 + \rho^2 + \dots + \rho^{2t-2}} = \frac{2\rho(1 - \rho^{2t-2} - (1 - \rho^2)(t-1)\rho^{2t-2})}{(1 - \rho^2)(1 - \rho^{2t})} > 0 \end{aligned}$$

The derivative of this expression with respect to time is

$$\frac{\partial}{\partial t} \left[\frac{2\rho(1-\rho^{2t-2} - (1-\rho^2)(t-1)\rho^{2t-2})}{(1-\rho^2)(1-\rho^{2t})} \right] = \frac{2\rho^{2t-1}(r^{2t} - 1 - \log \rho^{2t})}{(1-\rho^{2t})^2} > 0,$$

where the inequality follows because $\rho^{2t} - 1 \geq \log \rho^{2t}$ for all $\rho^{2t} > 0$. Since $\left(\frac{1}{\sigma_t^2} \frac{\partial \sigma_t^2}{\partial \rho} \right)$ is positive and monotonically increasing over t , we know the first term is negative. From Equation (22) we know that $\frac{\partial \Omega_{\rho,t}}{\partial \rho} > 0$, so $-\sum_{t=1}^T \frac{1}{(1+r)^t} \frac{\sigma_t^2}{c_t^2} b_t \Omega_\delta \frac{\partial \Omega_{\rho,t}}{\partial \rho} < 0$. Therefore, $\frac{\partial \Delta^T}{\partial \rho} < 0$ and it follows that $\frac{\partial(b_0/c_0)_{\text{crit}}}{\partial \rho} < 0$.

g_c . Rewriting Equation(17):

$$(26) \quad \Delta_t = \frac{\sigma_t^2}{2c_0 \exp(g_c t)} (b_0/c_0)_{\text{crit}} \exp((g_b - g_c)t) \Omega_\delta \Omega_{\rho,t},$$

and we have

$$\begin{aligned} \frac{\partial \Delta_t}{\partial g_c} &= -t \Delta_t + t \frac{\sigma_t^2}{2c_1 \exp(g_c t)} (b_0/c_0)_{\text{crit}} \exp((g_b - g_c)t) \Omega_\delta \Omega_{\rho,t} \\ &\quad + \frac{\sigma_t^2}{2c_1 \exp(g_c t)} (b_0/c_0)_{\text{crit}} \exp((g_b - g_c)t) \Omega_\delta \frac{\partial \Omega_{\rho,t}}{\partial g_c} \end{aligned}$$

and

$$\frac{\partial \Delta^T}{\partial g_c} = -\sum_{t=1}^T \frac{t \Delta_t}{(1+r)^t} + \sum_{t=1}^T \frac{1}{(1+r)^t} (\text{second + third terms}) > 0,$$

where the inequality follows from the fact that the second and third terms are both positive and

that t is a non-negative, monotonically increasing function. Since $\frac{\partial \Delta^T}{\partial g_c} > 0$, it follows that

$$\frac{\partial(b_0/c_0)_{\text{crit}}}{\partial g_c} > 0.$$

g_b . From Equation (26) we have

$$\frac{\partial \Delta_t}{\partial g_b} = -\frac{\sigma_t^2}{2c_1 \exp(g_c t)} (b_0/c_0)_{\text{crit}} \exp((g_b - g_c)t) \Omega_\delta \Omega_{\rho,t} \left(t + \frac{(1-\delta)^2}{1+r - (1+g_b)(1-\delta)^2} \right) < 0,$$

where the second term follows from computing $\partial\Omega_\delta/\partial g_b$ using Equation (16). Therefore,

$$\frac{\partial\Delta^T}{\partial g_b} < 0 \text{ and it follows that } \frac{\partial(b_0/c_0)_{\text{crit}}}{\partial g_b} < 0.$$

r. From Equation (26) we have

$$\frac{\partial\Delta_t}{\partial r} = -\frac{\sigma_t^2}{2c_1 \exp(g_c t)} (b_0/c_0)_{\text{crit}} \exp((g_b - g_c)t) \Omega_{\rho,t} \frac{\partial\Omega_\delta}{\partial r}.$$

From Equation (16) $\frac{\partial\Omega_\delta}{\partial r} < 0$ since Ω_δ accounts for the discounted value of future stock

variation, which declines with higher discount rates. Therefore, $\frac{\partial\Delta_t}{\partial r} > 0$. Now,

$$\frac{\partial\Delta^T}{\partial r} = -\sum_{t=1}^T \frac{t\Delta_t}{(1+r)(1+r)^t} + \sum_{t=1}^T \frac{1}{(1+r)^t} \frac{\partial\Delta_t}{\partial r} > 0, \text{ where the first term is positive because } t/(1+r)$$

is a non-negative monotonically increasing function. Since $\frac{\partial\Delta^T}{\partial r} > 0$ it follows that

$$\frac{\partial(b_0/c_0)_{\text{crit}}}{\partial r} > 0.$$

δ. From Equation (26) we have

$$\frac{\partial\Delta_t}{\partial\delta} = -\frac{\sigma_t^2}{2c_1 \exp(g_c t)} (b_0/c_0)_{\text{crit}} \exp((g_b - g_c)t) \left(\Omega_{\rho,t} \frac{\partial\Omega_\delta}{\partial\delta} + \Omega_\delta \frac{\partial\Omega_{\rho,t}}{\partial\delta} \right).$$

From Equation (16), we know that $\frac{\partial\Omega_\delta}{\partial\delta} < 0$, since a higher decay rate will reduce the future

effects of current period stock variation. Similarly, from Equation (22) $\frac{\partial\Omega_{\rho,t}}{\partial\delta} < 0$ since a higher

decay rate decreases the potential correlation with past shocks. Therefore, $\frac{\partial\Delta_t}{\partial\delta} > 0$, $\frac{\partial\Delta^T}{\partial\delta} > 0$

and it follows that $\frac{\partial(b_0/c_0)_{\text{crit}}}{\partial\delta} > 0$.

We now consider the derivatives of q_1^*/\bar{q}_1 with respect to its arguments, finding the following:

$$q_1^* = \bar{q}_1 \cdot q_1^*/\bar{q}_1 \left(S_0/\bar{q}_1, b_0/c_0, g_b, g_c, g_q, r, \delta \right).$$

We do not have an analytic expression for q_1^* because it is the solution to an optimal control problem with time-varying parameters. However, we can use the Euler equation (24) to establish the sign of several derivatives using proofs by contradiction similar to the one presented in section C.

S_0 . Suppose $\tilde{S}_0 > S_0$ and in response $\tilde{q}_1^* \geq q_1^*$ where tildes indicate an alternative parameterization and solution. The left side of the Euler equation (24) becomes more positive and the right side more negative. In response, $\tilde{q}_2^* > q_2^*$, and from the accumulation equation $\tilde{S}_1 = (1 - \delta) \tilde{S}_0 + \tilde{q}_1^* > S_0$. Applying this argument iteratively, we find $\tilde{q}_t^* > q_t^*$ and $\tilde{S}_t > S_t$ for all t (except $\tilde{q}_1^* \geq q_1^*$). But this violates Equation (25), where the left side is more positive and the right more negative. Therefore $\tilde{q}_1^* < q_1^*$ and $\frac{\partial q_1^*}{\partial S_0} < 0$.

b_0/c_0 . Section C of the Appendix demonstrates that $\frac{\partial q_1^*}{\partial (b_0/c_0)} < 0$.

g_b . Suppose $\tilde{g}_b > g_b$ and in response $\tilde{q}_1^* \geq q_1^*$ where tildes indicate an alternative parameterization and solution. Using the same argument given for the effect of S_0 , this leads to a violation of the conditions for optimality. Therefore $\tilde{q}_1^* < q_1^*$ and $\frac{\partial q_1^*}{\partial g_b} < 0$.

g_c and g_q . These parameters cannot be signed since they affect both sides of the Euler equation in the same way.

r . Suppose $\tilde{r} > r$ and in response $\tilde{q}_1^* \leq q_1^*$ where tildes indicate an alternative parameterization and solution. Using the same line of argument given for the effect of S_0 , this leads to a violation of the conditions for optimality. Therefore $\tilde{q}_1^* > q_1^*$ and $\frac{\partial q_1^*}{\partial r} > 0$.

δ . Suppose $\tilde{\delta} > \delta$ and in response $\tilde{q}_1^* \leq q_1^*$ where tildes indicate an alternative parameterization and solution. Using the same argument given for the effect of S_0 , this leads to a violation of the conditions for optimality. Therefore $\tilde{q}_1^* > q_1^*$ and $\frac{\partial q_1^*}{\partial \delta} > 0$.

References

Anderson, E. 1986. Taxes vs. Quotas for Regulating Fisheries Under Uncertainty: A Hybrid Discrete-Time Continuous-Time Model. *Marine Resource Economics* 3:183–207.

Androkovich, R. A. and K. R. Stollery. 1991. Tax Versus Quota Regulation: A Stochastic Model of the Fishery. *American Journal of Agricultural Economics* 73(2):300–308.

Arrow, K. J., W. R. Cline, K-G. Maler, M. Munasinghe, R. Suitieri, J. E. Stiglitz. 1996. Intertemporal Equity, Discounting, and Economic Efficiency. In *Climate Change 1995: Economic and Social Dimensions of Climate Change* (eds Bruce, J., H. Lee, and E. Haites). p 125–144. Cambridge, UK: Cambridge University Press.

Bazelon, C. and K. Smetters. 1999. Discounting Inside the Washington, D.C. Beltway. *Journal of Economic Perspectives* 13(4): 213-228.

Council of Economic Advisors. 1999. *1999 Economic Report of the President*. Washington, DC: Council of Economic Advisors.

Dasgupta, P., P. Hammond, and E. Maskin. 1980. On Imperfect Competition and Pollution Control. *Review of Economic Studies* 47:857–981.

Energy Information Administration. 1998. *International Energy Outlook 1999*. Washington, DC: Energy Information Administration.

Energy Information Administration. 1999. *International Energy Database*. Washington, DC: Energy Information Administration.

Falk, I. and R. Mendelsohn. 1993. The Economics of Controlling Stock Pollutants: An Efficient Strategy for Greenhouse Gases. *Journal of Environmental Economics and Management* 25:76–88.

Hoel, M. and L. Karp. 1998. Taxes Versus Quotas for a Stock Pollutant. Unpublished manuscript, University of Oslo, Oslo.

International Monetary Fund. 1998. *World Economic Outlook*. Washington, DC: International Monetary Fund.

Kattenberg, A., Giorgi, F., Grassl, H., Meehl, G. A., Mitchell, J., Stouffer, R. J., Tokioka, T., Weaver, A. J., and Wigley, T. M. 1996. Climate Models—Projections of Future Climate. In *Climate Change 1995: The Science of Climate Change* (eds. Houghton, J. T., Meiro Filho, L. G., Callendar, B. A., Harris, N. Kattenberg, A., and Maskell, K.). p 285–358. Cambridge, UK: Cambridge University Press.

Keeler, E., M. Spence, and R. Zeckhauser. 1971. The Optimal Control of Pollution. *Journal of Economic Theory* 4:19–34.

Keeling, C. D. and T. P. Whorf. 1999. Atmospheric Carbon Dioxide Record from Mauna Loa. In *Trends: A Compendium of Data on Global Change*. Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, U.S. Department of Energy, Oak Ridge, TN.

Kitabatake, Y. 1989. Optimal Exploitation and Enhancement of Environmental Resources. *Journal of Environmental Economics and Management* 16:224–241.

Koenig, E. F. 1984a. Controlling Stock Externalities in a Common Property Fishery Subject to Uncertainty. *Journal of Environmental Economics and Management* 11:124–138.

Koenig, E. F. 1984b. Fisheries Regulation Under Uncertainty: A Dynamic Analysis. *Marine Resource Economics* 1:193–208.

Kolstad, C. D. 1996. Learning and Stock Effects in Environmental Regulation: The Case of Greenhouse Gas Emissions. *Journal of Environmental Economics and Management* 31(1):1–18.

Kwerel, E. 1977. To Tell the Truth: Imperfect Information and Optimal Pollution Control. *Review of Economic Studies* 44(3):595–601.

Laffont, J. J. 1977. More on Prices vs. Quantities. *Review of Economic Studies* 44:177–182.

Leggett, J., Pepper, W. J., and Swart, R. J. 1992. Emission Scenarios for the IPCC. In *Climate Change 1992: The Supplementary Report to the IPCC Scientific Assessment* (eds. Houghton, J. T., Callander, B. A., and Varney, S. K.). p 75–95. Cambridge, UK: Cambridge University Press.

Marland, G., T. A. Boden, R. J. Andres, A. L. Brenkert and C. Johnston. 1999. Global, Regional, and National CO₂ Emissions. In *Trends: A Compendium of Data on Global Change*. Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, U.S. Department of Energy, Oak Ridge, TN.

McKibbin, W. J. and P.J. Wilcoxen. 1997. A Better Way to Slow Global Climate Change. Policy Brief #17, The Brookings Institution, Washington, DC.

Neftel, A., H. Friedli, E. Moor, H. Lütscher, H. Oeschger, U. Siegenthaler, B. Stauffer. 1999. Historical Carbon Dioxide Record from the Siple Station Ice Core. In *Trends: A Compendium of Data on Global Change*. Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, U.S. Department of Energy, Oak Ridge, TN.

Nordhaus, W. D. 1994a. *Managing the Global Commons*. Cambridge: MIT Press.

Nordhaus, W. D. 1994b. Expert Opinion on Climate Change. *American Scientist* 82:45–51.

Pizer, W. A. 1997. Prices vs. Quantities Revisited: The Case of Climate Change. Discussion Paper 98-02, Resources for the Future, Washington, DC.

Plourde, C. 1972. A Model of Waste Accumulation and Disposal. *Canadian Journal of Economics* 5(1):119–125.

Plourde, C. and D. Yeung. 1989. A Model of Industrial Pollution in a Stochastic Environment. *Journal of Environmental Economics and Management* 16(2):97–105.

Roberts, M. J. and M. Spence. 1976. Effluent Charges and Licenses Under Uncertainty. *Journal of Public Economics* 5(3,4):193–208.

Roughgarden, T. and S. H. Schneider. 1999. Climate Change Policy: Quantifying Uncertainties for Damages and Optimal Carbon Taxes. *Energy Policy* 27:415–429.

Schimel, D.S. et al. 1995. Radiative Forcing of Climate Change and an Evaluation of the IPCC IS92 Emissions Scenarios. In *Climate Change 1995: The Science of Climate Change* (eds Houghton, J. T. et al). p 35–71. Cambridge, UK: Cambridge University Press.

Smith, V. L. 1972. Dynamics of Waste Accumulation: Disposal Versus Recycling. *Quarterly Journal of Economics* 86:600–616.

Stavins, R. N. 1996. Correlated Uncertainty and Policy Instrument Choice. *Journal of Environmental Economics and Management* 30:218–232.

Weitzman, M. L. 1974. Prices vs. Quantities. *Review of Economic Studies* 41(4):477–491.

Weitzman, M. L. 1978. Optimal Rewards for Economic Regulation. *American Economic Review* 68:683–691.

Weyant, J. P. and J. Hill. 1999. Introduction and Overview. In *The Costs of the Kyoto Protocol: A Multi-Model Evaluation*, a Special Issue of *The Energy Journal*.

Yohe, G. W. 1978. Towards a General Comparison of Price Controls and Quantity Controls Under Uncertainty. *Review of Economic Studies* 45:229–238.