NBER WORKING PAPER SERIES

INDEXED REGULATION

Richard G. Newell
William A. Pizer

Working Paper 13991
http://www.nber.org/papers/w13991

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2008

We thank Drew Baglino for research assistance and Arun Malik, two anonymous referees, Carolyn Fischer, Ulf Moslener, John Parsons, Wally Oates, Brian McLean, David Evans, and participants in seminars at RFF, FEEM, HEC Montréal, and the Southern Economic Association Annual Meetings for useful comments on previous versions of the paper. We acknowledge funding from MISTRA, the Swedish Foundation for Strategic Environmental Research. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

Indexed Regulation
Richard G. Newell and William A. Pizer
NBER Working Paper No. 13991
May 2008
JEL No. C68,D81,H41,Q54,Q58

ABSTRACT

Seminal work by Weitzman (1974) revealed prices are preferred to quantities when marginal benefits are relatively flat compared to marginal costs. We extend this comparison to indexed policies, where quantities are proportional to an index, such as output. We find that policy preferences hinge on additional parameters describing the first and second moments of the index and the ex post optimal quantity level. When the ratio of these variables' coefficients of variation divided by their correlation is less than approximately two, indexed quantities are preferred to fixed quantities. A slightly more complex condition determines when indexed quantities are preferred to prices. Applied to climate change policy, we find that the range of variation and correlation in country-level carbon dioxide emissions and GDP suggests the ranking of an emissions intensity cap (indexed to GDP) compared to a fixed emission cap is not uniform across countries; neither policy clearly dominates the other.

Richard G. Newell
Nicholas School of the Environment
Duke University
Box 90227
Durham, NC 27708
and NBER
richard.newell@duke.edu

William A. Pizer
Resources for the Future
1616 P Street, NW
Washington, DC 20036
pizer@rff.org

Indexed regulation

Richard G. Newell and William A. Pizer

1. Introduction

The literature on policy instrument choice under uncertainty historically has focused on the relative performance of prices, quantities, and hybrid price–quantity instruments [9; 12; 15; 18; 21; 22; 24]. In practice, however, the decision for policymakers often comes down to choosing among different types of quantity-based instruments, not choosing between prices and quantities. Nowhere is this better illustrated than in the debate on the form and implementation of measures to address global climate change. While the Kyoto Protocol and current U.S. congressional proposals for a greenhouse gas cap-and-trade system focus on a quantity-based system with largely fixed emissions targets, the United States' Bush administration and Australia have so far embraced targets based on emission intensity—a quantity target indexed to economic activity.[1] Canada has committed to a quantity target under the Kyoto Protocol, but is pursuing an intensity-based approach in its domestic program. Japan has similarly pursued voluntary sectoral emission intensity goals. Few countries have chosen the relevant price instrument, a carbon tax, and this option is largely taboo in the United States.[2]

[1] The EU Emission Trading Scheme (ETS) is a somewhat intermediate case because not all sectors are covered under the system and allocations can be adjusted in response to output growth (i.e., new installations) and decline (i.e., installation shutdowns).

[2] Finland and Sweden are currently the only countries that have a broad-based carbon tax, in addition to participating in the EU ETS. New Zealand initially proposed then rejected a carbon tax in favor of a cap-and-trade system. Japan has considered but not enacted proposals for a carbon tax.

This paper considers the welfare implications of this indexed versus fixed quantity distinction and reveals a simple condition for preferring one to the other. When the coefficient of variation in the index divided by the coefficient of variation in the ex post optimal quantity level is less than approximately twice their correlation, indexed quantities are preferred. The question is whether the variance of the index is "too much," with the factor of two reflecting adjustment by double the desired adjustment (given the useful information in the index) to the point where the over-adjustment erases any gain. Applying these results to the question of indexed versus fixed emissions limits to address global climate change, we find that the empirical range of variations and correlations for emissions and GDP is such that a cap on emission intensity (emissions indexed to GDP) may be preferred to a fixed emission cap for some countries, but not for others.

Of course, interest in indexed and fixed quantity regulation is not limited to climate change. Environmental policy in the United States is replete with examples of both kinds of quantity policies. The most familiar example to many is the U.S. sulfur dioxide tradable permit or "cap-and-trade" system for electricity generators [2; 17]. Since 1995, this system has set a fixed limit on the tons of sulfur dioxide emitted from power plants, while allowing sources to trade emissions allowances in order to minimize compliance costs. The NOx Budget Program, the Clean Air Interstate Rule, and Clean Air Mercury Rule (all promulgated by the Environmental Protection Agency under the Clean Air Act), along with a host of regional air and water trading programs, round out the U.S. experience with fixed targets.

Despite these examples, performance standards are a more common form of environmental regulation, typically set in terms of an allowable emissions rate per unit of product output (i.e., emission intensity) [5; 16]. The phase-down of lead in gasoline—the first large-scale experiment with market-based environmental policy—employed a tradable performance standard [7; 11]. Over 25 U.S. states now have (and the U.S. Congress has proposed) renewable portfolio standards that require a certain share of electricity generation from renewable sources. Corporate Average Fuel Economy (CAFE) standards are a less flexible performance standard that can be traded and banked within but not across firms (e.g., across vehicle lines within a firm). Even less flexible are traditional command-and-control style regulations, such as New Source Performance Standards.

Like an intensity target for greenhouse gases, these forms of regulation allow the effective emissions cap to adjust in response to changes in output. This feature has political appeal because it provides a way to set environmental standards that are less likely to be, or perceived to be, constraints on economic growth, either within a regulated sector or across the economy. Intuition suggests that the responsiveness of intensity-based quantity regulation to output changes also may have economic appeal. Such adjustments could lower the expected costs of achieving a particular environmental target by loosening the cap when costs are unexpectedly high and tightening it when costs are low.[3] This is analogous to the cost advantage of prices over quantities identified in the literature.

On the other hand, including an index in the policy formula introduces another uncertain variable and potentially unrelated noise, which could have negative consequences for efficiency. The purpose of this paper is to clarify these tradeoffs by identifying the key features of the regulatory problem, modeling the relative performance of fixed versus indexed quantity targets, and demonstrating how the resulting framework can be applied using the case of climate policy.

[3] Linking an emissions limit to output raises questions of both subsidizing output and creating pro-cyclical costs, points we consider at the end of this paper.

We note that all of our theoretical results apply, almost without modification, to the case of indexed price policies.[4]

Several recent papers have looked both theoretically and empirically at the relative advantages of intensity targets in the case of climate change policy. Quirion [14] finds that strong positive correlation between a cost shock and the index favors indexed quantities, similar to our results.[5] His analytic model differs from ours in that he considers a particular functional form where marginal costs depend on emissions as a share of baseline emissions, and where baseline emissions are the index, rather than a more general linear approximation. He argues that indexed quantities typically rank between prices and quantities in terms of expected net benefits, a result we can replicate by applying the restrictions implied by his functional form to our model. However, we find such restrictions at odds with empirical evidence, as discussed further in sections 2.5 and 3.

Focusing solely on costs (and ignoring benefits), Ellerman and Sue Wing [3] employ a simulation model to argue that partial indexing—what we call a general indexed quantity—is more sensible. Such an approach sets the mean emissions level and the rate of adjustment to the index separately, rather than allowing a single parameter (the emissions rate) to determine both.[6]

---

4 The obvious exception being the large number of policies and contracts with nominal values indexed to inflation, as well as natural gas contracts that are sometimes indexed to crude oil prices.

5 In Quirion's model, stronger positive correlation between the cost shock and the index is implicit in a higher variance in baseline emissions (see his proposition 4). His index is baseline emissions, which is assumed to be perfectly correlated with marginal costs. Therefore, a higher variance in baseline emissions means a larger portion of uncertainty in costs (comprised of both uncertainty about baseline emissions and an independent shock) is due to the index, and thus correlation between the two is higher.

6 Compare $q = a + rx$ with and without the parameter a, where q is emissions, x is the index, and r is the rate of adjustment to the index. We address the distinction between general and ordinary (proportional) indexing further below.

Similar to Ellerman and Sue Wing, Jotzo and Pezzey [6] derive an optimal indexing rule based on minimizing expected costs and use a simulation model to evaluate both a general indexed quantity policy and an ordinary, proportional indexed quantity policy. They conclude that for climate policy, indexes of either type are better than fixed quantity policies at a global level, that the more general index is considerably better than ordinary indexing, and that the rate of indexing varies greatly among countries.

Finally, Sue Wing et al. [19] conclude that sufficiently small GDP variance and high correlation favor indexed quantities for climate policy—as we do—but based on minimizing expected costs rather than maximizing expected benefits minus costs (note the ranking would be same if marginal benefits are flat). They also present empirical evidence supporting indexed quantities over fixed quantities, with strong support for developing countries and more equivocal results for industrialized countries.

Our work ties together and clarifies this literature by deriving simple analytic expressions for the advantage of indexed quantities relative to both fixed quantity and price controls. Two factors tend to tilt the preference toward indexed quantities relative to fixed quantities: stronger positive correlation between the index and marginal cost uncertainty and relatively small index variance. The intuition is straightforward. With low correlation between the index and the cost shock, indexing introduces unwanted noise in the target without reducing cost uncertainty. Further, a large index variance, relative to the marginal cost variance, will over-adjust quantities even if there is perfect correlation. Whether indexed quantities also dominate prices depends, in addition, on marginal benefits being sufficiently steep relative to marginal costs—as in the original Weitzman result.

Applying the analytic results to the case of greenhouse gas mitigation, we find that the ranking of carbon dioxide caps indexed to GDP compared to fixed quantity controls varies across countries. Those countries with a strong correlation between output and emissions and relatively low output variance tend to favor indexed quantities, while those with low correlation and/or high output variance favor fixed quantities.

Our motivation and application relate to cases where the government is seeking to regulate a market constrained by both an information asymmetry (between the moment a policy is determined and the horizon over which it applies) and a limit on regulatory complexity. Similar features characterize other mechanism design problems where these results may be helpful. Sales contracts, for example, face an asymmetry of information between the moment a contract is agreed and when it is executed and, similarly, a limit to the complexity of contract contingencies. Like our regulatory example, subjecting delivery quantity and/or price to indexing rules can enhance contract performance [see, e.g., 1; 8]. In the case of monetary policy, the uncertain link between the instrument (current interest rate or money supply) and outcome (future inflation and output) mimics our information asymmetry. The literature on monetary policy considers ways to use all available information to improve performance but often ends up, like our regulatory example, with fairly simple linear index rules to minimize squared errors [see 20 for a recent summary]. These examples suggest a potentially broad application of our analysis beyond the regulatory arena.

The remainder of the paper is organized as follows. First, we set up our model and review the original Weitzman [21] result. Next, we introduce the notion of indexed quantities and derive results for both a general indexed quantity—where the mean quantity level and rate of adjustment to the index are distinct—as well as a simple proportional indexed quantity. Finally, we present an application to climate change and conclude.

2. Model and analytic results

Our modeling approach follows Weitzman [21] with quadratic cost and benefit functions for a generic market, q. The functions can be viewed as local approximations about an arbitrary point. Maximizing net benefits based on these functions, we determine expected net benefits for optimal price and fixed quantity controls, as in Weitzman, and then for an indexed quantity policy. We consider two types of indexed policies, "ordinary" indexed quantities where the regulated quantity equals a fixed rate times the index, and "general" indexed quantities where the rate of adjustment is distinct from the mean level of control. We derive expressions for the difference in net benefits for pair-wise comparisons of the policies, evaluate the dependence of these policy rankings on key parameters, and summarize these rankings in a two-dimensional space defined by key parameter values.

2.1. Review of prices versus quantities

We start by replicating the basic Weitzman [21] results with costs and benefits measured as quadratic functions about the expected optimal quantity q^*.[7] Note that for the case of pollution control, we take q to represent emissions, not abatement. Costs are given by

(1) $$C(q) = c_0 + (c_1 - \theta_c)(q - q^*) + \frac{c_2}{2}(q - q^*)^2,$$

[7] Like Weitzman, we make the approximation around the optimal fixed quantity q^* without loss of generality for the resulting comparative advantage expressions. A more general approximation simply adds constant terms to all of the expected net benefit expressions, which cancel out when they are compared.

where θ_c is a mean-zero random shock to marginal costs with variance σ_c^2, and the c_n parameters capture constant, linear, and quadratic behavior. We assume $c_2 > 0$; that is, costs are strictly convex. Note that we have defined the cost shock such that a positive value of θ_c reduces the marginal cost of *producing q*, but increases the marginal cost of *reducing q*. We chose this specification to ease the interpretation for our application to pollution control, where the regulator typically is seeking emissions reductions.

Similar to costs, benefits are given by the form

(2) $$B(q) = b_0 + b_1\left(q - q^*\right) - \frac{b_2}{2}\left(q - q^*\right)^2,$$

with $b_2 \geq 0$; that is, benefits are weakly concave (marginal benefits are non-increasing). We do not include benefit uncertainty because, unless it is correlated with cost uncertainty, it does not affect net benefits in this quadratic setting [18; 21].[8] The remaining parameters, particularly the linear terms, can be negative. This is relevant for our motivating example of pollution where marginal benefits are negative, and increasingly so, for increases in *q*.

Differentiating (1) and (2) with respect to *q* to obtain marginal costs and benefits, taking expectations, and equating the expressions, yields the condition

$$b_1 - b_2\left(q - q^*\right) = c_1 + c_2\left(q - q^*\right)$$

[8] Uncertainty in marginal benefits would influence the price-quantity policy instrument ranking if it is correlated with marginal costs, tilting the preference toward (from) quantities in the case of positive (negative) correlation [18; 21]. Regarding indexed quantities, positive (negative) correlation between the *index* and marginal benefits would similarly tilt the preference toward (from) fixed quantities. The reasoning is the same: both prices and indexed quantities allow a less stringent policy in states of the world with high costs / high index values—a feature that is less desirable when benefits also tend to be high in those same states of the world (and vice versa). Uncertainty in marginal benefits will also influence the *stringency* of policy if marginal benefits are non-linear.

for the optimal fixed quantity policy, a condition satisfied at $q = q^*$ if and only if $b_1 = c_1$. We now see the implication of approximating benefits and costs around the optimal fixed quantity policy (q^*). That is, marginal benefits equal expected marginal costs ($b_1 = c_1$) at the optimum. Taking this into account and subtracting (1) from (2) we have a general equation for net benefits as a function of q:

(3) $$NB(q) = b_0 - c_0 + \theta_c\left(q - q^*\right) - \frac{b_2 + c_2}{2}\left(q - q^*\right)^2 .$$

Expected net benefits under the optimal fixed quantity policy, NB_Q, are therefore given by

(4) $$E\left[NB_Q\right] = b_0 - c_0$$

because the optimal fixed quantity policy is q^*.

An arbitrary price policy p equates marginal cost to price, ex post. That is, $p = c_1 - \theta_c + c_2\left(q - q^*\right)$, with an associated quantity response function of $q_p\left(\theta_c\right) = q^* + \left(p - c_1 + \theta_c\right)/c_2$. The optimal price policy, p^*, equates marginal benefits and expected marginal costs, given the response function $q_p\left(\theta_c\right)$. It is straightforward to show that

(5) $$q_p^*\left(\theta_c\right) = q^* + \frac{\theta_c}{c_2},$$

with the implication that the optimal price equals the expected price at the optimal fixed quantity and yields the optimal fixed quantity (q^*) in expectation. Substituting (5) for q in (3) and taking expectations we find the expected net benefits under the optimal price policy, NB_P, are

(6) $$E\left[NB_P\right] = b_0 - c_0 + \frac{\sigma_c^2(c_2 - b_2)}{2c_2^2}.$$

Taking the difference between $E[NB_P]$ and $E[NB_Q]$ yields the familiar Weitzman [21] relative advantage expression for prices versus fixed quantities

(7) $$\Delta_{P-Q} = \frac{\sigma_c^2 (c_2 - b_2)}{2c_2^2}.$$

Prices therefore outperform fixed quantities if the slope of marginal benefits is less than the slope of marginal costs, and vice versa.

At this point, it is useful to define both the first-best policy and the associated net benefits. Setting marginal costs equal to marginal benefits after the shock θ_c is revealed yields what we refer to as the optimal ex post quantity, q_O:

(8) $$q_O(\theta_c) = q^* + \frac{\theta_c}{c_2 + b_2}.$$

Intuitively, the sum of the cost and benefit slopes $c_2 + b_2$ reflects the rate at which a deviation in quantities translates into a deviation in net benefits, where here the net benefit deviation equals θ_c. This in turn leads to an expected value for net benefits under the first-best policy, NB_O of

(9) $$E[NB_O] = b_0 - c_0 + \frac{\sigma_c^2}{2(c_2 + b_2)}.$$

Graphically, we can visualize the outcomes under the first-best, price, and fixed quantity policies in Figure 1 for a particular realization of θ_c. The larger hatched area represents the loss under the fixed quantity policy and the smaller shaded area represents the loss under the price policy, both relative to the first best. From the figure, we can see that while the price policy misses the optimum because it over-adjusts the expected quantity target, the fixed quantity policy misses the optimum because it fails to adjust at all in response to the cost shock.

The divergence in performance of price and quantity controls from the optimum and from one another arises because of an information asymmetry. The regulator does not observe the cost shock θ_c that, in contrast, is known to the regulated firms at the time q is chosen. Once the information is revealed, it is not possible to rapidly adjust the policy, and we find that fixed policies lead to second-best outcomes with the well-known distinction between prices and quantities.

An important observation at this point is that an alternate policy could improve upon both fixed prices and quantities if somehow it adjusted the ex post quantity level in a way that was closer to the optimum than either of these instruments. In this regard, three things are immediately necessary: the adjustment should be correlated to the cost shock; the adjustment should not be too small; and it especially should not be too large.[9] We now turn to how indexed quantities might achieve this end.

2.2. General indexed quantities

We consider a random variable, x, that is used to index the otherwise fixed quantity policy. In the pollution case, it is useful to think of x as some observed index of activity, such as output (e.g., kWh of electricity or gallons of fuel), that is correlated with the level of unregulated pollution. More generally, it could be anything related to the object of regulation, q, including weather or prices in related markets. We assume a linear policy of the form

(10) $$q(x) = a + rx,$$

where a and r are policy parameters, $E[x] = \bar{x}$, $\text{var}(x) = \sigma_x^2$, and $\text{cov}(x, \theta_c) = \sigma_{cx}$. That is, the index has mean $\bar{x}$, standard deviation σ_x, and covariance σ_{cx} with the cost shock θ_c.

9 It is also necessary that information about the index become available alongside information about the shock. Learning about an index adjustment after firms have made their final decision about q is of little use.

As an example of an indexed quantity policy, the U.S. phase-down of lead in gasoline established rate limits in terms of grams of lead per gallon of gasoline, with the eventual quantity limit equaling the fixed rate times the volume of gasoline produced. The volume of gasoline represented an unknown, random variable to the regulator at the time of regulation and introduced variation in the ex post quantity of lead released into the environment.

We consider two types of policies, a general indexed quantity (GIQ) policy where no restrictions are placed on the parameters *a* and *r*, and an ordinary indexed quantity (or just indexed quantity, IQ) policy where we constrain $a = 0$. In practice, the latter is the much more common form of regulation, where the regulated level of *q* is simply a multiple *r* of random variable *x*, as in the U.S. phase-down of lead in gasoline.

Substituting the indexed quantity rule (10) into (3), maximizing net benefits with respect to *r* and *a*, and taking expectations yields the optimal form of the GIQ policy:

(11) $$q^*_{GIQ}(x) = q^* + r^{**}(x - \bar{x}),$$

where

(12) $$r^{**} = \left(\sigma_{cx}/\sigma^2_x\right)\big/\left(b_2 + c_2\right).$$

Repeating these steps but maximizing only with respect to *r*, while constraining $a = 0$, yields the optimal form of the ordinary indexed quantity (IQ) policy:

(13) $$q^*_{IQ}(x) = r^* x,$$

where

(14) $$r^* = \frac{\left(q^*/\bar{x}\right) + v_x^2 r^{**}}{1 + v_x^2},$$

and $v_x = \sigma_x/\bar{x}$ (the coefficient of variation in *x*). Note that while the optimal GIQ policy delivers the optimal fixed quantity q^* in expectation (as with the other policies), the optimal IQ policy does not generally yield q^* in expectation, although it will be quite close in typical applications where v_x is small.

Intuitively, using two parameters permits the GIQ policy to match both the mean and variance of the optimal quantity adjustment. Using only one parameter (as in the IQ policy) requires a trade-off between matching the mean and variance. This points to an underlying question of the paper: does this trade-off on net improve welfare compared to the simpler approach of using fixed quantities to set the mean at the optimum ignoring *x*? For the remainder of this section, we focus on the GIQ policy, returning to discuss the ordinary IQ policy in the next section.

The parameter r^{**} equals the coefficient of a regression of the first-best optimal adjustment ($\theta_c/(b_2+c_2)$ from (8)) on *x*. Therefore, we can interpret the GIQ policy as the best linear predictor of the first-best adjustment, given *x*. If *x* and θ_c are jointly normal, the GIQ policy is also the minimum variance unbiased predictor (i.e., including the possibility of non-linear predictors). This result is extended easily to the case of multiple index variables, where *x* would be a vector of index variables and r^{**} would be a vector of regression coefficients.

Substituting (11) for *q* in (3) and taking expectations, we derive expected net benefits of the GIQ policy equal to

(15) $$E\left[NB_{GIQ}\right]=b_0-c_0+\frac{\sigma_c^2}{2(c_2+b_2)}\rho_{cx}^2,$$

where $\rho_{cx}=\sigma_{cx}/(\sigma_x\sigma_c)$, the correlation of *x* and θ_c. Comparing this to the expected net benefits under the first-best policy given in (9), we can see that the GIQ policy achieves the first best if $\rho_{cx}=1$, that is, if the index and cost shock are correlated perfectly. In other words, if we have an exogenous, observable index variable that perfectly reveals the cost shock, the information problem that creates our second-best setting vanishes and we can implement a first-best policy.

More generally, the gain from the GIQ policy will depend on the squared correlation, which can be interpreted as the goodness of fit (R^2) of a regression of the first-best optimal adjustment on the index. Thus, the degree to which the index can predict the underlying cost shock, in terms of predicted versus residual variation, determines the degree to which the indexed policy achieves the first-best result given in (9). Meanwhile, the GIQ policy is always at least as good as the fixed quantity policy, with the relative advantage given by

(16) $$\Delta_{GIQ-Q} = \frac{\sigma_c^2}{2(c_2 + b_2)} \rho_{cx}^2 \geq 0 .$$

This expression is always non-negative and tends to zero as the correlation goes to zero. Similar observations about the ability of the GIQ policy to always outperform fixed quantities are made by both Jotzo and Pezzey [6] and Sue Wing et al. [19].

Subtracting (6) from (15) yields the relative advantage of the GIQ policy over the price policy:

(17) $$\Delta_{GIQ-P} = \frac{\sigma_c^2}{2(c_2 + b_2)} \left(\rho_{cx}^2 + \left(\frac{b_2}{c_2} \right)^2 - 1 \right).$$

The GIQ policy will therefore be preferred if benefits are sufficiently steep (as with a fixed quantity policy) or if correlation between the index and the cost shock is high. Put another way, the preference for the GIQ policy versus prices is a competition between the relative flatness of marginal benefits (pushing Δ_{GIQ-P} negative) and the correlation between the index and the cost shock (pushing Δ_{GIQ-P} positive).

Figure 1 shows the loss under the GIQ policy as thickly outlined for a case where it adjusts for roughly half of the observed cost shock. As indicated, the GIQ policy will have an expected loss no larger than the quantity policy, but its advantage relative to the price policy hinges on the relative slopes and degree of correlation between the index and shock.

While not the focus of this paper, we note that a generalized indexed price (GIP) policy of the form $p^*_{GIP}(x) = p^* + u^{**}(x - \bar{x})$ also is possible, equaling the optimal fixed price policy plus an adjustment rate u^{**} times the deviation in the index from its expectation. As in the GIQ case, the optimal adjustment rate equals a regression coefficient but this time for the optimal ex post price regressed on the index. Similar to the relative advantage of the GIQ policy over fixed quantities, the relative advantage of the GIP policy over fixed prices equals the difference between the first-best welfare gain and the fixed price policy, times the correlation squared, $\Delta_{GIP-P} = 1/2\left(\sigma_c^2/(c_2+b_2)\right)\left(b_2/c_2\right)^2 \rho_{cx}^2$. As the correlation goes to unity, the GIP policy achieves the first-best outcome; as it tends to zero, it becomes the same as the fixed price policy.

While it is easy to imagine general indexed quantities of the form (10), or even the general indexed prices noted above, in practice we see very few (if any)—much as we see very few price–quantity hybrid policies along the lines of Roberts and Spence [15] or Pizer [12]. For that reason, we now focus on the ordinary indexed quantity policies given by (13).

2.3. (Ordinary) indexed quantities

Consider the more common case in practice where the regulated quantity is strictly equal to a rate times the index variable, $q_{IQ}(x) = r\,x$, where we have imposed the constraint that $a = 0$ in (10). We noted above that the optimal indexed quantity is

(13) $$q^*_{IQ}(x) = r^* x\,,$$

where

(14) $$r^* = \frac{\left(q^*/\bar{x}\right) + v_x^2 r^{**}}{1 + v_x^2}.$$

Like the optimal rate for the GIQ policy (r^{**}), the optimal indexed quantity rate r^* can be interpreted as a regression coefficient when the first-best adjustment $\theta_c/\left(c_2 + b_2\right)$ is regressed on *x*—but this time with the constraint that a constant term is not included—a point we return to below.

Expression (14) for r^* is a weighted average of two terms, $q^*/\bar{x}$ and r^{**}, with the relative weight depending on the coefficient of variation of *x*. As the index variation becomes large, r^* tends to the GIQ regression coefficient r^{**} and the *variance* of the IQ and GIQ policies converge. As variation in the index tends to zero, r^* tends to $q^*/\bar{x}$, and the *mean* of the IQ and GIQ (and fixed quantity) policies converge. Because r^* cannot simultaneously match both the mean and variance of the GIQ policy (unless it happens that $a = 0$ even when unconstrained), (14) represents the minimum variance solution given that there is only one rather than two flexible parameters. An implication, nonetheless, is that r^* does not generally yield the optimal fixed quantity q^* in expectation, as do the other policies, although it will be quite close in typical applications where v_x is small.

Substituting (13) for *q* in (3) and taking expectations, we derive the expected net benefits of the optimal IQ policy,

(18) $$E\left[NB_{IQ}\right] = b_0 - c_0 - \frac{b_2 + c_2}{2}\frac{\left(q^*/\bar{x}\right)^2}{1 + v_x^2}\sigma_x^2 + \frac{\sigma_c^2}{2\left(b_2 + c_2\right)}\frac{2\left(v_{q^*}^2/v_x\right) + v_x^2\rho_{cx}}{1 + v_x^2}\rho_{cx},$$

where we have defined $v_x = \sigma_x/\bar{x}$ and

(19) $$v_{q^*} = \left(\sigma_c/\left(b_2 + c_2\right)\right)/q^*,$$

the coefficient of variation in the index and the ex post optimal quantity from (8), respectively.

Arranged this way, the expression (18) highlights two important results. First, if there is no correlation between the index and the cost shock, the last term vanishes and variance in the index reduces expected net benefits based on the third term. This follows from Jensen's inequality applied to the fact that net benefits (costs minus benefits) are a concave function of the regulated quantity level and that higher variance in the index implies higher variance in the indexed quantity level and lower expected net benefits. Second, for a given index variance—that is, holding the third term constant—correlation between the index and cost shock improves net benefits based on the fourth term.

If $\rho_{cx} \neq 0$, it is useful to rearrange (18) to yield,

$$E\left[NB_{IQ}\right] = b_0 - c_0 + \frac{\sigma_c^2}{2(c_2 + b_2)}\rho_{cx}^2\left(1 - \frac{1}{1+v_x^2}\left(1 - \frac{v_x}{\rho_{cx}v_{q^*}}\right)^2\right), \tag{20}$$

Note that $v_x / \left(\rho_{cx} v_{q^*}\right) = \left(q^* / \overline{x}\right) / r^{**}$, the ratio of the two terms being averaged to determine the index rate in (14).

Comparing (15) and (20), we can see that the net benefit expression for ordinary indexed quantities is the same as for the GIQ policy, except that it contains an extra factor multiplying the third term. Holding other parameters fixed, increases in v_x improve the performance of the IQ policy up to the point where

$$\frac{v_x}{\rho_{cx} v_{q^*}} = 1, \tag{21}$$

beyond which further increases in v_x worsen the performance of IQs. When this condition is met exactly, the term in (outer) parentheses of (20) equals one and the expected welfare gain from the IQ policy equals the expected gain from the GIQ policy.

How do we interpret the condition given in (21)? One way is to recall, as noted above, that $v_x / (\rho_{cx} v_{q*}) = (q^* / \bar{x}) / r^{**}$, yielding the condition $(q^* / \bar{x}) / r^{**} = 1$ and, therefore, $r^* = r^{**}$ from the definition of r^* in (14). That is, the rate of adjustment is the same under the IQ policy and the GIQ policy. As noted above, both r^{**} and r^* are regression coefficients in models predicting the first-best adjustment as a function of *x*, the former with a constant and the latter without.[10] If the two regression coefficients happen to be equal thanks to a lucky or thoughtful choice of the index variable, it implies also that the freely estimated constant *a* in the GIQ policy equals zero, the IQ policy yields the same response function $q(x)$ as the GIQ policy, and it performs just as well. However, as $(q^* / \bar{x}) / r^{**}$ diverges from unity, r^* diverges from r^{**}. This divergence reflects an increased importance of the non-zero constant term in the regression model, and the ordinary IQ policy does increasingly worse than the more flexible GIQ policy.

As an alternative interpretation of (20) and the resulting condition (21), we can think about the "desired" value of v_x for an IQ policy, given the values of v_{q*} and the correlation ρ_{cx}. How much variation should there be in the index in order to maximize net benefits? If an index's correlation with the underlying cost shock is perfect, it makes sense to have the index vary by just as much as the ex post optimal quantity (i.e., $v_x = v_{q*}$). At the other extreme, when the correlation is zero, it is preferable to have an index with no variation because the index is all noise with respect to the cost shock and optimal quantity. Likewise, for cases between these two

10 In general, a regression coefficient with and without a constant will be the same if the coefficient of variation of the explanatory variable equals the dependent variable's coefficient of variation, times the correlation between the variables—exactly the condition in (21).

extremes, expression (21) reveals that as the correlation declines the desired variation in the index should also decline.

As the level of index variation deviates from the desired level, the performance of the IQ policy deteriorates. In particular, very noisy indexes will tend to a limiting net benefit given by

$$\lim_{v_x \to \infty} E\left[NB_{IQ}\right] = b_0 - c_0 + \frac{\sigma_c^2}{2(c_2 + b_2)} \rho_{cx}^2 \left(1 - \left(\frac{1}{\rho_{cx} v_{q*}}\right)^2\right).$$

In contrast, indexes with too small a coefficient of variation tend toward the fixed quantity result, a point we now confirm.

2.4. The advantage of indexed quantities relative to prices and quantities

We can now calculate the relative advantage of indexed quantities to prices and fixed quantities. The relative advantage of indexed quantities over fixed quantities is

(22) $$\Delta_{IQ-Q} = \frac{\sigma_c^2}{2(c_2 + b_2)} \rho_{cx}^2 \left(1 - \frac{1}{1 + v_x^2}\left(1 - \frac{v_x}{\rho_{cx} v_{q*}}\right)^2\right).$$

In terms of the parameter v_x, this expression equals zero when v_x equals zero, reaches a maximum of Δ_{GIQ-Q} when $v_x = \rho_{cx} v_{q^*}$ and then declines to $\Delta_{GIQ-Q}\left(1 - \left(\rho_{cx} v_{q*}\right)^{-2}\right)$ as $v_x \to \infty$. Based on these tendencies, if $\rho_{cx} v_{q*} \geq 1$, this expression is non-negative for all values of v_x and indexed quantities are always at least as good as, and usually better than, fixed quantities.

In all practical cases, however, $\rho_{cx} v_{q*} < 1$, because $\rho_{cx} \leq 1$ by definition and $v_{q*} < 1$ unless the first-best optimum is highly variable relative to its mean—an unusual case that would in any event be inappropriate for the modeling framework we have set out. With $\rho_{cx} v_{q*} < 1$, indexing is preferred to fixed quantities so long as $v_x \big/ \left(\rho_{cx} v_{q*}\right) < 2 \Big/ \left(1 - \left(\rho_{cx} v_{q*}\right)^2\right)$. We can simplify this condition by further focusing on cases where the variation is not only less than one, but relatively small (i.e., $v_{q*} \ll 1$), which leads to the approximate condition $v_x \big/ \left(\rho_{cx} v_{q*}\right) < 2$ for indexed quantities to be preferred. Such a focus already is implicit given the framing of our problem as a local quadratic approximation around the expected optimum and likewise seems reasonable for practical targets of regulation.

The intuition for this latter condition is straightforward to understand. We previously observed that for parameter values satisfying $v_x \big/ \left(\rho_{cx} v_{q*}\right) = 1$ the indexed quantity matches the GIQ policy. That is, the expression can be re-written as $\left(q^* / \overline{x}\right) \big/ r^{**} = 1$ and the indexed quantity rate of adjustment $r*$ equals the GIQ rate of adjustment r^{**} based on (14). Now imagine parameter values whereby $v_x \big/ \left(\rho_{cx} v_{q*}\right) = \left(q^* / \overline{x}\right) \big/ r^{**} = 2$, corresponding to the threshold condition for preferences between indexed and fixed quantities to flip. Under the noted assumption $v_{q*} \ll 1$, we know $v_x \ll 1$ over the relevant range near v_{q*} and therefore $r^* \approx q^* \big/ \overline{x}$ based on (14). Under this assumption, the adjustment based on this rate is approximately double what it should be compared to the GIQ policy (i.e., $r^* \big/ r^{**} \approx 2$) and all the expected gains relative to the fixed quantity policy are squandered by overshooting the new expected optimum, as shown in Figure 2. The surprisingly simple result that the point of indifference between indexing and fixed quantities occurs at $r^* \big/ r^{**} \approx 2$ is attributable to the linear marginal form assumptions, which imply that equal-sized positive and negative deviations from the optimum have equal and opposite effects on marginal net benefits.[11]

When v_{q*} and v_x are closer to one, r^* is an average of $q^*/\bar{x}$ and r^{**} based on (14), reflecting the fact that we are willing to trade off higher mean error to better match variance and reduce the mean-squared error. Therefore, $v_x/(\rho_{cx} v_{q*})$ can actually be slightly larger than 2 before indexed quantities squander their gain over fixed quantities. Specifically, for values of $2/\left(1-(\rho_{cx} v_{q*})^2\right)$ or smaller, indexed quantities continue to be preferred.

What about prices? The relative advantage of indexed quantities over prices is given by

(23) $$\Delta_{IQ-P} = \frac{\sigma_c^2}{2(c_2+b_2)}\left(\left(\frac{b_2}{c_2}\right)^2 - 1 + \rho_{cx}^2\left(1-\frac{1}{1+v_x^2}\left(1-\frac{v_x}{\rho_{cx} v_{q*}}\right)^2\right)\right).$$

Based on the expression in outer parentheses, the sign of Δ_{IQ-P} is positive or negative depending on whether $(b_2/c_2)^2$ is greater or less than $1-\rho_{cx}^2\left(1-(1+v_x^2)^{-1}\left(1-\left(v_x/(\rho_{cx} v_{q*})\right)\right)^2\right)$. As illustrated further below, the expression can be viewed as a parabola-like function in $v_x/(\rho_{cx} v_{q*})$. with a maximum at $v_x/(\rho_{cx} v_{q*}) = 1$, where it equals $(b_2/c_2)^2 - 1 + \rho_{cx}^2$ and where $\Delta_{IQ-P} = \Delta_{GIQ-P}$ (matching the general indexed quantity policy comparison). Thus indexed quantities are preferred to prices when marginal benefits are relatively steep and/or when correlation with the index is high, as was the case comparing the GIQ policy to prices.

---

11 Note that if marginal costs are convex, the critical value would be more than 2; conversely, if marginal benefits are convex, the critical value would be less than 2.

As $v_x/(\rho_{cx}v_{q*})$ deviates from 1, the relative performance of the indexed quantity policy worsens and the expression in outer parentheses of (23) tends to $(b_2/c_2)^2-1+\rho_{cx}^2\left(1-(1/\rho_{cx}v_{q*})^2\right)$ for values of $v_x\to\infty$ (treating $\rho_{cx}v_{q*}$ as fixed) and equals $(b_2/c_2)^2-1$ when $v_x=0$. Whether indexed quantities prevail over prices depends on the degree of correlation between the index and cost shock and/or relative steepness of marginal benefits, as well as on whether $v_x/(\rho_{cx}v_{q*})$ is sufficiently close to 1.

2.5. Summary of the relative advantages of alternate policies

We can summarize the relative advantages of indexed quantities, prices, and fixed quantities in a two-dimensional space. The space is defined by the squared ratio of marginal benefit /cost slopes along the *y*-axis, and the expression $v_x/(\rho_{cx}v_{q*})$, measuring how closely indexed quantities match the GIQ policy, along the *x*-axis. Figure 3 shows three relations plotted in this space (with dashed lines), with each relation separating the space into two regions where one or another of two policies is preferred. Together the relations distinguish six regions where different policy rankings among the three policy options occur. Given five parameters—the marginal cost and benefit slopes, the coefficients of variation for the index and ex post optimal quantity, and the correlation between the latter two—one can identify a point in Figure 3 and determine the relative ranking of policies.

First, above the horizontal line where $(b_2/c_2)^2=1$ fixed quantities are preferred to prices and below it prices are preferred to fixed quantities—this is the Weitzman (1974) result. Second, note the vertical line at $v_x/(\rho_{cx}v_{q*})=2/\left(1-(\rho_{cx}v_{q*})^2\right)\approx 2$ based on (22). For cases (to the left of the vertical line) where $v_x/(\rho_{cx}v_{q*}) < 2/\left(1-(\rho_{cx}v_{q*})^2\right)$ indexed quantities are preferred to fixed quantities; for the reverse, fixed quantities are preferred. The rough intuition for the fixed versus indexed quantity result is that indexed quantities are an improvement unless they adjust by more than about twice the desired amount conditional on *x*—with the desired amount arising where $v_x/(\rho_{cx}v_{q*}) = 1$ (i.e., where indexed quantities replicate the GIQ policy).

Third, the nearly parabolic function defined by (23) describes the boundary between a preference for prices over indexed quantities (below the curve) and a preference for indexed quantities over prices (above the curve). Here, losses relative to the first-best outcome under the price policy depend on the distance from the *x*-axis—the ratio $(b_2/c_2)^2$—with relatively steep marginal benefits disfavoring prices. Meanwhile, losses under indexed quantities depend on a more complex relationship involving both the correlation between index and the cost shock (ρ_{cx}) and the difference between $v_x/(\rho_{cx}v_{q*})$ and its most favorable value of 1—with high values of ρ_{cx}^2 and values of $v_x/(\rho_{cx}v_{q*})$ close to 1 favoring indexed quantities. The locus of points where these losses are equivalent, and where prices and indexed quantities generate the same expected net benefits, defines the parabola-like function shown in the figure, passing through the points $(0,1)$, $(1,1-\rho_{cx}^2)$, and (approximately) $(2,1)$. Inside this parabola, $(b_2/c_2)^2$ is sufficiently large and $v_x/(\rho_{cx}v_{q*})$ is sufficiently close to 1 to favor indexed quantities. Note that the effect of ρ_{cx} on the policy comparison arises from both its scaling of $v_x/(\rho_{cx}v_{q*})$ and its movement of the minimum of the parabola at $1-\rho_{cx}^2$. Based on (6) and particularly (18), however, we know that the unambiguous effect of higher values of ρ_{cx} (other things equal) is to tilt preferences towards indexed quantities.

Note that for the GIQ policy, we can look along a vertical line where $v_x / (\rho_{cx} v_{q*}) = 1$ to determine policy rankings. In that case, when $(b_2/c_2)^2 < 1 - \rho_{cx}^2$, we have prices preferred to general indexed quantities preferred to quantities. When $1 > (b_2/c_2)^2 > 1 - \rho_{cx}^2$, general indexed quantities are preferred to prices are preferred to quantities. Finally, when $(b_2/c_2)^2 > 1$, indexed quantities are preferred to quantities are preferred to prices. In no instance is the fixed quantity policy preferred to index policies when $v_x / (\rho_{cx} v_{q*}) = 1$ as indexed quantities match the performance of the GIQ policy, and the GIQ policy is always (weakly) preferred to fixed quantities.

These results differ from Quirion [14] in that he finds indexed quantities almost always rank between quantities and prices; that is, indexed quantities are hardly ever the best or worst policy in his model. However, we find that this is due to his specific model for marginal costs—namely, that marginal costs are a function of emissions relative to baseline emissions, where baseline emissions is also the index. If we linearize such a model around the optimal quantity policy q^*, mean index $\bar{x}$, and mean-zero cost shock we have

$$
\begin{aligned}
MC &= f(\theta_c, q/x) \\
&\approx E\left[f(\theta_c, q/x)\right] - \left.\frac{\partial f}{\partial \theta}\right|_{(0, q^*/\bar{x})} \theta_c + \left.\frac{\partial f}{\partial (q/x)}\right|_{(0, q^*/\bar{x})} \left(\frac{q}{x} - \frac{q^*}{\bar{x}}\right), \\
&\approx c_1 - \omega + c_2 (q - q^*) - \frac{c_2 q^*}{\bar{x}} (x - \bar{x})
\end{aligned}
\tag{24}
$$

where ω is the linearized cost shock and we have written the new model to preserve the previous interpretation of the c_1 and c_2 parameters as the expected marginal cost and slope of marginal costs with respect to emissions, respectively, as in (1).[12] Setting expected marginal costs equal to marginal benefits from (2) yields

$$q = q^* + \left(\frac{c_2 q^* / \bar{x}}{b_2 + c_2} \right)(x - \bar{x}),$$

revealing that the optimal GIQ policy is $r^{**} = (c_2 q^* / \bar{x}) / (b_2 + c_2)$ under this specification of the index. Comparing the latter expression to (12) we see that while the effect of the index on optimal emissions in our more general model is governed by its *empirical* relationship with emissions, the Quirion approach *assumes* a particular relationship between optimal emissions and the index—specifically that the index perfectly measures baseline emissions and that costs hinge on fractional reductions.

What does this imply? Given that $r^* \approx q^* / \bar{x}$ from (14) and assuming v_x is small, we have $v_x / (\rho_{cx} v_{q^*}) = r^* / r^{**} = 1 + b_2 / c_2$. From (22) and (23), this implies that indexed quantities are preferred to fixed quantities when $b_2 < c_2 \sqrt{1 + v_x^2}$ and indexed quantities are preferred to prices when $b_2 > c_2 \Big/ \sqrt{1 + \left(\rho_{cx}^2 / (1 - \rho_{cx}^2)\right)\left(v_x^2 / (1 + v_x^2)\right)}$. Again assuming v_x is small, the preceding two conditions reduce to $b_2 < c_2(1 + \varepsilon_1)$ and $b_2 > c_2(1 - \varepsilon_2)$, respectively, for some small ε. In other words, indexed quantities are always better than *either* prices *or* fixed quantities, but only

[12] Quirion assumes a more specific functional form; however, the critical assumption is simply that the relationship between marginal costs and emissions is in terms of emissions relative to baseline emissions, where baseline emissions is also the index. Another difference between our model and Quirion's is that he incorporates marginal cost uncertainty as a multiplicative factor on the entire marginal cost function (including the slope), rather than solely in the intercept. In (24) we have approximated any uncertainty in the marginal cost slope as the linear term ω, finding this assumption is not the key driver of differences in our results.

better than both prices and fixed quantities in some narrow space $c_2(1-\varepsilon_2) < b_2 < c_2(1+\varepsilon_1)$, which is exactly the Quirion result.

But how plausible are these restrictions? It certainly seems plausible that marginal costs could be linear in the ratio of controlled emissions to baseline emissions—climate policy studies often compare marginal costs from models with different baselines in this way [23]. However, baseline emissions are not observable, and therefore are not a practical index; instead most proposals suggest using aggregate (e.g., GDP) or sectoral (e.g., electricity) economic output as an index. Therefore, we have to worry about imperfect correlation between the index and the level of baseline emissions. When we turn to our climate policy application below, we find that indexed quantities turn out to be worse than both fixed quantities and prices in some cases, suggesting the restrictive model is not always appropriate. For other applications, we see no reason why it is implausible that indexed quantities should be preferred to both fixed quantities and prices, as is possible in our theoretical framework.

3. Applying the model: the case of climate policy

In order to demonstrate how the model can be applied, as well as to highlight the importance of our modeling assumptions vis-à-vis Quirion [14], we briefly consider a simplified representation of national climate policy targets. Based on the results in section 2, the necessary parameters for understanding the relative advantage of the price, quantity, and indexed quantity policies, given by equations (7), (22), and (23), are the marginal cost and benefit slopes (c_2 and b_2), the variance of the cost shock (σ_c^2), the coefficients of variation of the index and ex post optimal quantity (v_x and v_{q*}), and the correlation of the cost shocks and the index (ρ_{cx}).

Given evidence in the literature that marginal benefits for greenhouse-gas mitigation are much flatter than marginal costs [9], we make the simplifying assumption that the slope of marginal benefits of climate change mitigation is approximately zero over the relevant range of emissions reduction. This implies that prices dominate quantities for climate policy, whether or not the quantities are indexed. This does not render the question of indexed quantities for climate policy uninteresting, as most countries are practically more interested in quantity policies of some sort, compared to the relevant price policy (a carbon tax). Coupled with an assumption about the nature of the index, this dramatically simplifies the analysis.

In particular, we also assume that the index (output or a similar measure) is being used to predict baseline emissions, which horizontally shifts the marginal cost schedule and is otherwise unrelated to remaining cost shocks.[13] In other words,

$$\theta_c = c_2\theta_{\bar{q}} + \upsilon$$

in (1) where $\theta_{\bar{q}}$ is the deviation in baseline emissions, υ is the remaining unexplained cost shock, and they are uncorrelated. Expressing deviations in baseline emissions ($\theta_{\bar{q}}$) as a prediction based on the index yields

$$\theta_c = c_2\beta_{\bar{q}x}\left(x-\bar{x}\right)+\upsilon' = c_2\sigma_{\bar{q}x}\left(\sigma_{\bar{q}}/\sigma_x\right)\left(x-\bar{x}\right)+\upsilon',$$

where $\beta_{\bar{q}x}$ is the regression coefficient of baseline emissions on the index (then expanded into its definition in terms of correlations and variances) and $\upsilon' = \upsilon + c_2\theta_{\bar{q}} - c_2\beta_{\bar{q}x}\left(x-\bar{x}\right)$ now includes the residual error from the prediction of $\theta_{\bar{q}}$. With the assumption that *x* and υ are uncorrelated,

[13] In their focus on the level of abatement (baseline emission minus the target), Sue Wing et al. [19] implicitly make this assumption. A more thorough application would need to consider a potentially more flexible relationship between the index and marginal costs (e.g., using data from cost modeling exercises), but this is unnecessary to provide a general sense of the results.

we have $\rho_{cx} = c_2 \rho_{\bar{q}x} (\sigma_{\bar{q}} / \sigma_x)(\sigma_x / \sigma_c)$. From (19) we have $v_{q*} = \sigma_c / (c_2 q^*)$ (making use of the assumption that $b_2 = 0$). This yields

$$\text{(25)} \qquad \frac{v_x}{\rho_{cx} v_{q*}} = \left(\frac{v_x}{\rho_{\bar{q}x} v_{\bar{q}}} \right) \left(\frac{q^*}{\bar{q}} \right).$$

Recalling the approximate condition that indexed quantities are preferred to fixed quantities if $v_x / (\rho_{cx} v_{q*}) < 2$, expression (25) indicates that under these simplifying assumptions this preference depends only on the coefficients of variation in baseline emissions and the index, their correlation, and the ratio of expected optimal to baseline emissions.[14]

Using country-level data on historic variation in carbon dioxide emissions and GDP (as the index), we estimate the following ranges for these variables across 19 high-emitting countries: $0.02 < v_{\bar{q}} < 0.07$, $0.01 < v_x < 0.07$, and $0.01 < \rho_{\bar{q}x} < 0.74$ (for details see [10]). As shown in Table 1, combining these values for each country yields estimates of $v_x / (\rho_{\bar{q}x} v_{\bar{q}})$ ranging from 0.9 to 4.0. Depending on the stringency of emissions reduction ($q^* / \bar{q}$), these values clearly indicate that for some countries indexed quantities will have higher expected net benefits while for other countries fixed quantities clearly dominate. For example, for a country with $v_x / (\rho_{\bar{q}x} v_{\bar{q}}) = 4$, indexed quantities will rank higher for emission reductions up to 50 percent, while for a country with $v_x / (\rho_{\bar{q}x} v_{\bar{q}}) < 2$, fixed quantities will dominate even for very

14 Pizer [13] compared the coefficient of variation in intensity and emissions to provide a crude argument in favor of targets based on one or the other. Such a comparison is equivalent to the condition in (25) only if emissions and output have roughly the same coefficient of variation.

small rates of control. These results are consistent with Sue Wing et al. [19], but contrast with Quirion's results for the reasons described above.

4. Conclusion

The relevance and importance of instrument choice for policy design never has been greater, particularly in the realm of environmental and energy policy. With the increasing acceptance in policy circles of market-based instruments, especially tradable permits, attention has turned to the more subtle design elements of these instruments and how they might be refined. In addition, interest has risen in the properties of more traditional instruments, such as performance standards, when flexibility is introduced through trading (e.g., renewable fuel standards and potential CAFE reforms in the United States). This interest is particularly intense in the realm of climate change, both in relation to the form that national commitments might take within an international framework and in the design of domestic implementing policies.

Our paper contributes to this debate by clarifying analytically how uncertainty in the costs of meeting particular policy targets might or might not be ameliorated by indexing fixed quantity policies to variables such as economic output. We find that the advantage of such indexing depends on a tradeoff between the introduction of an additional source of uncertainty—which lowers expected net benefits—and the benefit-raising effect of adjusting the policy target ex post thanks to correlation of the index with the object of regulation. For typical cases where uncertainty is relatively small (variation in the ex post optimum of less than about 10 percent of its mean), the preference for indexed over fixed quantities reduces to a question of whether the ratio of coefficients of variation of the index and the ex post optimal quantity, divided by their correlation, is less than approximately two. This fundamentally is an empirical question for ordinary indexed quantity policies, where the quantity is strictly proportional to the index. A general indexed quantity policy, however, allows separate setting of the mean quantity level and rate of adjustment to the index, and such a policy will always dominate a fixed quantity policy from the perspective of maximizing expected net benefits. Comparisons to a price policy are more complex and involve the ratio of the slopes of marginal benefits and costs.

These conclusions are subject to the caveat that we have chosen a deliberately simple model to focus on what we believe to be one of the most important elements of the instrument choice question, namely cost uncertainty. We have abstracted from other relevant concerns, including the potential for an indexed quantity policy to create undesirable incentives if firms perceive that they can gain additional emissions rights by increasing their output.[15] While we do not think this is a concern for national-level policies, it could be for indexed policies at the sectoral or product level. We also have not addressed the fact that quantities indexed to output, even if they reduce overall expected costs, may lead to worse outcomes when output is low and better outcomes when output is high. This type of pro-cyclical behavior may be undesirable from a macroeconomic perspective, although we suspect this concern is not large.[16]

Applying these conceptual insights and analytic formulae to the case of climate change policy across the biggest international emitters of CO_2, we find (consistent with previous literature) that while prices (i.e., carbon taxes) dominate both fixed and indexed quantities from an efficiency perspective, indexing quantities to economic output could yield higher expected net benefits than fixed quantity policies for some but not all countries. Applying this framework to

---

[15] This is the typical form of a performance standard, such as the U.S. lead phase-down in gasoline. Alternatively, emissions rights could be increased for all firms based on aggregate output, diluting the effect. See Fischer [4] for a discussion of these issues.

[16] This point is discussed in Ellerman and Sue Wing [3].

more in-depth empirical analyses of specific environmental, resource, and other policies would be an interesting focus of further research.

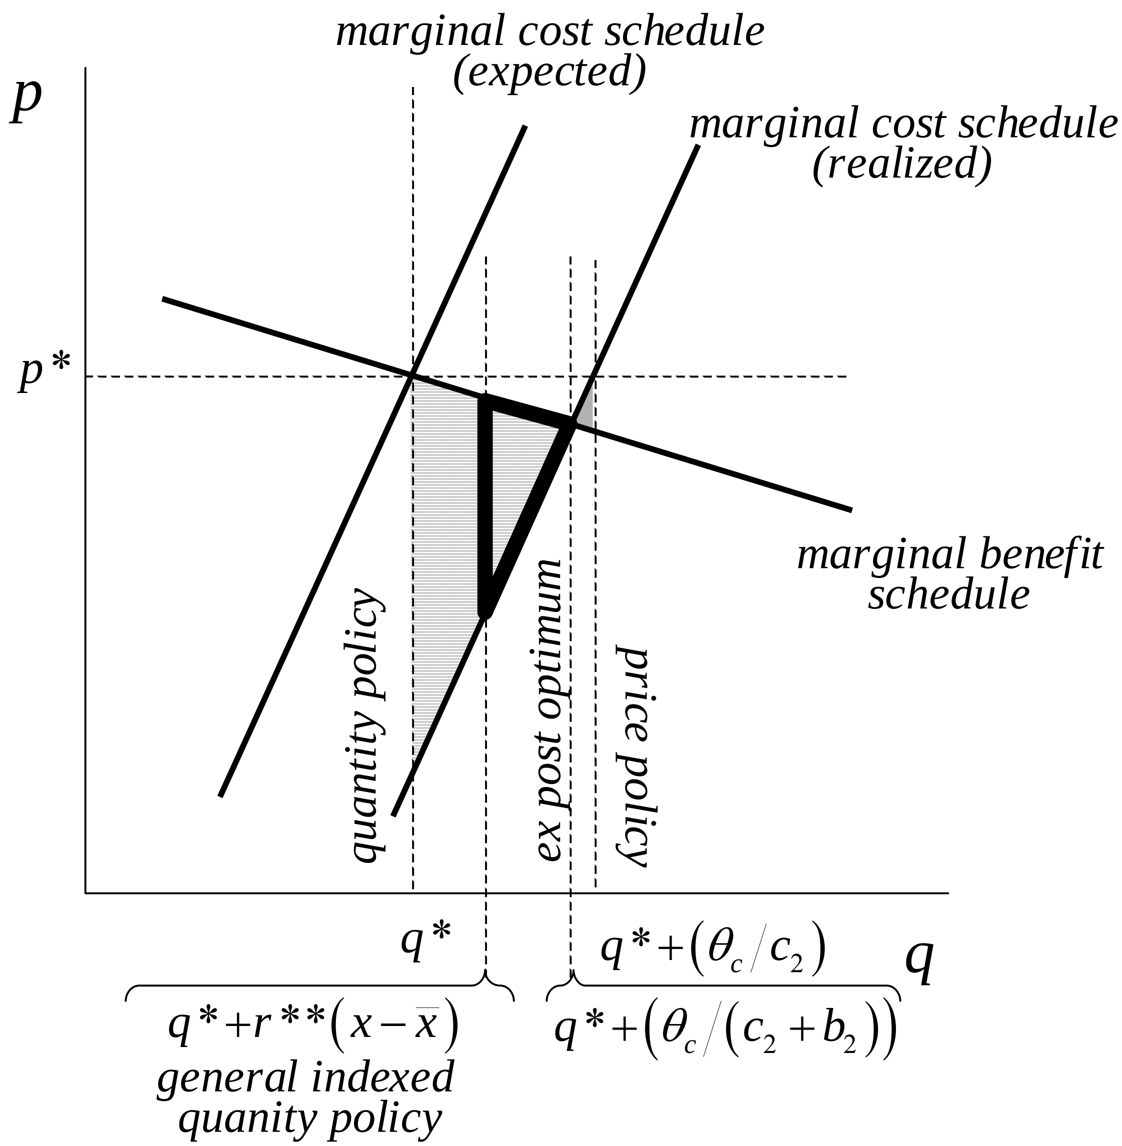

Figure 1. Welfare losses under quantity, price, and general indexed quantity policies (quantity loss is hatched; price loss is shaded; indexed quantity loss thickly outlined)

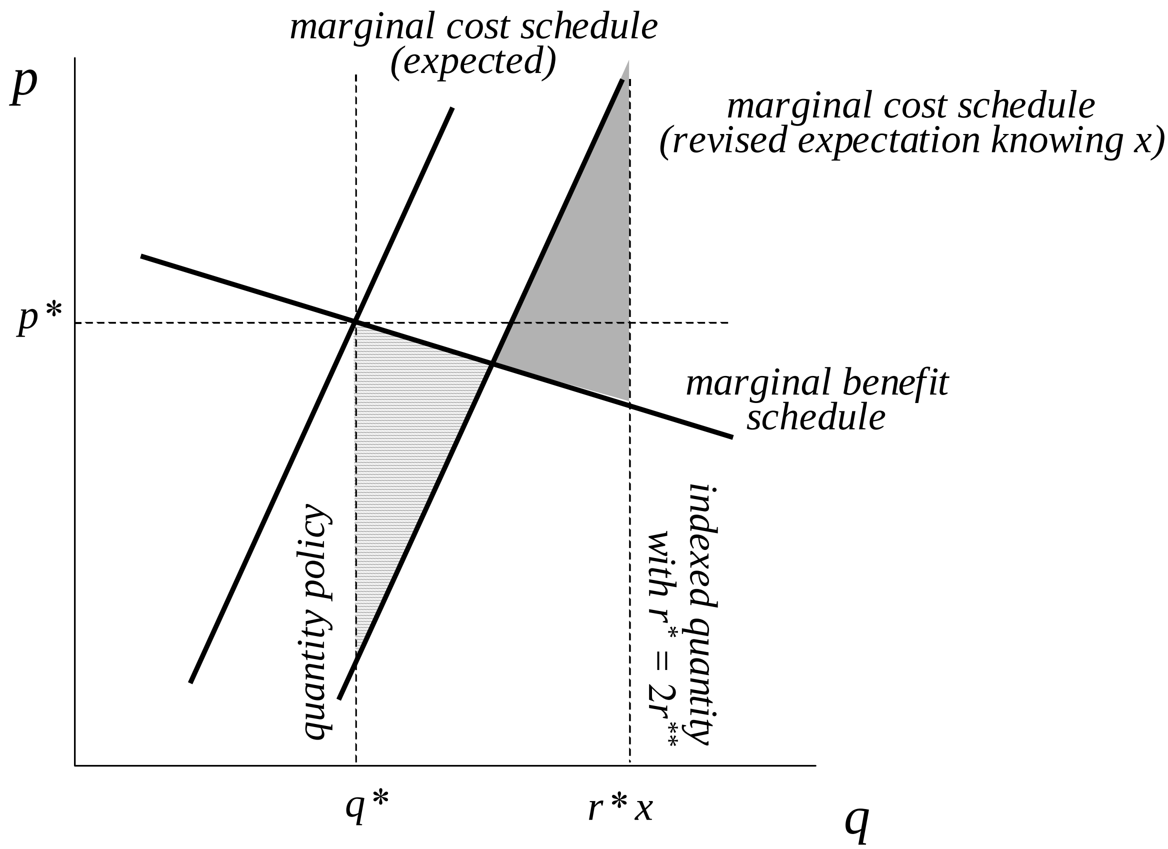

Figure 2. Point of indifference between quantity and (ordinary) indexed quantity policies (quantity loss is hatched; indexed quantity loss is shaded)

Note: The figure shows how, under an assumption of linear marginal costs, the gains from ordinary indexing relative to fixed quantities is exactly offset by the losses at the point where the ordinary indexed quantity adjusts by twice the optimal amount.

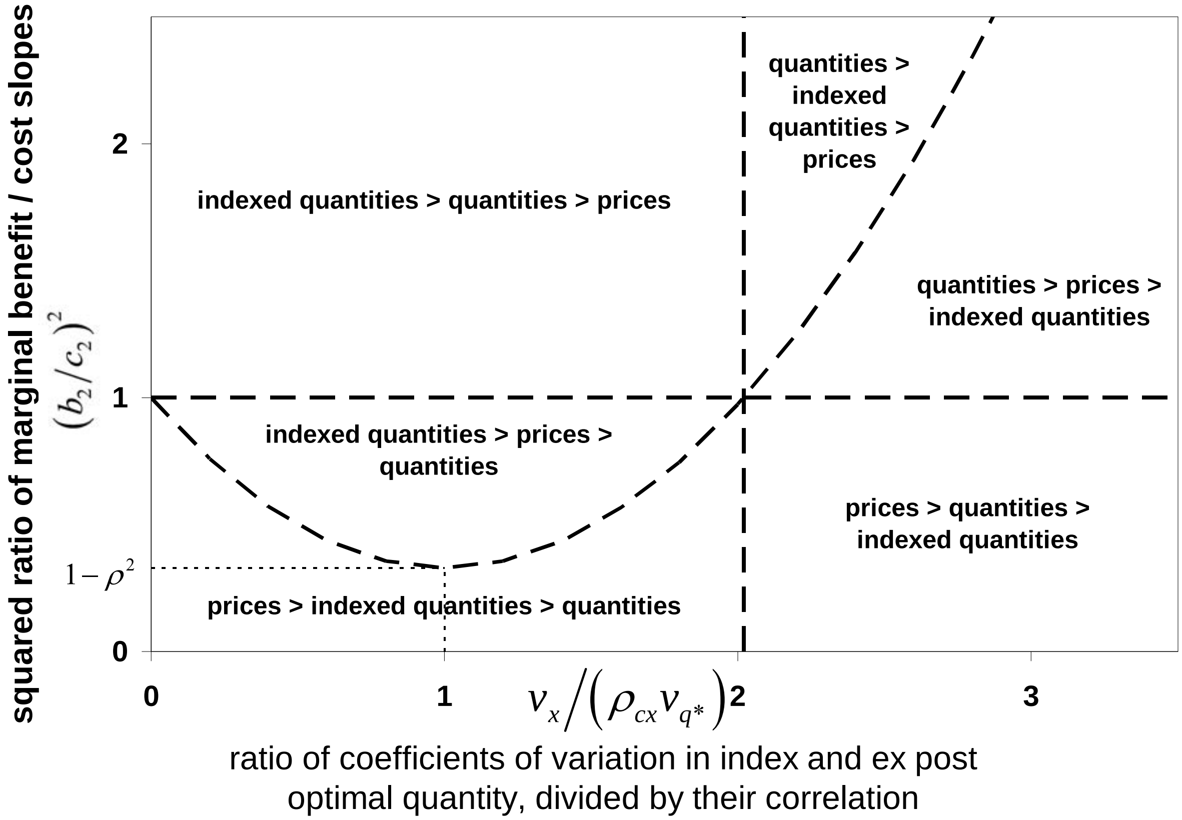

Figure 3. Regions of relative advantage for indexed quantities, prices, and quantities

Table 1. Application to climate change policy

Country	$\frac{v_x}{\rho_{\bar{q}x} v_{\bar{q}}}$
Australia	3.1
Brazil	1.0
Canada	1.2
China	1.4
France	3.9
India	2.6
Indonesia	2.5
Iran	2.5
Italy	1.3
Japan	4.0
Korea (South)	0.9
Mexico	1.3
Netherlands	1.3
Poland	2.4
Saudi Arabia	2.4
South Africa	3.0
Spain	1.3
United Kingdom	2.7
United States	1.1
U.S. Electricity	0.9
World	1.4

Note: Under the simplifying assumptions described in section 3 of the text, indexed quantities are preferred to fixed quantities for CO_2 mitigation if $v_x / (\rho_{\bar{q}x} v_{\bar{q}})$ multiplied by the ratio of optimal to baseline emissions ($q^* / \bar{q}$) is approximately less than 2.

References

[1] K.K. Aase, A pricing model for quantity contracts. Journal of Risk and Insurance 71 (2004) 617-642.

[2] C. Carlson, D. Burtraw, M. Cropper, and K.L. Palmer, SO2 Control by Electric Utilities: What are the Gains from Trade. Journal of Political Economy 108 (2000) 1292-1326.

[3] A.D. Ellerman, and I.S. Wing, Absolute versus Intensity-Based Emission Caps. Climate Policy 3 (2003) S7-S20.

[4] C. Fischer, Combining rate-based and cap-and-trade emissions policies. Climate Policy (2003) S89-S109.

[5] G.E. Helfand, Standards versus Standards: The Effects of Different Pollution Restrictions. American Economic Review 81 (1991) 622-634.

[6] F. Jotzo, and J.C.V. Pezzey, Optimal intensity targets for greenhouse gas emissions trading under uncertainty. Environmental and Resource Economics 38 (2007) 259-284.

[7] S. Kerr, and R.G. Newell, Policy-Induced Technology Adoption: Evidence from the U.S. Lead Phasedown. Journal of Industrial Economics 51 (2003) 317-343.

[8] C.-L. Li, and P. Kouvelis, Flexible and risk-sharing supply contracts under price uncertainty. Management Science 45 (1999) 1378-1398.

[9] R.G. Newell, and W.A. Pizer, Regulating Stock Externalities Under Uncertainty. Journal of Environmental Economics and Management 45 (2003) 416-432.

[10] R.G. Newell, and W.A. Pizer, Indexed regulation, RFF Discussion Paper 06-32, Washington, DC, 2006.

[11] A.L. Nichols, Lead in Gasoline. in: R.D. Morgenstern, (Ed.), Economic Analyses at EPA: Assessing Regulatory Impact, Resources for the Future, Washington, DC, 1997, pp. 49-86.

[12] W.A. Pizer, Combining price and quantity controls to mitigate global climate change. Journal of Public Economics 85 (2002) 409-434.

[13] W.A. Pizer, The case for intensity targets. Climate Policy 5 (2005) 455-462.

[14] P. Quirion, Does uncertainty justify intensity emission caps. Resource and Energy Economics 27 (2005) 343-353.

[15] M.J. Roberts, and M. Spence, Effluent Charges and Licenses Under Uncertainty. Journal of Public Economics 5 (1976) 193-208.

[16] C.S. Russell, W. Harrington, and W.J. Vaughan, Enforcing Pollution Control Laws, Resources for the Future, Washington, DC, 1986.

[17] R. Stavins, What Can We Learn from the Grand Policy Experiment? Positive and Normative Lessons from SO2 Allowance Trading. Journal of Economic Perspectives 12 (1998) 69-88.

[18] R.N. Stavins, Correlated Uncertainty and Policy Instrument Choice. Journal of Environmental Economics and Management 30 (1996) 218-232.

[19] I. Sue Wing, A.D. Ellerman, and J. Song, Absolute vs. Intensity Limits for CO_2 Emission Control: Performance under Uncertainty. in: H. Tulkens, and R. Guesnerie, (Eds.), The Design of Climate Policy, MIT Press, Cambridge, in press.

[20] L.E.O. Svensson, What is wrong with Taylor rules? Using judgement in monetary policy through targeting rules. Journal of Economic Literature XLI (2003) 426-477.

[21] M.L. Weitzman, Prices vs. Quantities. Review of Economic Studies 41 (1974) 477-491.

[22] M.L. Weitzman, Optimal Rewards for Economic Regulation. American Economic Review 68 (1978) 683-691.

[23] J.P. Weyant, and J. Hill, The Costs of the Kyoto Protocol: A Multi-Model Evaluation, Introduction and Overview. Energy Journal Special Issue (1999).

[24] G.W. Yohe, Towards a General Comparison of Price Controls and Quantity Controls under Uncertainty. The Review of Economic Studies 45 (1978) 229-38.