Absolute vs. Intensity Limits for CO_2 Emission Control: Performance Under Uncertainty

Ian Sue Wing*
Center for Energy & Environmental Studies
Dept. of Geography & Environment, Boston University
Joint Program on the Science & Policy of Global Change, MIT

A. Denny Ellerman and Jaemin Song
Joint Program on the Science & Policy of Global Change
Center for Energy & Environmental Policy Research, MIT

Abstract

We elucidate the differences between restricting CO_2 emissions via absolute limits and intensity-based caps which are indexed to GDP. We demonstrate that the two instruments are identical under certainty, rigorously establish the conditions under which each generates less variance in abatement and compliance costs under uncertainty, and derive the optimal degree of indexation for a "hybrid" instrument which combines an absolute and an intensity cap. Empirical tests of the resulting conditions using both historical data and past forecasts indicate a general preference for intensity-based—and particularly hybrid—caps. But they also make clear that the conditions favoring the latter instruments are far from universal, and that a country might even shift back and forth between absolute and intensity limits as its circumstances change.

Keywords: climate change policy, emission caps, intensity targets, dynamic targets

JEL Classification: Q540, Q580

---

* Corresponding author. Address: Rm. 461, 675 Commonwealth Ave., Boston MA 02215. Email: isw@bu.edu. Phone: (617) 353-5741. Fax: (617) 353-5986. ISW was supported by U.S. Department of Energy Office of Science (BER) Grants Nos. DE-FG02-02ER63484 and DE-FG02-06ER64204. ADE and JMS were supported by the MIT Joint Program on the Science and Policy of Global Change, funded through a government-industry partnership including U.S. Department of Energy Office of Science (BER) Grant No. DE-FG02-94ER61937, U.S. Environmental Protection Agency Cooperative Agreement No. XA-83042801-0, and a group of corporate sponsors from the U.S. and other countries. This paper has benefited from suggestions by two anonymous referees, as well as helpful comments by Jake Jacoby, John Reilly, John Parsons, Dick Eckaus, Philippe Quirion, and participants at the Cambridge-MIT Electricity Policy Forum Spring Research Seminar, the EPRI Global Climate Change Research Seminar, and the David Bradford Memorial Conference on the Design of Climate Policy. Errors and omissions are solely the authors' responsibility.

“Prediction is very hard... particularly of the future...”

—Niels Bohr

1. Introduction

This paper addresses an environmental policy question which is simple and fundamental: if a country makes a commitment now to constrain emissions at or below a target level in some future time-period, does it make a difference if the commitment is expressed as a limit on the absolute level of emissions or on the intensity of emissions? Nowhere is this question more relevant than the design of policies to mitigate the emissions of greenhouse gases (GHGs), because of the widespread concern that attempts to cut GHG emissions will cause significant increases in energy prices and reductions in economic output and welfare. The GHG emission limits negotiated under the Kyoto Protocol have been criticized, for example, as contributing to this unfavorable outcome because they are expressed as fixed caps on countries’ ability to emit. The absolute character of these caps, it has been argued, fails to account for the possibility that economies and their emissions might grow more quickly than was expected at the time the targets were negotiated, and would therefore inflict larger-than-anticipated economic losses on the Kyoto signatories.

Several generic proposals have been advanced in response to these concerns. A “safety valve” would set an upper bound on the marginal costs of abatement and thereby truncate the upper end of the distribution of outcomes (Kopp et al 2000; Jacoby and Ellerman 2002; Philibert 2005). We do not engage in further discussion of these proposals here, but focus instead on another alternative: intensity limits. Although rare in the domain of GHG emissions control, limits on the pollution intensity of output are by far the more common method of constraining emissions in the field of environmental regulation.[1] Nevertheless, "relative" or intensity-based targets have been adopted as a component of climate policy in the U.K. Emissions Trading Scheme (U.K. DEFRA 2001),[2] and in 2001 the Bush administration proposed a voluntary target to reduce the GHG emission intensity of the U.S. economy by 18 percent by 2012.

Implicit in the adoption of all these measures is the recognition of the general principle that the pollution from a source can be limited by specifying either an absolute cap on the quantity of emissions that it generates or by setting a maximum allowable intensity of emissions relative to some measure of output or input. Examples are the units of output or the amount of energy input required by some production process at the firm level, and the volume or value of commodities purchased by consumers at the level of an economic sector, or even GDP at the national level. Such an intensity limit can be imposed either directly, as an emission rate limit or an efficiency standard, or indirectly, by means of technology mandates that have the same effect.

In this paper we elucidate the differences between absolute and intensity limits under uncertainty. Our guiding assumption is that the variance in the intended environmental and economic effects of an emission constraint is an important consideration in deciding how it is to be implemented. For simplicity, our discussion will focus on the economic costs of this choice, without consideration of the associated environmental benefits.[3] As well, we restrict our analysis

---

[1] Familiar examples of intensity limits are the emission rate limits imposed on nearly all sources under State Implementation Plans in the US, best available control technology mandates, such as in the U.S. New Source Performance Standards or the EU Large Combustion Plant Directive, and the Corporate Average Fuel Economy standards in the US and similar programs in Europe. Although many of the latter do not explicitly specify an emission rate, the effect of these programs is to reduce emission (or energy) intensity and to allow emissions to vary with the level of output. However, absolute emission caps can also be found in several programs controlling conventional pollutants, for example the SO_2 trading (acid rain), RECLAIM, and the Northeastern NO_x Budget programs in the U.S. Rosenzweig and Varilek (2003) review experience with these and other rate-based emission regulations.

[2] The UK Emissions Trading Scheme is unique in having two sectors, an absolute sector containing firms with absolute limits on GHG emissions and a relative sector containing firms with intensity limits, and allowing trading (with some restrictions) between the two sectors.

[3] The latter has been the focus of studies by Quirion (2005), Pezzey and Jotzo (2005) and Newell and Pizer (2006), who use a cost-benefit framework to analyze the performance of intensity limits relative to other instruments.

to the setting of a single-economy, and leave the interaction of absolute and intensity limits in an international emission trading system to future research. Also left to future investigation is a rigorous comparison of the relative merits of intensity limits safety valves, or intertemporal banking and borrowing as means to reduce the variance of outcomes. We concentrate instead on laying the groundwork for such assessments by elaborating the conceptual and theoretical foundation introduced in Ellerman and Sue Wing (2003).

These sacrifices in terms of scope allow us to make three contributions. First, we demonstrate that an emissions constraint can be expressed equivalently as an absolute or intensity limit on emissions when there is no uncertainty about the future. On its face, this may appear to be a trivial point, but there seems to be much misunderstanding in policy circles on this issue, and we are hard-pressed to find analyses that rigorously establish this basic fact.

Second, we demonstrate rigorously the conditions under which an absolute or indexed intensity limit would be preferred on the basis of reduced variance and we also discuss the characteristics of an optimal degree of indexing where an intensity limit would produce less variance.

Third, we explore the policy implications of these conditions using time series data on nations’ actual CO_2 emissions and GDP, as well as historical forecasts of these variables. We do this by conducting a backcasting analysis that considers an alternate state of the world in which countries decided to limit their emissions of CO_2 in earlier decades, which allows us to investigate what would have been the optimal choice for the form of their emission cap.

The plan of paper is as follows. Section 2 provides a review of the literature on intensity limits. In section 3 we build on the arguments in Ellerman and Sue Wing (2003) to demonstrate that emission limits based on the emission intensity of GDP are equivalent with intensity limits in a world in which there is no certainty. We derive the main theoretical results of the paper in section 4, where we establish the conditions for the relative attractiveness of fixed versus flexible instruments and indicate the conditions for optimal indexing. In section 5 we present and discuss the results of the empirical simulations and analysis. Section 6 concludes with a summary of the implications for climate policy and the design of emission trading programs.

2. Literature Review

A recent and diverse literature has developed concerning the use of intensity-based and indexed caps in the context of climate policy.[4] The nearly uniform motivation is the widespread perception that developing countries would not accept absolute caps because of the perceived limit on economic growth. The proposal by Argentina in November 1999 at the 5th Conference of the Parties to the Kyoto Protocol first drew official attention to this subject (Argentina 1999; Barros and Conte Grand 2002). Shortly thereafter, one of President Clinton's economic advisors proposed indexing GHG emission targets to GDP growth as a means of making Kyoto-type caps more acceptable to developing countries (Frankel 1999). Key early papers by Baumert et al (1999) and especially Lutter (2000) introduced the idea of intensity targets as a hedge against uncertainty—in particular their potential to mitigate excess abatement costs incurred by higher-than-expected business-as-usual (BAU) emissions. Subsequently, the Bush administration's announcement of a target to reduce U.S. GHG intensity by 18 percent by 2012 (White House 2002), and its advocacy of intensity limits for developing countries prompted a spate of analyses

---

[4] We use the terms intensity-based and indexed as virtually interchangeable. A conventional intensity limit is automatically indexed to whatever is the denominator by which the intensity is stated. By the same token, an otherwise fixed absolute cap can be indexed to vary the level of allowed emissions according to movements of some denominated quantity, such as output or GDP.

concerned both with the adequacy of the U.S. target and the more general merits of intensity-based caps. [5]

While analysts appear united in finding that the target set by the Bush administration is indistinguishable from BAU emissions, opinion on the attractiveness of intensity limits is less uniform. Gielen, Koutstaal and Vollebergh (2002) and Fischer (2003) draw on an old literature in environmental economics, going back to Spulber (1985) and Helfand (1991), which criticizes intensity limits because of the incentive they give producers to use larger quantities of the input or output in which the intensity index is denominated. Compared to absolute limits, intensity caps are a “subsidy” to firms’ use of the denominated input or to their production of the denominated output, thereby giving rise to an inefficient allocation of resources.

The output subsidy critique of intensity limits applies only in so far as the limit is faced by individual producers. It does not apply at the country level because the indexation variable, aggregate output, does not figure in firm-level decisions. Within a nation, producers could be expected to take into account the fact that an indexed emissions cap would be adjusted upward (or downward) but the practical incentive they would face is a lower (or higher) cost for the use of allowances than would otherwise be the case. Individual firms would not face any greater or lesser constraint as a result of variations in the output of, or the inputs to, their production processes.

---

[5] The stated intent of the Bush Administration’s espousal of intensity targets is to take future economic growth into account: “This new approach focuses on reducing the growth of GHG emissions, *while sustaining the economic growth* needed to finance investment in new, clean energy technologies.” (White House 2002) [our emphasis]. The stringency of the Bush target (or lack thereof) is a legitimate concern. The 18 percent reduction in the GHG emission intensity of the U.S. economy by 2012 is to be compared with the contemporary DOE/EIA (2004) forecast that projects a decline in the CO_2-GDP ratio of 15 percent by 2010. By contrast, the reduction in the CO_2 emission *intensity* over the same period implied by the U.S. Kyoto target is greater than 40 percent. Moreover, the Bush target is specified not as a legally-binding limit, but as a goal to be achieved through an array of voluntary actions, creating the potential for little or no abatement to take place.

Another persistent critique of intensity limits is that, relative to an absolute ceiling on emissions that is fixed ex ante, an intensity cap creates the potential for an environmentally adverse outcome if GDP is higher than expected since the cap would be adjusted upward, making the target less stringent in absolute terms (cf. Dudek and Golub 2003). Comparisons between Kyoto's absolute emission targets and intensity limits that characterize the latter as economically advantageous while being environmentally disadvantageous reflect this criticism. What happens if GDP growth declines or is *less* vigorous than expected is, however, rarely noted. The level of an intensity cap will adjust downward, making the target *more* stringent than an unchanging absolute cap. In this case, an intensity cap is environmentally advantageous and economically disadvantageous. Ellerman and Sue Wing (2003) argue intuitively that an intensity limit trades off less stringent control of emissions in a state of the world with higher than expected economic growth for more stringent control in a state of the world with lower than expected growth. Mirroring this ex post divergence in stringency and environmental outcome will be ex post divergence in the quantity and cost of abatement. For this reason, the presumption that intensity-based limits are inherently less stringent is wrong.

This presumption is rife among the negative reactions to the Bush climate change plan, which uniformly argue that indexing future emission constraints to GDP would allow GHG emissions to continue to rise when GDP is increasing, as it is generally expected to do.[6] Such criticism belies confusion of the *stringency* of the target with the *form* of the instrument employed in its execution. In spite of the fact that these are two separate issues, the unstated implication appears to be that intensity limits allow emissions to continue growing unabated, while absolute caps do not. The flaw in this argument is that it ignores the counterfactual no-policy path of emissions which can as easily be higher than that of an indexed cap, which would

---

[6] See e.g. "Blowing Smoke", The Economist, Feb. 14 2002, p. 27.

allow emissions to increase over time, as it can be higher than an absolute cap, as shown by the experience of Russia and the East European countries under the Kyoto Protocol. The former case would produce real reductions despite being an intensity limit and the latter would impose no constraint despite being an absolute limit. A nation's decision to set an absolute cap on emissions is invariably informed by a sense of the limit's expected effects, which typically incorporate a forecast of GDP in the future period when that instrument is slated to enter into force. Given this set of expectations, there are numerous schemes for specifying GDP-indexed emission targets that are entirely equivalent to the absolute limit, a point which we demonstrate in section 3 and appendix A.

The essential caveat to this equivalence is uncertainty about the future. Of principal concern is the ex-post level of the emission limit that results from imposing either instrument ex ante. Different instruments whose effects are predicted to be equivalent based on ex ante expectations of GDP may turn out not to hold if actual GDP in the target period diverges from its expected level. In particular, the level of an intensity-based cap will fluctuate in proportion to the ratio of actual to expected GDP.

A third critique of intensity limits can be found in Müller and Müller-Furstenberg (2003), who in addition to the preceding arguments cite problems of implementation in choosing appropriate indices and avoiding biases in these indices. While these concerns are legitimate they are also typical of many forms of indexing that are commonly accepted such as indexing wages and benefits, and more recently, inflation protected bonds.

On the other side of this debate, Baumert (1999), Lutter (2000), Kim and Baumert (2002), Strachan (forthcoming) and Kolstad (2005) all find merit in the concept of intensity-based caps because of the reduction in uncertainty in the economic outcome gained by indexing the cap to GDP and, crucially, the effect on the countries' willingness to participate in international agreements. Baumert (1999), Lutter (2000), and Lisowski (2002) also see intensity limits as a means to avoid the "hot air" resulting from overly generous absolute caps that might be needed to reassure acceding countries that the emissions limit would not place undue costs on them in the event of greater than expected economic growth. Along these lines, Jotzo and Pezzey (2005) provide a theoretical analysis and simulations of binding absolute and intensity caps in which parties are assumed to posses varying degrees of risk aversion to unexpectedly high-cost outcomes with particular attention given to developing country participation. They find that intensity-based caps are superior to absolute caps for circumstances where all parties to a treaty place some positive value on global abatement, face positive abatement costs, and are risk-averse in varying degrees to high cost outcomes. For individually-varying but positive valuations on global abatement, parties would be willing to embrace tighter binding targets in return for the removal of some of the uncertainty relating to high cost outcomes.

Our own contribution to this debate (Ellerman and Sue Wing 2003) treats absolute and intensity limits with equanimity while focusing squarely on the nature of the relevant uncertainties. Under conditions of certainty, equivalent absolute and intensity-based caps would have identical effects, and the outcomes between the two forms differ only to the extent that realized values for GDP or other indices diverge from expectation. Our aim in this paper is to develop the implications of uncertainty in baseline emissions and GDP for policy makers' choice between an absolute and an intensity cap. In particular, we rigorously establish the conditions under which one or the other form of emission limit will give rise to smaller variance in cost outcomes, and we test which form of the limit would have produced less variance in abatement cost using historical and forecast data.

All of the analyses focusing on the merits of intensity limits as a means of reducing uncertainty assume that emissions and GDP are positively correlated. Like Jotzo and Pezzey, we find that the positive correlation between emissions and GDP is often large enough that intensity caps reduce the variance of cost outcomes. However, we also demonstrate that this result has failed to hold for some countries over varying periods of time.

An important assumption in our analysis is that policy makers care about variance in outcomes. As we have stressed, if expectation were all that mattered the form of the limit can be treated with indifference so long as the limits being compared are ex ante equivalent. Accordingly, the concern is whether by choosing one form or the other, the policy maker can reduce or even minimize the expected variance in outcomes.

At least two different motivations can be offered for seeking to minimize variance, which can be characterized as preserving initial expectations and avoiding undue adjustment costs. Since policy makers will tend to set the level of an emission constraint based on their expectations of the economic and environmental conditions that will prevail when that target enters into force, they might seek a limit that would result in less deviation from the initially expected environmental and economic outcomes as a result of the inevitable changes that will occur over time. Also, in a non-putty-putty world in which investments cannot be instantly made and undone to insure optimal responses to the constraint, they would prefer the form of the limit that would reduce the adjustment costs that would be associated with over- or under-investment in abatement capability. Both of these motivations would lead to an interest in minimizing variance.

3. Absolute and Intensity Limits: Equivalence Under Certainty

Our first task is to establish the equivalence of absolute and intensity limits under certainty. Our analytical approach builds on Ellerman and Sue Wing (2003). We consider a nation which commits to limit its emissions, but is undecided whether to express this limit as a constraint on the absolute level of emissions or on the intensity of emissions indexed to GDP.

Let Q denote emissions, Y denote GDP, and γ the emission intensity of the economy:

(1) $$\gamma = Q / Y .$$

Suppose that the country chooses to limit its emissions to an absolute level, $\underline{Q}$. We assume that this decision is made conditional on an initial information set, θ, which we represent using the conditional expectation operator, E_θ. With expected baseline emissions $E_\theta[Q^{BAU}]$, if the country chooses a binding absolute cap on emissions, the level of abatement (A^A) is, in expectation:

(2) $$E_\theta[A^A] = E_\theta[Q^{BAU}] - \underline{Q} > 0 .$$

Eq. (1) implies that this fixed limit can be transformed into an emission intensity cap according to the expectation of GDP, $E_\theta[Y]$. If the emission target is expected to bind, there exists a corresponding ceiling on emission intensity:

(3) $$\underline{\gamma} = \underline{Q} / E_\theta[Y] < E_\theta[\gamma^{BAU}] .$$

Therefore, under stable expectations, the expected level of abatement with an intensity limit (A^I), is the same as in (2), i.e.,

(4) $$E_\theta[A^I] = E_\theta[Q^{BAU}] - \underline{\gamma} E_\theta[Y] = E_\theta[A^A] .$$

The condition expressed by equation (4) expresses what we later refer to as ex ante equivalence. Thus, given an abatement cost schedule $C(A)$ which we assume is positive, monotonic increasing and known with certainty, the expectation at time zero of the cost of reducing under either instrument is $C\left(E_\theta[A^A]\right) = C\left(E_\theta[A^I]\right)$. The policy maker would be indifferent between the two and the form of the emission limit would be irrelevant, if expectation were all that mattered.

4. The Choice Between Absolute and Intensity Limits Under Uncertainty

In keeping with the motivation of this paper, we imagine a policy maker who is concerned about variance in abatement and cost outcomes. The actual levels of abatement and cost under the two forms would correspond to:

(5) $A^A = Q^{BAU} - \underline{Q}$, and

(6) $A^I = Q^{BAU} - \underline{\gamma} Y$.

Since it will be generally true that $E_\theta[Q] \neq Q$, and $E_\theta[Y] \neq Y$, different levels of abatement and cost will be associated with the two limits. Since the emission target expressed by the intensity limit adjusts to changes in GDP, whereas that determined by absolute cap does not, the difference in actual abatement will be:

(7) $A^I - A^A = \underline{\gamma} Y - \underline{Q}$.

Any rational policy maker would know that things will change. Not knowing the future changes, he or she might well ask whether some form of the emissions constraint might reduce variance in outcomes so that the actual outcomes not deviate too much from the initial set of expectations.

To evaluate variance more formally, we use the "hybrid" GDP-indexed limit introduced by Ellerman and Sue Wing (2003), which specifies the indexed cap on emissions, $\tilde{Q}$, as the convex combination of a fixed cap and a pure intensity target:

(8) $\tilde{Q} = (1-\eta)\underline{Q} + \eta \underline{\gamma} Y$.

The form of the emission limit combines an absolute limit with an intensity target specified by the product of the intensity limit in eq. (3) and actual GDP. The coefficient $\eta \in [0,1]$ is an indexation parameter which represents the degree to which the limit accommodates changes in GDP from its expected level, and it is under the policy maker's control. When $\eta = 0$ the limit is absolute, and when $\eta = 1$ it is a pure intensity limit that adjusts fully to the change in GDP. The result is a more general form of eq. (6):

$$(6') \qquad \tilde{A}^I = Q^{BAU} - \tilde{Q} = Q^{BAU} - (1-\eta)\underline{Q} - \eta\,\underline{\gamma}Y .$$

In keeping with the result of the previous section, if $\underline{Q}$ and $\underline{\gamma}$ are set initially to be ex ante equivalent, such that $\underline{Q} = \underline{\gamma}E_\theta[Y]$, and it be further assumed that $E_\theta[Y] = Y$, then $\tilde{Q} = \underline{Q}$ regardless of the value of η. Further results, with different forms of the emissions limit, are provided in an appendix to the paper.

4.1. Indexed limits

We now establish the conditions under which an indexed limit will be preferred. Our criterion in making this determination is minimization of the variance in the cost of abatement. Given the monotone increasing character of the cost function C, it therefore suffices to demonstrate which instrument generates the smaller variance in abatement.

From eq. (5), the variance of abatement under the absolute cap is simply:

$$(9) \qquad \text{var}[A^A] = \text{var}[Q^{BAU}],$$

while (6′) implies that the variance of an indexed intensity cap ($\eta > 0$) is:

$$(10) \qquad \text{var}[\tilde{A}^I] = \text{var}[Q^{BAU}] + (\eta\underline{\gamma})^2\,\text{var}[Y] - 2\eta\,\underline{\gamma}\,\text{cov}\left[Q^{BAU}, Y\right].$$

The key question is whether the variance in the expected effect of the latter instrument is less than that of the former. This can be determined by subtracting (9) from (10) and rearranging. The variance of expected abatement and cost is smaller for the indexed intensity limit if

$$\eta \underline{\gamma} / 2 < \mathrm{cov}[Q^{BAU}, Y] / \mathrm{var}[Y].$$

The intuition behind this expression becomes clearer if we multiply both sides by $E_\theta[Y]$ / $E_\theta[Q^{BAU}]$ to express the target, covariance and variances in normalized form:

$$(11) \qquad \frac{\eta}{2}(\underline{Q} / \overline{Q}) < \rho v[Q] / v[Y] = Z,$$

where $\overline{Q} = E_\theta[Q^{BAU}]$, ρ is the correlation between BAU emissions and GDP, $v[Q]$ and $v[Y]$ are the coefficients of variation of baseline emissions and GDP, and $\underline{Q} / \overline{Q}$ expresses the ex ante equivalent absolute limit as a fraction of expected BAU emissions. The left-hand-side is the product of two important policy variables: the form of the limit, given by the value of indexation parameter, and its stringency, expressed as the ratio of the constrained emissions to expected BAU emissions. By contrast, the quantity on the right-hand side, Z, is a function solely of stochastic properties of the economy, none of which are subject to manipulation by the policy maker.

Eq. (11) is the main result of the paper and it provides a mathematical statement of the conditions under which an intensity limit indexed by η would result in less variation of outcomes. The implication of eq. (11) is that the conditions under which an indexed intensity limit better preserves initial expectations about the level of actual abatement and cost are more likely to obtain the higher the correlation between Q and Y and the greater the variation in Q relative to that in Y. For a given emission target, consider first the case of a fully indexed intensity cap ($\eta = 1$). Since the left-hand side of (11) is always positive, as are the coefficients of variation, a necessary condition for the intensity limit to exhibit less variation is that the correlation between emissions and GDP be positive. This is however not a sufficient condition. If the variation in emissions were very small relative to the variation in GDP, indexing to GDP would produce more variance than an absolute limit. Therefore, the sufficient condition is that either the degree of indexation or the level of the emission limit (or both) be small enough for the inequality to hold.

Also, if emissions and GDP are perfectly correlated and have similar degrees of variability, so that the right-hand side of (11) equals unity, then any indexed intensity limit will always exhibit less variability and be preferred since the left-hand-side will always be less than half of unity, or 0.5. For any value of $Z < 0.5$, it is possible that an absolute cap might generate less variability, and be preferred. More generally, where there is a sufficiently weak positive correlation between emissions and GDP ($0 < \rho < 1$) or the volatility of emissions is sufficiently small relative to GDP ($v[Q] / v[Y] < 1$), Z may be small enough that $\eta \underline{Q} / \overline{Q} > 2\rho v[Q] / v[Y]$, in which case an absolute cap would produce less variance in outcomes and be preferred to an intensity limit. Obviously, for any non-positive correlation (and therefore a non-positive value of Z), an absolute limit would always be preferred since the left-hand-side variables cannot be negative.

The intuition behind these results can grasped by considering first the case of negative correlation. If emissions decline when GDP increases (and vice versa), any amount of indexing to GDP will cause the emissions constraint to vary inversely with deviations in emissions, and will thereby produce greater variance in abatement and cost than an absolute limit would. Alternatively, if correlation is positive, it is still possible that indexing would produce more variance. For instance, if there were no variation in emissions ($v[Q] = 0$) but variation in Y, any amount of indexation to *Y* would create variation in abatement and cost where an absolute limit would produce none. Where there is variation in *Q*, the choice of form of the limit depends on both the magnitude of its fluctuations and the correlation between *Q* and *Y*. Where either of these is sufficiently small, an absolute limit can exhibit less variance than an indexed cap.

The relationships between the stochastic properties of *Q* and *Y*, the desired emission constraint, and the degree of indexation to GDP are illustrated in Figure 1. The target's fraction of initially-expected baseline emissions ranges from zero to one, and is plotted on the horizontal axis. The value of *Z* is given on the vertical axis. The horizontal line *HH′* indicates the value of *Z* for the economy in question ($\overline{Z}$), which represents the boundary between the regions where less variance in outcomes is produced by an absolute or indexed intensity limit ($Z > \overline{Z}$ and $Z < \overline{Z}$, respectively). The diagonal ray *OJ* gives the locus of values of $\frac{1}{2}\eta \underline{Q} / \overline{Q}$ over the range of possible reduction fractions for some value of η. Its maximum slope, which is attained when η = 1, corresponds to the ray OJ^{Max}, which intersects the BAU emission level (i.e., $\underline{Q} / \overline{Q} = 1$) at *Z* = 0.5 on the vertical axis.

The point *K* where *HH′* and *OJ* intersect represents the equality of both sides of eq. (11), and defines the level of an emission target $\underline{Q}*$ below (above) which an intensity limit will exhibit the lower (higher) variance, and thus will (will not) be the preferred instrument. For example, if the degree of indexation is as indicated by the ray *OJ* and the emission target is given by the vertical line *LL′*, the intensity cap would be associated with higher variance and would not be preferred. Conversely, if the constrained level of emissions was much lower, say at *MM′*, an intensity limit with the degree of indexing implied by ray *OJ* would generate less variance, and be preferred.

4.2. Optimal indexation

The preceding section identifies, and Figure 1 illustrates, the relationships between variables that are under the control of policy makers and those that are fundamentally exogenous. If we assume that the stringency of the emissions constraint is determined exogenously without regard to concerns about variance, and that the policy maker desires to minimize variance of outcomes, then the degree of indexation can be used to achieve this goal. It can be easily shown that as long as there is any degree of positive correlation between emissions and GDP ($\rho > 0$) and any variation in emissions ($v[Q] > 0$), there exists a partially indexed cap that will *always* generate less variance in abatement, and will therefore be preferred to an absolute limit, since it is always possible to choose a sufficiently small value for η to shift the sign of the inequality in eq. (11) so that: $\eta \underline{Q} / \overline{Q} < 2Z$.

For a given $\underline{Q}$, the value of the index that minimizes the variance in abatement can be found by differentiating eq. (10) with respect to η and solving the first-order condition to yield:

$$(12) \qquad \eta^{Opt} = \frac{E_\theta[Y]\mathrm{cov}[Q^{BAU}, Y]}{\underline{Q}\,\mathrm{var}[Y]} = \frac{Z}{\underline{Q}/\overline{Q}}.$$

Substituting this expression into eq. (10) yields the minimized variance of abatement: $\mathrm{var}[\tilde{A}^I] = (1 - \rho^2)\,\mathrm{var}[Q]$ for the optimally indexed limit. Any non-zero value of ρ creates the possibility of an indexed limit that will exhibit less variance than the absolute limit. The optimal index will have the same sign as Z—which is to say ρ, so that for feasible values of the indexation parameter, i.e. $0 < \eta^{Opt} \leq 1$, indexing will be attractive only in the presence of positive correlation, as was previously demonstrated.

Eq. (12) implies that for any emission target $\underline{Q}$, there is a level of indexation given by η^* = 2 η^{Opt}, which equalizes the variance of the indexed and absolute forms of the limit, making the policy maker indifferent between them. However, whether η^* or η^{Opt} lie between zero and one depends on the particular values of Z and $\underline{Q}$. This outcome is also captured by Figure 1, where $\eta = \eta^* \in [0,1]$ is indicated by the OJ locus passing through the intersection of HH' and $\underline{Q}/\overline{Q}$. The corresponding optimal value of η is associated with the ray OJ' that would intersect $\underline{Q}/\overline{Q}$ at $\overline{Z}/2$. For any $Z \in$ (0, 0.5], both η^* and η^{Opt} will fall within the range [0,1] so long as $2Z \leq \underline{Q}/\overline{Q} \leq 1$. For $Z \leq 0$, there is no point in considering an indexed limit, while for $Z > 0.5$, a fully indexed limit will always exhibit less variance than an absolute cap and the minimum variance will be achieved by a partially-indexed cap so long as $Z < \underline{Q}/\overline{Q} < 1$.

4.3. Measurement issues and their implications for instrument choice

An unstated assumption which underlies the foregoing results is that policy makers have the capability to choose among instruments based on the true moments of the distribution of GDP and BAU emissions. But the true moments are not observed—rather they are inferred from a finite sample of data. Consequently, to give empirical content to the results obtained thus far, the *population* variances and covariance in (9) and (10) should be replaced by their *sample* counterparts, which are conditional on θ. Consequently, the right-hand side of (11) becomes:

(13) $$Z_\theta = \rho_\theta \, v_\theta[Q] \,/\, v_\theta[Y]$$

To clarify the implications of this expression, consider the effect of new information about the conditions which affect output and emissions. The latter represents a shift in the information set to θ' (say). This might not only induce a revision of the expectations which led to the setting of the emission target (i.e., the denominators of $v_\theta[Q]$ and $v_\theta[Y]$)—a sufficiently large structural shift could also affect policy makers' estimates of the variances of these quantities, as well as the perceived correlation between emissions and GDP. The upshot is that the quantity Z is not immutable, rather, its value evolves as conditions change.

Two consequences follow. First, the optimal degree of indexation will no longer be constant: by eq. (12), even an arbitrarily small shift $Z_\theta \to Z_{\theta'}$, will induce a change $\eta_\theta^{Opt} \to \eta_{\theta'}^{Opt}$ for any emission target $\underline{Q}$. Second, a large enough shift in Z can switch the direction of the inequality in (11), with the result that the even the binary choice between an absolute and a fully-indexed intensity cap may not remain constant over time. Policy makers should be concerned about such outcomes because of the often substantial errors which attend forecasts of emissions (Lutter 2000) and the CO_2 intensity of GDP (Strachan, forthcoming; Philibert 2005), which could lead to drastic revisions of expectations. We undertake an assessment of this issue in the following section.

5. Empirical Tests

We illustrate the practical importance of the foregoing theoretical results by examining the properties of Q and Y, and their implications for the choice of the form of an emission target, using two different approaches. In the first of these, we make assumptions about the character of policy makers' information set and the procedures they follow in using such information to estimate future values of Z. We employ historical data on emissions and GDP for a large number of countries, and define the information set on the assumption that policy makers' will invariably make projections of Z based on recently available data. Our second approach attempts to proxy for the conditional moments of Z directly by using a sequence of forecasts of emissions and GDP for a fixed future year. In both cases, the changes in the projections of Z yield insights into whether countries will tend to stick with an absolute or an intensity limit, or will be more likely shift back and forth between the two instruments as circumstances change.

5.1. Using Historical Time Series Data

Our first experiment focuses on observed values of Q and Y, for which there is an abundance of data. Using statistics on carbon emissions from Marland et al (2003) and real GDP from the Penn World Table 6.1 (Heston et al 2002) we compile a dataset of 30 developed and developing countries over period 1950-2000, from which we compute the value of Z.

Our use of these historical statistics attempts to recreate the kind of prospective assessment and data availability lags that are characteristic of climate policy. We therefore assume that a constraint which is in effect in a particular year (say t') is determined based on data which become available with a five-year lag and are observed over the course of a decade—i.e., the interval $(t'–15, t'–5]$. Moreover, since at t' current emissions and GDP are not observed, we approximate the denominators of $v[Q_{t'}]$ and $v[Y_{t'}]$ using forecasted quantities, which we estimate based on the growth rates of these variables over the lagged observation period.[7] Thus, for a constraint that is assumed to take effect in 1990, we use the data from 1975-85 to determine the value of Z. Table 1 provides these values for 14 countries for the 1990 experiment as well as for constraints in 1980 and 1999 (where we use data from 1965-75 and 1986-94, respectively).

The most striking feature of Table 1 is the strong positive correlation between emissions and GDP for developing countries over the length of the entire sample period, and for developed

---

[7] Formally, $v_{t'}[\zeta] = \left\{ \frac{\sigma_{\zeta}}{\zeta_{t'-5} \cdot \exp\left(5 \cdot E[g_{\zeta}]\right)} \middle| \zeta \in (\zeta_{t'-15}, \zeta_{t'-5}] \right\}$, where $\zeta = Q, Y$, and σ_ζ and g_ζ denote the historical standard deviation and historical average annual growth rate of each of these quantities.

countries before 1975 and after 1985. By contrast, OECD nations exhibit weak or even negative emissions-GDP correlation throughout the decade of high energy prices. The coefficients of variation of emissions and GDP are an order of magnitude smaller and similar in size, and show no trend in the dominance of one type of volatility over the other.[8] The values of Z mostly exceed 0.5. Of the 42 data points in the table, 31 indicate an unambiguous preference for an intensity limit, six indicate an unambiguous preference for an absolute limit, and the remaining five instances could go either way depending on the stringency of the emission target and the degree of indexation. The unambiguous choice of an intensity cap is far more characteristic of the developing countries than the developed countries due mostly to the consistently high temporal correlations between emissions and GDP. The upshot is that intensity caps are unequivocally preferable for developing countries, and may be generally preferable for developed countries. The qualification to the latter conclusion arises from the potential for rapid energy price increases to decouple emissions and GDP.

We conduct a more systematic exploration of these outcomes by computing annual values for the indifference point $Q^*/\overline{Q} = 2Z$ over the period 1965-1999 on a rolling basis for a sample of 22 developed and 7 developing countries.[9] Figure 2 presents these results as probability density functions (PDFs). In both panels, the shaded region corresponds to the range of values in which the choice of an absolute or indexed limit depends on the values of η and $Q^*/\overline{Q}$. In panel A, the bulk of the probability masses of both developed and developing

---

[8] The exceptions are India, South Korea and Mexico, whose emissions are persistently more variable than their GDP.

[9] The OECD country panel (N = 790) is made up of Australia, Austria, Belgium, Canada, Denmark, Finland, France, Greece, Iceland, Ireland, Italy, Japan, Luxembourg, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, UK, and USA. The developing country panel (N = 247) is made up of Brazil, China, India, Mexico, South Korea, South Africa, and Turkey.

countries lie to the right of this range (which we henceforth refer to as the equivocal region).[10] In terms of the geometry of Figure 1, this means that the point K lies completely to the right of the 0-1 scale, so that binding emission limits will tend to be positioned to its left. These results echo our previous findings, and imply a clear preference for the use of an indexed intensity limit, especially in developing countries.

The box plot in Panel B illustrates the substantial inter-country heterogeneity which underlies the foregoing conclusion. While the entire PDFs of $2Z$ for India, South Africa, Mexico and Korea lie to the right of the equivocal region, portions of the first quartiles of the distributions for Brazil and especially China overlap with the feasible region, indicating that in some (albeit rare) circumstances these countries might prefer an absolute cap. Even among developed countries, the medians of the distributions of the indifference point almost always exceed unity, again indicating a preference for intensity limits. Nevertheless, their lower quartiles intersect the equivocal region and the negative orthant to a greater degree than is the case for the developing countries, indicating that there are more occasions when an absolute cap might be preferred, especially in countries such as France and the U.K.

For each of the foregoing observations of countries in a given time-period, we also calculate the optimal degree of indexation for emission targets set at 95 and 75 percent of BAU levels using eq. (12). The box plots in Figure 3 give the PDFs of the corresponding values of η^{Opt} for each country. There are broad similarities with the results for the indifference levels of the emission target, with slight differences for individual countries. The bulk of the probability masses for large non-OECD emitters lies to the right of the range of allowed values of η^{Opt}

---

[10] The probability of the indifference point falling in the range [0,1] is less than unity is 28 percent for OECD countries and only 3 percent for non-OECD countries, while the probability of it being negative is 14 percent for OECD countries and only 1.2 percent for non-OECD countries. As in Table 1, the long lower tail of the distribution for OECD countries reflects the influence of the period of high energy prices from 1974-84, and the consequent negative correlation between emissions and GDP over this period.

(denoted by the shaded area), indicating that fully indexed intensity limits would produce the least variance in outcomes for these countries. Although the PDFs of OECD countries overlap the shaded region to a greater degree, the results for some of these countries, such as Australia and Spain, are similar to the non-OECD patterns.

5.2. Using Historical Forecasts

While historical data are plentiful, for our purposes it suffers from the defect of assuming that policy makers are purely extrapolative in their expectations and that they would not incorporate expected changes from past experience in their set of expectations. Historical forecasts would remedy this problem, but there is a dearth projections of emissions and GDP. Nevertheless, forecasts were made annually for a small number of regions for the years 2000 and 2010 by the DOE/EIA for the International Energy Outlook.

We focus first on the year 2000, for which there are the longest series of comparable historical forecasts over the broadest range of countries. EIA prepared forecasts of emissions and GDP in this year for four developed economies (USA, Japan, Canada and OECD Europe), one economy in transition (the Former Soviet Union—FSU) and two industrializing economies (China and Mexico).[11] These country series were used to compute values for ρ, $v[Q]$, $v[Y]$ and Z, with the source of variability being the changes in expectations captured by the revisions to the DOE/EIA forecasts of the relevant variables. The results for 2000 are shown in panel A of Table 2. A first result to note is that none of the values of Z are negative and two of the seven exceed 0.5. Therefore, in no region does an absolute limit generate less variance than a partially indexed limit, and in only two regions (OECD Europe and the FSU) does an indexed limit

---

[11] The date of the last forecast is 1999 for all of these regions, but the date of the first forecast differs by region. Complete data were available for Canada, Europe, Japan and the U.S. from 1987, for China from 1990, for FSU from 1994, and for Mexico from 1995.

unambiguously generate less variance than an absolute cap. For the remaining countries, the question of which limit exhibits less variance depends on the emissions target and the degree of indexation. The last two rows of panel A provide the optimum values η* for emissions constraints of 0.95 and 0.75. A fully indexed limit is indicated only for the FSU; for all others, a partially indexed limit would minimize variance.

The defining characteristic of this result is not so much the values of the correlations between Q and Y (which, except for China, are all comparable in magnitude and small), but the variability of emission forecasts relative to that of GDP forecasts. For OECD Europe and the FSU, the variability of emission projections exceeds that of GDP forecasts by a factor of three, so that a high degree of indexation is warranted, despite a relatively low Q-Y correlation. By contrast, China is an example of a case where indexation has *less* of a tendency to reduce variance, despite high correlation, because the variability in emissions forecasts is so much less than that for of GDP forecasts.

To test the robustness of these findings we re-compute values using the forecast data for the year 2010, projections for which are available from 1990 onward. The results, shown in panel B, exhibit some interesting differences but the conclusions are broadly the same. The values of Z for OECD Europe and the FSU exceed 0.5, which continues to argue unequivocally for an intensity cap. However, there are now two countries, Japan and Mexico, with negative values of Z, which points unequivocally to the use of an absolute cap. The remaining countries fall in the interval $0 < Z < 0.5$, for which the choice of instrument could go either way. For the five regions for which indexing is indicated, OECD Europe has joined the FSU as a region that would choose a fully indexed intensity cap to reduce variance because the correlation of emissions and GDP is considerably stronger for the 2010 forecasts than for those for 2000. For the remaining three regions, a partially indexed intensity cap would be optimal.

5.3. Comparing the two sets of experiments

The results from the forecast tend to support those from the historical data, namely, conditions that would suggest a general preference for indexed intensity limits. But they also provide clear evidence that these conditions are far from universal. More importantly, the results highlight the dependence of the choice between an absolute and an intensity cap of the expected statistical relationships between emissions and the GDP. The much larger sample for the historical data might argue for placing more confidence in those results than the few instances of actual repeated forecasts of emissions and GDP that we could find. But even though the latter are restricted to one source and a fairly narrow period of time, they have the advantage of being able to indicate how actual expectations evolve, whereas the experiments based on historical data suffer from the assumption of extrapolative expectations which remain constant as conditions change moving forward in time.

Moreover, for any given country, what may be preferred for one interval of time may not be for another period. For instance, for many of the developed countries, an intensity limit would have been the wrong choice for late 1970s and early 1980s, but then would have returned to being the right choice when energy prices declined after 1985. Thus, a policy maker faced with such a choice of limit would need to pay close attention to factors that might shift the historical relationship between Q and Y. For instance, at the time of this writing, when energy prices are once again at high levels and are expected to remain there, intensity limit might not be as strongly preferred as past data from the low-energy-price 1990s might suggest.

6. Conclusion

This paper has sought to elucidate the differences under uncertainty between absolute and intensity-based limits as they may be applied to CO_2 emissions. We demonstrate that the two are identical when there is no uncertainty about the future, and analyze the choice between them on the assumption that the policy maker wishes to reduce the variance in environmental and economic outcomes from the application of the limit. This analysis consists of identifying the conditions under which an intensity-based limit would be preferred to an absolute limit and of specifying the optimal index when an intensity-based limit is preferred. We also investigate the frequency of the conditions for preferring an intensity-based limit using historical data and past forecasts, as well as the distribution of the optimal level of indexing conditional on the emissions constraint.

The main result of the mathematical analysis is that positive correlation between emissions and GDP (or whatever other index is chosen) is a necessary but not sufficient condition for an intensity limit to be preferred. In addition, the variability of emissions relative to income must be sufficient to make indexation variance-reducing. Otherwise, intensity-based limits will increase the variance of outcomes. Alternatively, there are conditions under which absolute limits would minimize variance and be preferred. The empirical part of the paper shows that conditions favoring intensity-based limits predominate but that the conditions in which absolute limits would be variance-reducing cannot be dismissed. Moreover, the choice of the optimal index, as well as the binary choice between an absolute or intensity-based limit, can change over time as conditions and expectations change.

In this paper, we do not wish to suggest that other means of limiting variance in outcomes are not available. Safety valves and temporal trading (banking and borrowing) have similar although not identical advantages in avoiding undesirable outcomes. Our purpose has been to clarify the differences between absolute and intensity-based emission limits that are often discussed as if used in pure form. An important underlying assumption of the paper is that the reduction of the variance in intended outcomes is an important consideration in policy choices. If policy makers are concerned mostly with expected effects, the form of the limit is not as important so long as the two are ex ante equivalent.

Appendix. Further Results on the Equivalence of Absolute and Intensity Limits

We consider a situation in which GDP and emissions are known with certainty at a particular reference point in time, given by *t*, and policy makers commit to an emission target $\underline{Q}$ which is to take effect in some future period, *t*+*k*. We further assume that expectations are conditioned on data on the economy in the reference period, and use the subscript *t* as a shorthand to represent the information set $\theta(t)$. In this setting, the projected emission intensity of the economy under the cap is given by the analogue of eq. (3):

(A-1) $\underline{\gamma}_{t+k} = \underline{Q}_{t+k} / E_t[Y_{t+k}]$.

An Emission Target Based on the Growth of GDP

An intensity cap may be expressed in terms of the rate of growth of emissions. In particular, policy makers may choose to limit the growth of emissions to some maximum allowable fraction, $\overline{\varpi}$, of the expected growth of GDP over the period *t* and *t*+*k*:

(A-2) $(\underline{Q}_{t+k} / Q_t - 1) = \overline{\varpi}(E_t[Y_{t+k}]/Y_t - 1)$.

For the indexed limit in eq. (8) to behave similarly to the growth target specified above, it must be the case that

(A-3) $(\tilde{Q}_{t+k} / Q_t - 1) = \tilde{\varpi}(E_t[Y_{t+k}]/Y_t - 1)$,

in which $\tilde{\varpi}$ specifies the fraction of the rate of GDP growth at which emissions are allowed to increase. It is obvious that $\tilde{Q}_{t+k} = \underline{Q}_{t+k}$ if $\tilde{\varpi} = \overline{\varpi}$, implying that emissions are allowed to grow by the same fraction of GDP under both the absolute and the intensity cap, so that the two instruments are ex ante equivalent.

This result does not generally hold under uncertainty. Using (A-1), (A-2) and (A-3) to substitute for $\underline{\gamma}$, $\underline{Q}$ and $\tilde{Q}$ in (8) allows us to solve for $\tilde{\omega}$ as follows:

$$\tilde{\omega} = \frac{1}{E_t[g_Y]} \left\{ \left[(1-\eta) + \eta \frac{Y_{t+k}}{E_t[Y_{t+k}]} \right] (1 + \overline{\omega} E_t[g_Y]) - 1 \right\},$$

where $E_t[g_Y] = E_t[Y_{t+k}] / Y_t - 1$ is the projected rate of GDP growth between t and $t+k$. This expression makes clear that $\tilde{\omega}$ will diverge from $\overline{\omega}$ as GDP at $t+k$ differs from its forecast value, and the gap between these parameters will increase the more accommodation is made for fluctuations in GDP (i.e., as $\eta \rightarrow 1$).

An Emission Target Based on the Growth of Emission Intensity

An intensity cap may also be expressed as an upper bound on the future rate of decline in the economy's emission intensity. Denoting this maximum rate by $\overline{\phi}$ we have:

$$\text{(A-4)} \quad \overline{\phi} = \frac{\underline{Q}_{t+k} / E_t[Y_{t+k}]}{\gamma_t} - 1.$$

For our indexed cap to behave in the same way, it must be the case that:

$$\text{(A-5)} \quad \tilde{\phi} = \frac{\tilde{Q}_{t+k} / E_t[Y_{t+k}]}{\gamma_{t+k}} - 1,$$

where $\tilde{\phi}$ specifies the rate of decline in the emissions intensity of the economy. As before, once $\tilde{Q}_{t+k} = \underline{Q}_{t+k}$, the limits produce identical effects if $\tilde{\phi} = \overline{\phi}$, thus proving their equivalence under certainty.

To illustrate the effect of uncertainty, we use (A-1), (A-4) and (A-5) to substitute for $\underline{\gamma}$, $\underline{Q}$ and $\tilde{Q}$ in (8) and then solve for $\tilde{\phi}$ to obtain:

$$\tilde{\phi} = \left[(1-\eta) + \eta \frac{Y_{t+k}}{E_t[Y_{t+k}]} \right] (1+\bar{\phi}) - 1.$$

As before, $\bar{\phi}$ and $\tilde{\phi}$ will diverge if actual GDP differs from its forecast value, and the gap between the two policy variables increases as $\eta \to 1$.

References

1. Argentina (1999). Revision of the First National Communication to the U.N. Framework Convention on Climate Change, ARG/COM/2 B.
2. Barros, V. and M. Conte Grand (2002). Implications of a Dynamic Target of Greenhouse Gas Emission Reduction: The Case of Argentina, Environment and Development Economics 7(3): 547-569.
3. Baumert, K., R. Bhandari and N. Kete (1999). What Might a Developing Country Commitment Look Like? Climate Notes, Washington DC: World Resources Institute.
4. DOE/EIA (1998). International Energy Outlook, DOE/EIA-0484(98), Washington DC.
5. DOE/EIA (2004). Annual Energy Outlook, DOE/EIA-0383(2005), Washington DC.
6. Dudek, D. and A. Golub (2003). “Intensity” targets: pathway or roadblock to preventing climate change while enhancing economic growth?, Climate Policy 3 (Supplement 2): S21-S28.
7. Ellerman A.D. and I. Sue Wing (2003). Absolute vs. Intensity-Based Emission Caps, Climate Policy 3 (Supplement 2): S7-S20.
8. Fisher , A.C. (2001). Uncertainty, Irreversibility, and the Timing of Climate Policy, paper prepared for the conference on the Timing of Climate Change Policies, Pew Center on Global Climate Change, Washington DC, October 10-12.
9. Fischer, C. (2003). Combining rate-based and cap-and-trade emissions policies, Climate Policy 3 (Supplement 2): S89-S103.
10. Frankel, J. (1999). Greenhouse Gas Emissions, Bookings Institution Policy Brief No. 52, Washington D.C.

11. Gielen, A.M., P.R. Koutstaal, and H.R.J. Vollebergh (2002). Comparing Emission Trading with Absolute and Relative Targets, paper prepared for the 2nd CATEP Workshop on the Design and Integration of National Tradable Permit Schemes for Environmental Protection, London, 25-26 March.
12. Hahn, R.W. and R.N. Stavins (1999). What Has the Kyoto Protocol Wrought?: The Real Architecture of International Tradable Permit Markets, Washington DC: American Enterprise Institute Press.
13. Helfand, G.E. (1991). Standards versus Standards: The Effects of Different Pollution Restrictions, American Economic Review 81: 622-634.
14. Heston, A., R. Summers and B. Aten (2002). Penn World Table Version 6.1, Center for International Comparisons at the University of Pennsylvania (CICUP).
15. Jacoby, H.D. and A.D. Ellerman (2002). The Safety Valve and Climate Policy, Energy Policy 32(4): 481-491.
16. Jacoby, H.D. and I. Sue Wing (1999). Adjustment Time, Capital Malleability, and Policy Cost, Energy Journal Special Issue: The Costs of the Kyoto Protocol, A Multi-Model Evaluation: 3-92.
17. Jotzo, F. and J.C.V. Pezzey (2004). Flexible Targets for Greenhouse Gas Emissions from Developing Countries under Uncertainty, Economics and Environment Network Working Paper, Australian National University, Canberra.
18. Kolstad, C.D. (2005). The simple analytics of greenhouse gas emission intensity reduction targets, Energy Policy 33: 2231-2236.

19. Kim Y.-G. and K.A. Baumert (2002). Reducing Uncertainty through Dual-Intensity Targets, in K.A. Baumert et al (eds.), Building on the Kyoto Protocol: Options for Protecting the Climate, Washington DC: World Resources Institute, 109-134.
20. Kopp, R., R.D. Morgenstern and W. Pizer (2000). “Limiting Cost, Assuring Effort, and Encouraging Ratification: Compliance under the Kyoto Protocol,” CIRED/RFF Workshop on Compliance and Supplemental Framework <http://www.weathervane.rff.org/features/parisconf0721/summary.html>.
21. Lisowski, M. (2002). The emperor's new clothes: redressing the Kyoto Protocol, Climate Policy 2 (2/3): 161-177.
22. Lutter, R. (2000). Developing countries' greenhouse emissions: Uncertainty and implications for participation in the Kyoto Protocol, Energy Journal 21: 93-120.
23. Marland, G., T.A. Boden, and R. J. Andres (2002). Global, Regional, and National Fossil Fuel CO_2 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.
24. Müller, B. and G. Müller-Fürstenberger (2003). Price-related sensitivities of greenhouse gas intensity targets, Climate Policy 3 (Supplement 2): S59-S74.
25. Newell, R.G. and W.A. Pizer (2006). Indexed Regulation, Resources for the Future Discussion Paper 06-32, Washington DC.
26. Philibert, C. (2005). New Commitment Options: Compatibility with Emissions Trading, Annex I Expert Group on the UNFCCC, prepared for the United Nations Climate Change Conference, Montreal, 5 Dec., 2005 (COM/ENV/EPOC/IEA/SLT(2005)9).

27. Quirion, P. (2005). Does uncertainty justify intensity emission caps? Resource and Energy Economics 27: 343–353.
28. Rosenzweig, R. and M. Varilek (2003). Key Issues to be Considered in the Development of Rate-Based Emissions Trading Programs: Lessons Learned from Past Programs, paper presented at the EPRI Workshop, April 29, Vancouver BC.
29. Spulber, D.F. (1985). Effluent regulation and long-run optimality, Journal of Environmental Economics and Management 12: 103-116.
30. Strachan, N. (forthcoming). Setting Greenhouse Gas Emission Targets Under Baseline Uncertainty: The Bush Climate Change Initiative, Mitigation and Adaptation Strategies for Global Change.
31. U.K. Department for Environment, Food and Rural Affairs (2001). Framework for the U.K. Emissions Trading Scheme. <http://www.defra.gov.uk/environment/climatechange/trading/pdf/trading-full.pdf>
32. White House (2002). U.S. Climate Strategy: A New Approach, Policy Briefing Book, Washington, D.C., 14 February.

Figure 1. The Tradeoff Between Absolute and Intensity Limits

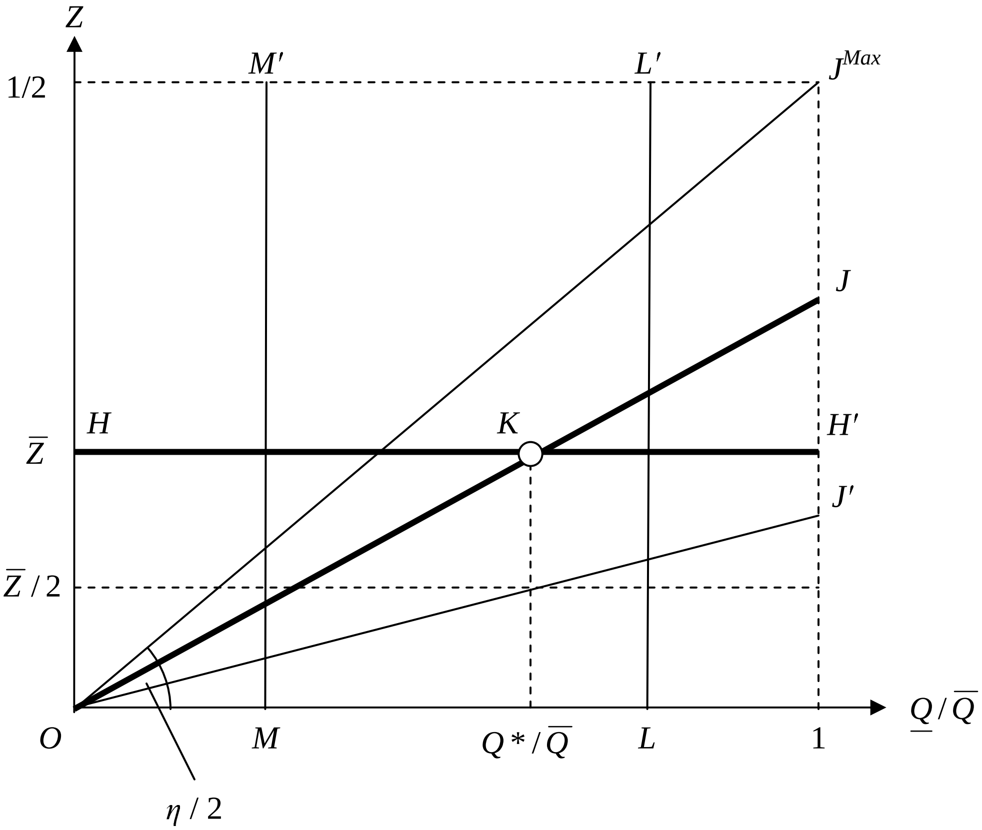

Figure 2. Probability Density Functions for Choosing a Fully-Indexed Intensity Limit

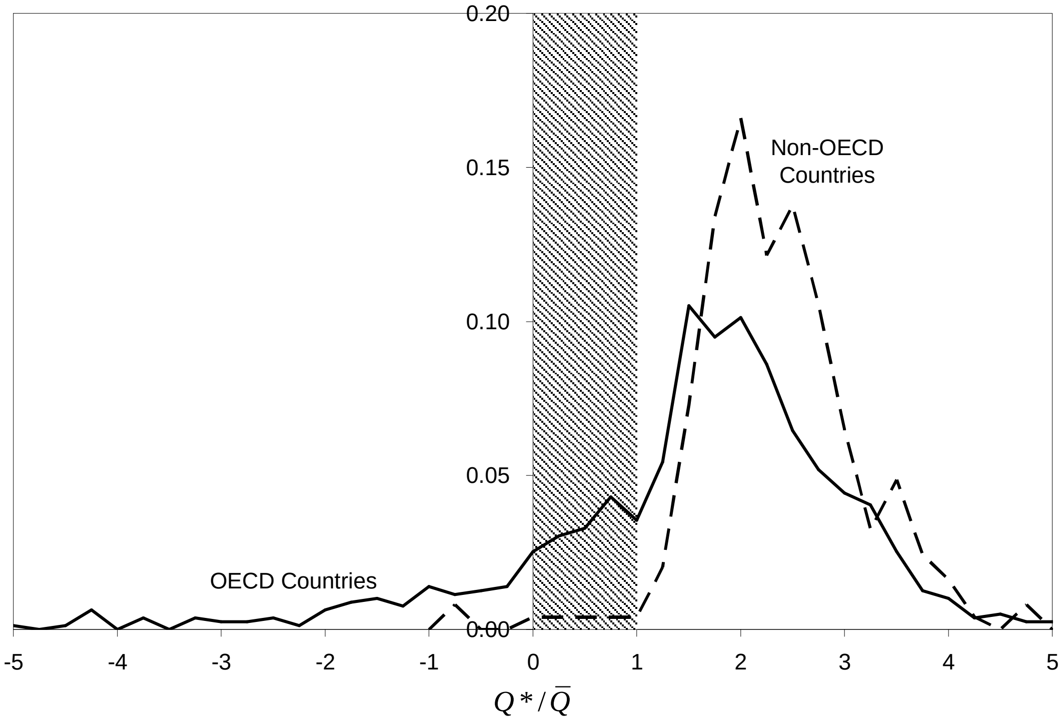

A. Global aggregates

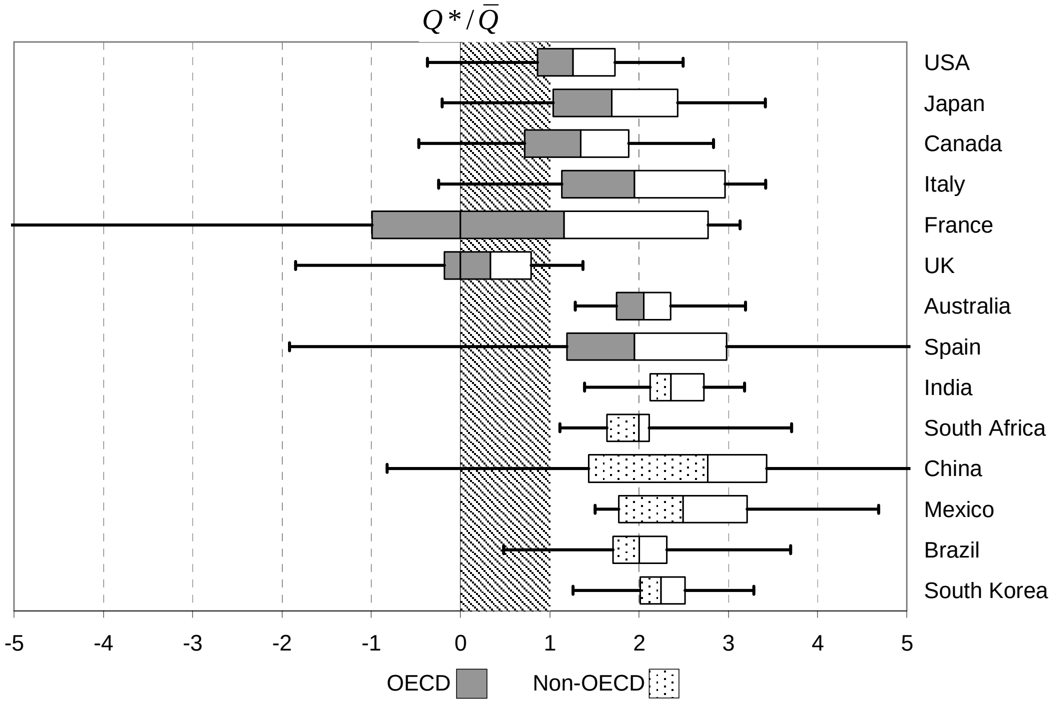

B. High-emitting countries

Figure 3. PDFs of the Optimal Index Conditional on the Level of the Emission Constraint

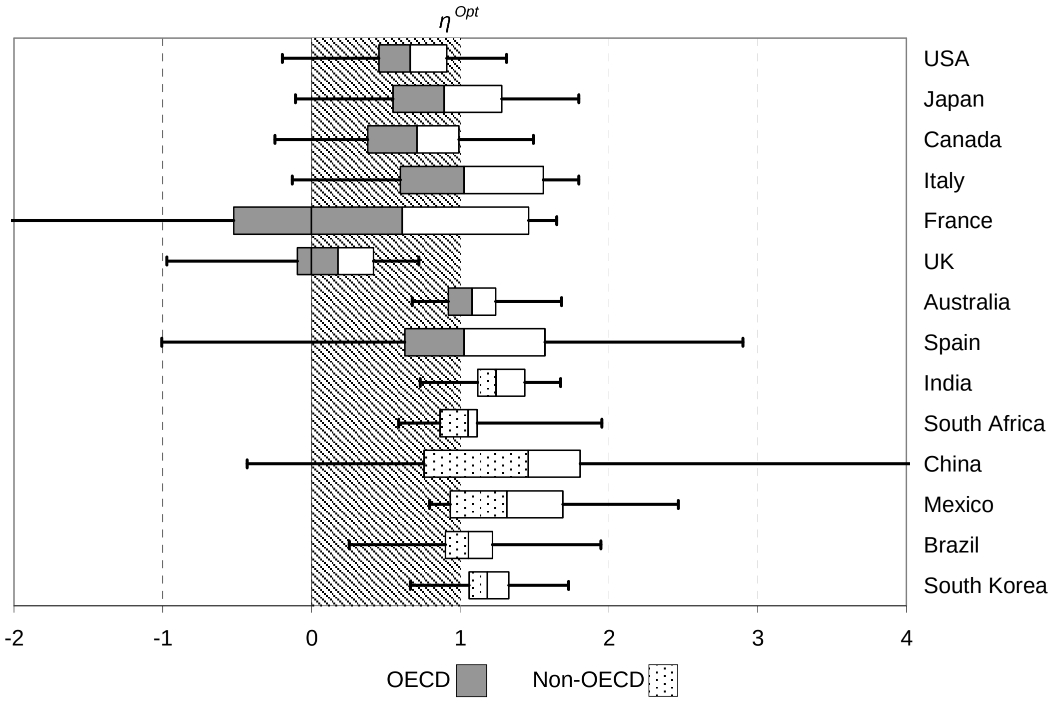

A. $\underline{Q}/\overline{Q}$ = 0.95

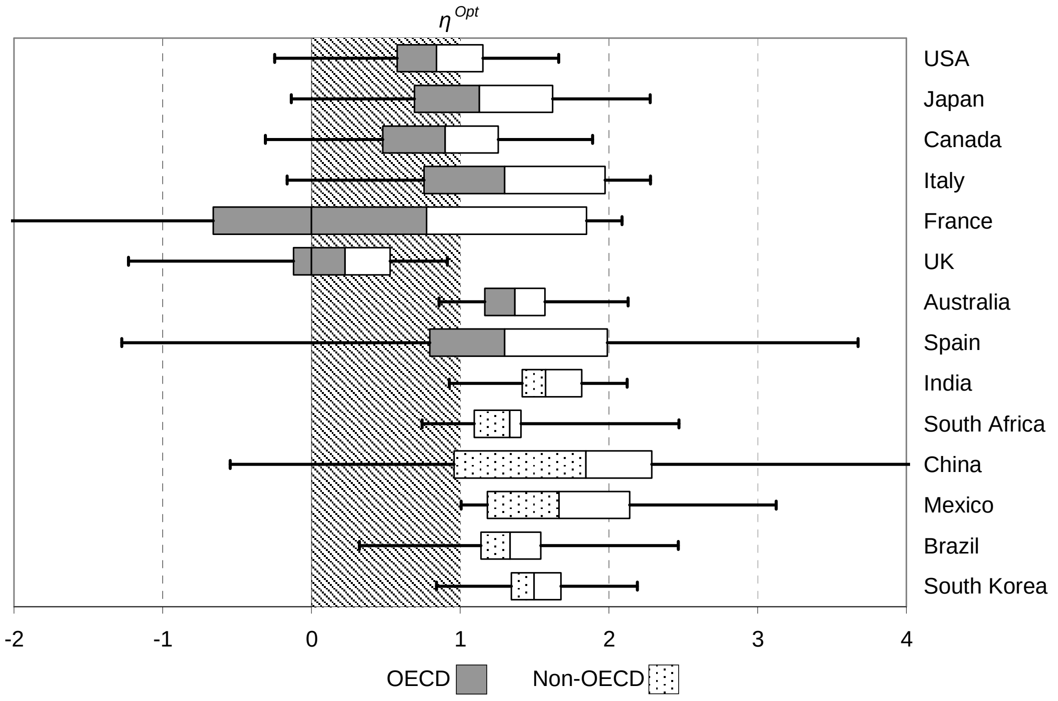

B. $\underline{Q}/\overline{Q}$ = 0.75

Table 1. Empirical Results: Historical Data

	Year	ρ	$v[Q]$	$v[Y]$	Z	Q^* (MTC)	$E[Q^{BAU}]$ (MTC)	Q^{BAU} (MTC)	Q^* / Q^{BAU} [a]
					Developed Countries				
USA	1980	0.929	0.064	0.070	0.847	2470	1458	1300	*1.900*
	1990	-0.276	0.031	0.063	-0.136	-393	1444	1374	**-0.286**
	1999	0.979	0.041	0.061	0.668	2241	1676	1567	*1.430*
Japan	1980	0.994	0.163	0.113	1.442	819	284	251	*3.259*
	1990	-0.046	0.026	0.071	-0.017	-10	286	292	**-0.033**
	1999	0.970	0.069	0.080	0.831	592	356	315	*1.877*
U.K.	1980	0.365	0.026	0.050	0.186	75	203	158	0.476
	1990	-0.583	0.053	0.044	-0.698	-242	173	155	**-1.555**
	1999	0.082	0.044	0.055	0.066	22	171	147	0.153
Canada	1980	0.944	0.092	0.094	0.916	234	128	115	*2.041*
	1990	0.119	0.029	0.067	0.051	13	124	113	0.111
	1999	0.860	0.041	0.049	0.719	207	144	120	*1.725*
Italy	1980	0.987	0.110	0.080	1.358	294	108	102	*2.893*
	1990	0.095	0.027	0.055	0.046	10	112	106	0.097
	1999	0.934	0.042	0.063	0.622	151	122	116	*1.311*
France	1980	0.954	0.083	0.090	0.886	261	147	132	*1.984*
	1990	-0.813	0.103	0.051	-1.658	-393	118	99	**-3.988**
	1999	-0.170	0.045	0.058	-0.131	-27	104	100	**-0.274**
Australia	1980	0.976	0.075	0.083	0.893	100	56	55	*1.801*
	1990	0.885	0.059	0.060	0.866	121	70	73	*1.670*
	1999	0.969	0.069	0.068	0.979	178	91	94	*1.896*
Spain	1980	0.979	0.155	0.101	1.510	170	56	55	*3.114*
	1990	-0.142	0.032	0.022	-0.209	-25	60	58	**-0.431**
	1999	0.869	0.073	0.073	0.866	118	68	75	*1.567*
					Developing Countries				
China	1980	0.983	0.180	0.092	1.909	1447	379	403	*3.590*
	1990	0.939	0.097	0.119	0.764	935	612	655	*1.428*
	1999	0.963	0.093	0.117	0.764	1424	932	771	*1.847*
India	1980	0.896	0.084	0.075	1.010	172	85	95	*1.816*
	1990	0.990	0.122	0.094	1.285	394	153	184	*2.138*
	1999	0.992	0.126	0.093	1.353	733	271	294	*2.494*
South Korea	1980	0.992	0.173	0.119	1.436	75	26	34	*2.202*
	1990	0.952	0.127	0.107	1.128	120	53	66	*1.817*
	1999	0.981	0.156	0.110	1.399	302	108	107	*2.813*
Mexico	1980	0.994	0.126	0.104	1.200	119	49	69	*1.721*
	1990	0.971	0.143	0.105	1.325	227	86	102	*2.212*
	1999	0.891	0.125	0.072	1.559	381	122	113	*3.383*
South Africa	1980	0.957	0.088	0.103	0.816	99	60	58	*1.711*
	1990	0.929	0.107	0.072	1.381	241	87	78	*3.087*
	1999	0.652	0.030	0.030	0.668	132	98	91	*1.449*
Brazil	1980	0.996	0.156	0.126	1.233	114	46	50	*2.285*
	1990	0.334	0.049	0.067	0.243	27	55	55	0.482
	1999	0.880	0.062	0.038	1.426	211	74	83	*2.551*

[a] Bold text indicates that absolute caps are unambiguously preferable, italics indicate that intensity caps are unambiguously preferable.

Table 2. Empirical Results: Forecast Data

	USA	OECD Europe	Japan	Canada	Former USSR	China	Mexico
				A. $T = 2000$			
ρ	0.297	0.233	0.135	0.158	0.313	0.644	0.192
$v[Q]$	0.026	0.090	0.127	0.037	0.446	0.061	0.101
$v[Y]$	0.023	0.031	0.175	0.029	0.147	0.148	0.056
Z	0.336	0.670	0.097	0.205	0.947	0.267	0.349
$E_0[Q^{BAU}]$ [a,b]	1491	672	284	144	786	840	115
$E[Q^{BAU}]$ range [a]	1471-1605	672-1235	273-401	143-161	249-810	840-1031	97-123
Q^* [a]	1001	902	55	59	1487	449	80
Q^* range [a]	987-1077	901-1655	53-78	58-65	471-1533	449-551	67-85
Q^{BAU} [a]	1619	787	323	119	185	762	116
Q^* / Q^{BAU} range	0.61-0.67	1.15-2.1	0.16-0.24	0.49-0.55	2.55-8.29	0.59-0.72	0.58-0.73
$\eta^{Opt}\left(\underline{Q}/\overline{Q} = 0.95\right)$	0.35	0.71	0.10	0.22	1.00	0.28	0.37
$\eta^{Opt}\left(\underline{Q}/\overline{Q} = 0.75\right)$	0.45	0.89	0.13	0.27	1.26	0.36	0.47
				B. $T = 2010$			
ρ	0.597	0.409	-0.042	0.140	0.837	0.575	-0.021
$v[Q]$	0.036	0.111	0.127	0.047	0.205	0.135	0.091
$v[Y]$	0.082	0.035	0.209	0.047	0.125	0.208	0.153
Z	0.260	1.306	-0.026	0.138	1.379	0.373	-0.012
$E_0[Q^{BAU}]$ [a, b]	1819	1101	309	168	1265	944	133
$E[Q^{BAU}]$ range [a]	1621-1835	982-1385	309-466	160-186	666-1265	944-1523	127-164
Q^* [a]	944	2877	-16	46	3490	703	-3
Q^* range [a, c]	842-953	2566-3619	(24)-(16)	44-51	1837-3490	703-1135	(4)-(3)
Q^{BAU} [a, c]	0.48-0.54	2.22-3.14	(0.07)-(0.04)	0.26-0.30	2.17-4.12	0.55-0.89	(0.03)-(0.02)
$\underline{Q}/E[Y]$ Kyoto range [d]	0.68-0.77	0.64-0.91	0.55-0.83	0.63-0.74	0.78-1.49	–	–
$\eta^{Opt}\left(\underline{Q}/\overline{Q} = 0.95\right)$	0.27	1.37	-0.03	0.15	1.45	0.39	-0.01
$\eta^{Opt}\left(\underline{Q}/\overline{Q} = 0.75\right)$	0.35	1.74	-0.03	0.18	1.84	0.50	-0.02

[a] Megatons of carbon.

[b] $E_0[Q^{BAU}]$ = initial emission forecast

[c] Figures in parentheses indicate negative values.

[d] Kyoto emission targets as specified in DOE/EIA (1998: Table 8).