Market-Based Environmental Regulation

Rough Draft Chapter for
Government Regulation: What Have We Learned? (Nancy Rose, ed.)

Roberton C. Williams III
University of Texas at Austin and NBER

August 29, 2005

Without any doubt, the most striking change in environmental regulation over the last two decades has been the rise of market-based regulation–primarily in the form of tradable emissions permits. Twenty years ago, while tradable permits were regarded as an interesting theoretical idea, there was a great deal of skepticism about their practical value for regulation. Only one significant regulatory program had allowed emissions trading, and it had been only modestly successful, falling far short of efficiency gains that simple theory had predicted. Even as recently as a decade ago, the U.S. sulfur dioxide trading program–the largest and most successful emissions trading program thus far–was just getting started, and was commonly viewed as a somewhat risky large-scale experiment (or worse–many environmentalists saw the entire concept of emissions trading as inherently immoral).

Contrast that with the situation today: emissions trading is widely accepted as a reliable and practical method of achieving environmental goals in a cost-effective manner. Most significant new environmental regulations and proposed regulations rely on emissions trading–and while many other aspects of those regulations are highly controversial, the use of emissions trading for the most part is not. Indeed, policymakers seem increasingly inclined to view emissions trading as a panacea–the ideal regulatory approach for any environmental problem. This represents a dramatic shift over a relatively brief span of time.

Prior to the introduction of market-based regulation, all environmental regulation was in the form of *command-and-control* mechanisms: policies that specify either a particular pollution control technology (a *technology mandate*) or a maximum legal threshold for pollution emissions (a *performance standard*).[1] Command-and-control regulation has some

[1] The distinction between technology mandates and performance standards often makes little difference in practice, because performance standards are frequently set by picking a

intuitive appeal. It is simple to set up and to enforce (particularly in the case of the technology mandate, because an inspector just needs to check whether the appropriate technology is installed and active). The regulator can take into account the relative costs and benefits of different control technologies when setting the standard. And such regulations ensure a predictable reduction in emissions from each pollution source.

This last point is particularly important, given that environmental legislation (e.g., the U.S. Clean Air Act) has typically directed a regulatory agency to calculate the level of each pollutant that would be harmful to human health, and to ensure that pollution concentrations remain below those levels. Such legislation leads regulators to adopt a particular pollution threshold as an objective, even for cases in which the marginal damage from pollution is relatively constant, and thus there is no distinct threshold in society's objective function. Partly as a consequence, regulatory agencies have strongly favored regulatory mechanisms that ensure a given pollution threshold will be met.[2]

However, regulation via command-and-control is generally not *cost-effective*: that is, it will not achieve a given overall reduction in pollution at the lowest possible cost, because it will not equalize marginal costs across sources.

I. A Simple Theoretical Model

To make this concept more concrete, suppose that there are a total of N pollution sources, and that the cost of reducing pollution emissions from source i by q_i units is given

---

particular technology and then setting the standard at the level of emissions that technology can reasonably be expected to achieve.

[2] A second reason is that many environmental regulators come from backgrounds in law, engineering, or the natural sciences, each of which focuses much more on direct controls than on incentives when trying to produce a given outcome.

by the function $C_i(q_i)$.[3] A *cost-effective* regulation that reduces total pollution emissions by Q units (i.e., a regulation that regulation that reduces pollution by Q units at the lowest possible cost) would set q at each source based on the solution to the problem

$$(1) \qquad \min_q \sum_{i=1}^{N} C_i(q_i) \quad s.t. \quad \sum_{i=1}^{N} q_i = Q.$$

The first-order condition for this problem is given by

$$(2) \qquad \frac{\partial C_i}{\partial q_i} = \lambda,$$

where λ is the shadow price (Lagrangian multiplier) on the quantity constraint, which represents the marginal cost of pollution control.[4] To be cost-effective, a regulation must equalize the marginal cost of pollution reduction across all sources; if a regulation does not equalize marginal costs, then one could achieve the same total reduction in pollution at a lower total cost by requiring a bit less pollution abatement at sources with higher marginal costs and requiring a bit more at sources with lower marginal costs.

Simple command-and-control regulation would obviously equalize marginal costs in certain special cases–e.g., if all pollution sources are identical, then imposing the same performance standard on all sources would equalize marginal costs. But in general, command-and-control regulation will be cost effective only if the regulator has an enormous amount of knowledge about pollution control costs at every pollution source–enough to

[3] <I'm assuming that it's ok to include some math in the text of the chapter (i.e., that the target audience for the book is sufficiently technical to handle simple multivariate calculus). But if I'm wrong in that assumption, all of the math in this draft could be moved to an appendix or omitted entirely.>

[4] Throughout this chapter, I implicitly assume all of the technical conditions necessary in order for the first-order conditions to define a unique interior optimum (e.g., the cost functions are convex, twice-differentiable, etc.), but for simplicity and brevity I omit the explicit representation of those conditions.

calculate the solution to the cost-minimization problem in (1)–and sets the regulation accordingly.

In contrast, consider the imposition of a pollution tax, as suggested by Pigou (1920). For simplicity, assume that each source is a separate firm. Let τ represent the pollution tax rate and $\overline{e}_i$ represent the uncontrolled level of pollution emissions from source *i* (i.e., the level when $q_i = 0$), which implies that the emissions from source *i* are given by $\overline{e}_i - q_i$. Each firm will want to minimize the sum of its pollution tax bill and its pollution control costs, so firm *i* faces the problem

(3) $$\min_{q_i} C_i\left(q_i\right) + \tau\left(\overline{e}_i - q_i\right),$$

with the corresponding first-order condition

(4) $$\frac{\partial C_i}{\partial q_i} = \tau \, .$$

The pollution tax gives each firm an incentive to reduce pollution up to the point where marginal cost equals the pollution tax rate. If all firms face the same pollution tax rate, then marginal costs will be equalized across firms–and thus the pollution tax will be cost-effective.

This result depends on a number of implicit assumptions made in this simplified model–one of which is that each firm knows its pollution control cost function $C_i\left(q_i\right)$. The total amount of information needed to achieve a cost-effective outcome is the same under a pollution tax as it is under command-and-control. However, because the pollution tax allows each firm to set its own level of emissions, that information requirement is now dispersed among the firms: under the tax, cost-effectiveness requires that each firm know its own pollution control costs, but does not require the regulator to know any of the cost functions, whereas under command-and-control, the regulator must know the pollution control costs for every firm. And in practice, that information requirement is much more likely to be met–each firm is likely to have a better idea of its own pollution control costs than does the regulator. This is one of the two key elements in an incentive-based regulation: that it allows the best-informed agent (which is usually the firm) to decide the level of emissions from any given source.

The second key element, of course, is to provide that agent with the proper incentives. Leaving pollution completely unregulated allows firms to choose their own emissions levels, but provides no incentive to reduce pollution. Imposing a pollution tax provides that incentive.

This is the fundamental idea of incentive-based environmental regulation: to provide firms–who have the best information about their costs of reducing pollution–with both the flexibility to choose their own emissions levels and the proper incentives to reduce emissions to an efficient level.[5]

Despite this cost advantage, pollution taxes have only very rarely been used in practice.[6] One major reason that regulatory agencies are reluctant to use pollution taxes is that a tax does not fix the quantity of pollution allowed; instead, the level of pollution is determined both by the tax rate and by firms' marginal costs of abatement. If the regulator underestimates those marginal costs, then firms will not reduce their emissions by as much as the regulator expected (and these errors are likely to be large in the cases that most strongly

[5] This discussion implicitly assumes that pollution is generated in production. But the concept is the same for the case in which pollution is generated by consumption, and consumers are best informed about the costs of reducing emissions. In that case, an incentive-based regulation would allow consumers to choose how much pollution to emit, and would provide an incentive to reduce emissions.

[6] The most common "pollution" taxes are taxes on gasoline or other energy goods. However, in most cases the primary motivation for such taxes is to raise revenue, not to reduce pollution. <other examples?>

favor incentive-based regulations, since the regulator's limited knowledge of firms' costs provides the rationale for such regulation). As noted above, regulatory agencies often strongly favor policies that ensure a given amount of pollution reduction.

A second reason is that even though pollution control costs are lower under a pollution tax, the burden on polluting firms is higher, because they must pay not only the cost of reducing pollution, but also the pollution tax on the remaining emissions: the cost to firm *i* of a command-and-control regulation is only $C_i(q_i)$, whereas the cost of a pollution tax is $C_i(q_i) + \tau(\overline{e}_i - q_i)$. The difference–the pollution tax–is not a social cost, because it is a transfer from the polluter to the government, but it is still a cost for the polluter. Consequently, polluting firms almost always strongly oppose pollution taxes, and that opposition is often sufficient to prevent the imposition of a tax.[7]

Tradable permits have been successful because, like pollution taxes, they are cost-effective (at least in this simple model), but they do not share the characteristics that make regulators reluctant to use taxes: tradable permits fix the aggregate amount of pollution, and they impose less of a burden on regulated firms than a tax does.

Consider a system in which source *i* is allocated r_i permits, is allowed to sell permits to and buy permits from other sources, and must have a permit for each unit of pollution that it emits. Under this system, each firm will try to minimize its pollution control costs plus the permit price times its net demand for permits (the quantity of permits it needs–which is equal

[7] Institutional features of the U.S. government also work against the implementation of a pollution tax. In particular, tax bills are handled by a different Congressional committee than are environmental regulations, thus making it more difficult to enact a pollution tax than to enact other regulations.

to the quantity of pollution it emits–minus its initial allocation).[8] If the market is large enough that all of the firms take the permit price p as given, then each firm faces the problem

$$\min_{q_i} C_i(q_i) + p(\bar{e}_i - q_i - r_i), \tag{5}$$

which implies the first-order condition

$$\frac{\partial C_i}{\partial q_i} = p. \tag{6}$$

Thus, each firm will equate its marginal cost of abatement with the permit price p. That price will adjust so that the permit market clears. This gives the equilibrium condition

$$\sum_{i=1}^{N}(\bar{e}_i - q_i - r_i) = 0 \Leftrightarrow \sum_{i=1}^{N}(\bar{e}_i - q_i) = \sum_{i=1}^{N} r_i, \tag{7}$$

which simply says that the aggregate net demand for permits equals zero, or, alternatively, that the total amount of pollution emitted equals the total number of permits handed out.

It is then straightforward to show the three key results just mentioned. First, because each firm equates its marginal cost of abatement with the permit price (by equation (6)), the permit system equates marginal cost across sources, and thus is cost-efficient. Second, because the total quantity of emissions must equal the number of permits issued (by equation (7)), the system ensures a given overall reduction in pollution (equal to $\sum_{i=1}^{N}(\bar{e}_i - r_i)$ units).

And third, the net cost to any particular firm is given by (5): the firm's pollution control cost plus its net demand for permits times the permit price. But in order for demand to equal supply in the permit market (for (7) to hold), the average firm must have a net demand of zero. So the net cost for the average firm is simply equal to pollution control costs: firms

[8] Note that a firm's net demand for permits may be either positive (if it buys permits) or negative (if it sells permits).

that buy permits bear a cost greater than their control costs, but that additional cost is paid to the firms that are selling permits–whereas in the case of the pollution tax, the tax payments go to the government and thus every polluting firm faces a cost that exceeds its direct control costs.[9] Thus, firms in the polluting industry will object much less strenuously to the imposition of a system of tradable permits than they would to the imposition of a pollution tax.

And for any given initial allocation, firms will prefer tradable permit regulation over command-and-control regulation (or, in other words, every firm i is at least as well off under a tradable permit system that assigns it r_i permits as it would be under a command-and-control regulation that requires it to reduce its emissions by $\overline{e}_i - r_i$ units–because if it chooses not to buy or sell permits, it will face the same cost as it would under the command-and-control regulation, and if it does buy or sell, then it must be better off).

But even given these attractive properties, the idea of permit trading faced a great deal of skepticism from regulators, who saw it as being attractive in theory, but not a practical solution for real-world environmental regulation. Nonetheless, despite that skepticism, the United States Environmental Protection Administration did implement three different emissions-trading programs between the late 1970s and the early 1990s.

---

[9] Of course, a pollution tax could be structured such that it would have the same distributional effect as this system of tradable permits (e.g., if the pollution tax revenues were returned lump-sum to firms in the polluting industry). Similarly, the permit system could be structured such that it would have the same distributional effects as the tax–if firms had to buy the emissions permits, rather than receiving them for free. In practice, however, pollution tax revenues are very rarely refunded to firms in the polluting industry, and emissions permits are nearly always given out for free.

II. Early Emissions Trading Programs

These first three U.S. emissions-trading programs differed widely in terms of the characteristics of the pollutant being regulated, the design of the program and rules for trading, and many other important details. But they were similar in that they were intended to address a dynamic regulatory problem, not just to minimize regulatory costs in a static setting. The first trading program, the "offset policy" (which was subsequently broadened to become the "Emissions Trading Program"), was originally intended to address the issue of how to permit new entrants into a polluting industry in a region that was already exceeding emissions targets. The next two trading programs–dealing with leaded gasoline and with ozone-depleting chemicals–were each designed to ease the phase-out of a good that was so pollution-intensive that its use needed to be eliminated entirely, but which could not be eliminated all at once.

In each of these three cases, this dynamic aspect made the problem especially difficult to address using a traditional command-and-control approach–and thus opened the door for an alternative.

A. The Offset Policy and Emissions Trading Program

The "Offset Policy" began in 1976 as a means of addressing the problem of accommodating new entrants into polluting industries in "nonattainment regions"–regions that had failed (or were obviously going to fail) to meet the deadlines established for achieving the air quality standards mandated by the Clean Air Act. These nonattainment regions posed a difficult challenge for the EPA. By law, the EPA was required to reduce air pollution in these regions, and to prevent any activity that would lead to further deterioration of the air quality.

But command-and-control regulation was ill-suited for this task. Tightening the legal requirements on existing sources would be politically very difficult, because it would drive up the costs of regulation for all of those existing sources, who could therefore be counted on to oppose the policy. Moreover, at least some of those existing sources would find it very expensive to comply with a stricter standard, because they had already installed pollution control equipment designed to comply with the existing rules. On the other hand, if the EPA did not reduce emissions from existing sources, then it would have to ban all new entrants, because allowing any new pollution sources without reducing pollution from existing sources would inevitably lead to an overall worsening of air quality. And given rapid economic growth in some of the nonattainment areas, banning new entrants was an unacceptable option.

The solution was to give existing sources the option of reducing their pollution emissions to a level below what was legally required. The EPA would certify that additional reduction, providing "emission reduction credits," which could then be sold to new sources. Under this policy, a new pollution source would be allowed to enter a nonattainment area only by acquiring enough credits to offset 120% of the emissions that it would produce. In effect, this policy imposed a 20% "tax" on credits purchased by new entrants, so that each new entrant would actually produce a reduction in overall emissions. This provision was not the most cost-effective way to achieve that reduction–it did not equalize marginal costs, providing too much incentive for new entrants to reduce emissions (or avoid entering altogether) and too little incentive for emissions reductions at existing sources–but it did provide a way to allow new entrants and to improve air quality without imposing additional costs on existing pollution sources.

The EPA then broadened the program, introducing three additional provisions, known as the bubble policy, netting, and banking–which, together with the offset policy, made up the Emissions Trading Program. The bubble policy allowed a firm to combine the emissions limits for across different sources–thus effectively allowing one-for-one transfers of credits among existing pollution sources. Netting allowed existing sources to be modified or expanded without meeting the standards imposed on new sources, as long as the resulting increase in emissions (net of emissions reductions elsewhere in the facility) was sufficiently small. And banking allowed emissions credits to be saved for the future, rather than requiring them to be used immediately.

While this program was a success, its value was severely limited by the bureaucratic requirements that the program imposed on credit trades. Credits were defined relative to what a given source would have emitted under the original command-and-control regulations–a baseline that was much more complicated than a simple emissions cap. Moreover, emission reductions needed to be certified by the government before they could be traded, and each credit transaction needed to be individually pre-approved by the regulator. These requirements imposed substantial costs on each transaction, thus discouraging credit trading–and likely preventing many cost-saving trades from occurring.

B. The Leaded Gasoline Phase-Out

<need to fill in>

C. The Phase-Out of Ozone-Depleting Chemicals

<need to fill in>

III. Fundamental Questions in Tradable Permit Program Design

The experience with these early programs provided important lessons about the practical issues that regulators need to address in order to implement an emissions trading program. In a simple theoretical model, setting up an emissions trading program is easy: just define the initial permit allocations and then allow firms to trade as they see fit. But implementing emissions trading in a more realistic setting presents the regulator with far more questions to address. And, as the early emissions trading programs demonstrated, the answers to these questions can have important implications for transaction costs, the scope of potential cost savings, and other factors that crucially influence the efficiency and equity effects of the program.

Which Emissions Sources Should Be Included in the Program?

Perhaps the most basic question that needs to be answered is which pollution sources will be included in the emissions trading program. Several factors come into play here, the first of which is physical location. On the one hand, the potential for cost savings from emissions trading depends on the difference in marginal costs across different sources–and thus the more sources are included in the program, the greater are the potential cost savings. But on the other hand, the simple theoretical analysis in Section I implicitly assumed that the regulator simply wants to minimize the cost of achieving a given reduction in pollution emissions–or in other words, the regulator only cares about how those emissions are distributed across the different pollution sources because that distribution affects the regulation's cost. In order for that objective to coincide with society's objectives, pollution damage must depend only on the total quantity of pollution emitted, and not on its location.

For some pollutants, such as greenhouse gases and ozone-depleting chemicals, this is true regardless of the geographic dispersion of the pollution sources–the effect of these pollutants is the same no matter where on Earth they are emitted. In this case, there is no reason to limit the geographic scope of the trading program.[10] But for other pollutants (e.g., air pollutants that contribute to smog formation, such as VOCs), location matters much more; a given quantity of pollution emitted from a source in one city may have a very different effect than would the same quantity of the same pollutant emitted from a source in a different city. In this case, it is likely most efficient to restrict the geographic scope of the trading program to just one city.[11]

Other considerations include which industries to include in the trading program, and what size firms should be included. Including a wider range of firms expands the potential for cost-saving trades. But emissions trading programs generally require detailed monitoring of emissions, which may be more costly in some industries than in others. Similarly, a small firm may not account for enough emissions to make it worthwhile to incur the fixed cost of monitoring that firm's emissions.

How Are the Permits Defined?

The next set of fundamental design questions relate to how the emissions permits are defined. Perhaps the most important of these questions is what type of baseline is used to define the reductions in emissions. One alternative is a credit system, which defines emissions reductions relative to the level of emissions allowed under a pre-existing regulation,

[10] A potential exception would involve extending the trading program to sources in different nations, which might make enforcement more difficult.

[11] Williams (2005) provides a detailed analysis of the optimal geographic scope of emissions trading (or temporal scope of emissions banking).

such as a technology mandate. This approach might seem logical, particularly for cases in which the permit trading system is replacing (or augmenting) an existing regulatory framework. However, the extra complexity and bureaucratic hurdles imposed by a credit system can substantially discourage cost-saving trades, as the U.S. experience with the offset program demonstrated.

The alternative, a cap and trade system, simply sets an aggregate cap on emissions, and then allocates the quantity of permits that corresponds to that cap. This provides a clear and simple baseline–each firm needs to hold one permit for each unit of pollution that it emits, and can sell any remaining permits–which minimizes transaction costs. It can also provide better incentives for emissions reduction in certain cases. For example, a credit system will typically not allow any credits in the case of a firm that exits the market entirely, whereas a cap-and-trade system typically would. Since having a polluting firm exit the market generally does reduce emissions, the credit system would not provide the proper incentive for a firm to exit.

Another question is whether permits should be defined in terms of a quantity (tons of emissions, for example) that can be emitted in a particular year or a flow rate (e.g., tons per year–in essence an infinite stream of single-year permits). This distinction matters mainly for transaction costs: if a firm emits less pollution than its permit allocation in one year, but expects to emit more in the following year, then if permits are defined in terms of a flow rate, the firm would have to sell the permit in the first year, and then buy it back in the subsequent year–whereas if permits are defined based on quantity in a particular year, then only one transaction is necessary.

<property rights to permits?

How Are the Permits Allocated?

<Allocation of permits?

-total quantity

-distribution of that quantity (auction? Grandfather? Based on output?)>

<safety valve?>

What Limits Are Imposed on Permit Trades?

<permit banking? Borrowing?>

<interpollutant trading?>

<geographic limits (e.g., Reclaim)>

IV. Potential Problems for Emissions Trading Programs

A. Transaction costs

-may prevent reaching least-cost allocation of emissions

-potential solutions:

-auction at least some permits, so market price is observable

-denominate permits based on quantity, not on flow rate

B. Non-uniformly mixed pollutant (location of pollution matters, not just quantity)

-in this case, trades may increase pollution damage ("hot spot" problem)

-potential solutions:

-"Regulatory tiering": source must have permit <u>and</u> must not cause violation of local standard

-Permit trading zones: sources can only buy/sell permits within same zone

-variant used in RECLAIM program (So. Cal.): two zones, coastal and inland; coastal sources can sell permits to inland sources, but not vice-versa

-Trading ratios: quantity of pollution equal to one credit varies by location

-problems with potential solutions: regulatory tiering and trading zones each may prevent efficiency-improving trades, and worsen transaction costs and market power. Trading ratios avoid this if set properly, but to set them properly in general, need to know efficient pollution allocation.)

C. Temporal hot spots–analogous to spatial hot spots

-potential solutions:

-Don't allow banking/borrowing of permits (but this prevents some cost-reducing trades–analogous to the trading zone approach for spatial hot spots)

-Banking/borrowing ratios, analogous to spatial trading ratios (and with the analogous advantages and disadvantages)

D. Market power:

-in permit market: sell(buy) fewer permits in order to drive the price up(down)

-in output market: manipulate permit price to drive up competitors' marginal costs

-potential solutions:

-If permits are grandfathered, set initial permit allocations close to cost-minimizing pollution allocation

-Auction at least some of the permits

-Avoid setting up markets where concentration will be a problem (e.g., very small trading zones)

E. Inefficient and/or inequitable distribution of rents

-Value of permit rents often exceeds cost of pollution emissions reduction. Furthermore, some of that cost is borne by consumers of the firm's output and suppliers of its inputs. So if all permits are grandfathered, firm receives substantial windfall profits

-Auctioning permits provides revenue for the government. This generates an efficiency gain if the government would otherwise have to use distortionary taxes to raise revenue.

-potential solutions:

-auction as large a share of the permits as is politically feasible

-impose tax to capture at least part of the permit rents (as U.S. did in the CFC case)

V. Recent Experience with Tradable Permit Programs

A. Sulfur dioxide (SO_2) trading program

-Nationwide program covering electric utilities' SO_2 emissions

-Implemented in two phases: dirtiest plants in 1995, the rest in 2000

-Most permits grandfathered; small fraction auctioned, but revenue goes to firms

-Allows banking, but not borrowing, of permits

-Results:

-Permit prices are far less than *ex ante* predictions, so program is viewed as a huge success, but:

-Not all of the difference is due to trading. The costs of reducing emissions by burning low-sulfur coal fell sharply during this period.

-The social marginal cost substantially exceeds the permit price, because of interactions with tax-distorted factor markets

-Still, even after taking those factors into account, cost savings are large

-Allowance trading market appears to be fairly efficient—transaction costs and strategic behavior seem to have little effect on the market

-Substantial allowance banking occurred in the early years of the program, in preparation for tighter standards to come

B. RECLAIM program: local air pollution (nitrogen and sulfur oxides) in the Los Angeles area

-Started in 1994, number of credits is being gradually reduced

-Credits are grandfathered. Banking and borrowing are not allowed.

-Results:

-Transaction costs appear to be significant

-Some concerns about market power for large electric utilities

-Still appears to be much more efficient than CAC regulation

C. NOx budget program

-Cap and trade program covering NOx emissions

-Began with Northeastern states, has since expanded to include a total of 21 states (East Coast and Midwest) plus DC.

-Includes only emissions during the "ozone season" (ozone isn't a problem in winter)

D. EU greenhouse gas allowance trading program

-Started at beginning of 2005, "warm-up" phase through end of 2007

-Covers the 25 countries of the expanded EU, with each country having its own cap

-Only four sectors (iron & steel, minerals, energy, and pulp & paper) included in trading program, but the rest of each country's economy is still covered by the national cap.

-Each country independently decides how to divide it's national cap between the trading sectors and non-trading sectors

-Most permits will be grandfathered, but countries can choose to auction up to 5% during the first phase, and up to 10% in the second phase (starting in 2008).

-"Safety valve" price of 40 Euros/ton in first phase, 100 Euros/ton in second phase

-Each country can choose whether to allow banking or not

VI. Open Questions

A. Practical importance of hot spots (spatial or temporal)? In theory, nontradable permits could be more efficient than tradable permits when pollution isn't perfectly mixed. But no clear empirical evidence yet.

B. Effect of tradable permits within a federal system (e.g., state- and national-level permit programs within the U.S., or national- and EU-level programs in Europe)

-when different levels of government regulate the same pollutant, or regulate two closely-linked pollutants

-when program is set by top-level government, but important details of implementation are left to lower-level government (as in the new EU carbon trading program)

References

<need to fill in>