THE GEORGE WASHINGTON UNIVERSITY

WASHINGTON, DC

Public Interest Comment[1] on

The Environmental Protection Agency's Advanced Notice of Proposed Rulemaking:

State Guidelines for Greenhouse Gas Emissions from Existing Electric Utility Generating Units

Docket ID No. EPA–HQ–OAR–2017–0545
RIN 2060–AT67

February 26, 2018

Brian F. Mannix, Research Professor[2]

The George Washington University Regulatory Studies Center

The George Washington University Regulatory Studies Center works to improve regulatory policy through research, education, and outreach. As part of its mission, the Center conducts careful and independent analyses to assess rulemaking proposals from the perspective of the public interest. This comment on the Environmental Protection Agency's (EPA) Advanced Notice of Proposed Rulemaking (ANPRM) for greenhouse gas emissions from electric generating units does not represent the views of any particular affected party or special interest, but is intended to assist EPA in developing economically efficient options for regulating these sources without exceeding its statutory authority or unnecessarily intruding on state autonomy.

---

[1] This comment reflects the views of the author, and does not represent an official position of the GW Regulatory Studies Center or the George Washington University. The Center's policy on research integrity is available at http://regulatorystudies.columbian.gwu.edu/policy-research-integrity.

[2] Brian F. Mannix served (2005 – 2009) as EPA's Associate Administrator for Policy, Economics, and Innovation. He is currently a research professor with the George Washington University Regulatory Studies Center, 805 21st St. NW, Suite 609, Washington, DC. He can be reached at BMannix@gwu.edu or (202) 994-7543.

Introduction

The EPA is considering whether and how to use the Clean Air Act (CAA) to regulate greenhouse gas emissions (GHGs), especially carbon dioxide (CO_2), as pollutants emitted from electric generating units (EGUs). This task is complicated by the fact that, unlike most other pollutants regulated under the CAA, CO_2 is not an inadvertent byproduct of an industrial process; it is the direct product of fossil fuel combustion. A coal-fired EGU is bound by the second law of thermodynamics, and produces electricity by releasing both heat and CO_2 into the environment. Improvements in efficiency of coal-fired EGUs are certainly possible, but their CO_2 emissions can never be comparable to those of a hydropower EGU.

Despite recent shifts in the economics of electricity production, coal-fired EGUs remain a major part of our energy infrastructure. It is possible to imagine that Congress might some day decide that coal combustion is too harmful and should be banned in favor of, say, nuclear power; but it is also possible to imagine them doing the opposite. So far it has not done either of these things, nor has it delegated the authority to make such momentous decisions to the EPA. What it has delegated to EPA under the CAA, as interpreted by the Supreme Court, is the authority to set standards that reduce the pollution-intensiveness of various categories of industrial facilities, and charged the agency to do so using the "best" (which I read to mean the most economically efficient) system of emissions control.

Toward that end, this comment offers four principles for designing a rule. Together they are intended to achieve an economically efficient outcome—similar to a Pigovian carbon tax—but one that works within EPA's existing authority.

First, EPA should continue to focus on an intensiveness standard rather than a mass-based standard. Arguably this is more congruent with EPA's legal authority, but it has other desirable properties as well. A carbon-intensiveness standard is more equitable in that it avoids the regressive effects of a mass-based standard, and it is likely to be more efficient because it reduces the opportunity for rent-seeking.

Second, EPA should set a standard that is tiered by technology, recognizing the inherent differences in the carbon footprint of coal vs. oil vs. gas vs. other EGU technologies.

Third, EPA should allow for emissions trading, across technology tiers as well as within them. The agency should consider setting targets for states using a "composite standard" of the type proposed by the Regulatory Analysis Review Group (RARG) in the early days of the corporate average fuel economy (CAFE) standards.

Fourth, EPA should consider mechanisms that would tie the trading price for carbon credits to the domestic SCC, rather than to any fixed quantitative goal. By using a safety-valve feedback mechanism to stabilize the price at the SCC, the standard can act as a "last-ton tax"—with many of the desirable efficiency properties of a Pigovian tax on carbon, but one that works within the limits of EPA's current statutory authority.

This comment has two attachments, which will be discussed below. The first is the "Composite Standard" section of the RARG report of March 31, 1980. The second is the author's earlier comment on EPA's 2014 proposal for the Clean Power Plan (CPP). That 2014 comment goes into much greater detail on the economic theory of constraints on the intensive and extensive margins.

We also plan to file a separate comment on EPA's NPRM withdrawing the 2015 final CPP. Meanwhile, we recommend that any replacement rule take the following form.

I. A Constraint on Carbon Intensiveness

In its 2014-15 CPP rulemaking, EPA sent mixed signals about whether states needed to comply with a rate-based or a mass-based emissions standard, and suggested that, to some degree, the two were equivalent or interchangeable. In fact, however, the two types of standard have very different effects.

A rate-based standard puts an enforceable constraint on the *intensive margin*, limiting the pollution-intensiveness of an industrial activity. You can recognize this type of standard because it typically has both a numerator and a denominator: grams per mile, parts per million, or tons per megawatt-hour. The vast majority of EPA standards take this form, which puts an upper bound on the amount of pollution per unit of economic output.

In contrast, a mass-based standard puts the regulatory constraint on the *extensive margin*, imposing a fixed cap on the total amount of a pollutant that may be emitted. There is a rich economic literature exploring cap-and-trade proposals for regulating CO_2 emissions, but EPA actually has limited experience with this type of constraint. Some authors confuse the two, and describe the 1982 – 1987 trading of lead allowances in gasoline as cap-and-trade, for example; but that program actually used a grams-per-gallon constraint on the intensive margin.

In many ways a mass-based standard has more in common with economic regulation than it does with health, safety, and environmental regulation. It would require EPA (or the states) not only to establish a cap for CO_2 emissions, but also to decide who is entitled to use those emissions allowances. The scarcity of emissions allowances would produce both the desired substitution effect, a reduction in CO_2 emissions, and an output effect—a reduction in the quantity of electricity produced and a concomitant increase in the price of electricity. The result of the output effect is a substantial transfer of income from electricity consumers to the winners of the rent-seeking contest for the initial allocation of allowances. This competition for emissions allowances can incur rent-seeking costs that raise the total cost of a mass-based standard by an order of magnitude.

The distinguishing feature of a constraint on the intensive margin is that it creates incentives in two distinct dimensions. To come into compliance, regulated entities have an incentive to decrease pollution (the numerator) and a simultaneous incentive to increase output (the denominator). The net result is that there is no output effect to raise prices for consumers. Consumers will need to pay the cost for reducing CO_2 emissions, but will not be burdened with also paying for income transfers to the owners of capped emissions allowances. *(See the attached 2014 comment for a more detailed explanation of compensated demand curves, and the absence of an output effect.)*

Mass-based standards are the equivalent ("dual," in economists' parlance) of an emissions tax. EPA does not now have the authority to impose an emissions tax, and it is not clear that the agency has the authority to set mass-based standards. Any replacement rule for the Clean Power Plan should make it clear that it is strictly a rate-based standard—i.e., a constraint on the intensive margin.

II. Tiered by Technology

The 2015 CPP rule set state-by-state targets for CO_2 emissions, based on EPA's judgment about how the electricity market in each state ought to change. That is not the task given to EPA by the CAA. A CO_2 rule for stationary sources under the CAA should follow the same pattern that EPA has used for every other pollutant: it should limit the emissions intensiveness for each defined category of source.

Tiered standards on the intensive margin are not just a routine practice under the CAA; they are also used by DOE for appliance efficiency standards, by DOT for CAFE standards, and by HUD and other agencies for building energy performance standards. Such energy efficiency standards typically will use categories of technology that are distinguished by some aspect of their output or features that warrant a distinction. Refrigerator-freezers with a through-the-door icemaker may be given a greater energy budget than those that lack that feature. CAFE standards for trucks will not be set at the same level as for cars. End-unit townhouses will get a larger energy budget than those in the middle.

But technology categories can be distinguished by the nature of their inputs, as well as their outputs. In the case of carbon emissions from EGUs, all of the electricity looks the same on the output side, but different technologies use very different inputs to produce that electricity. The United States as a whole is blessed with a mixed and resilient energy resource base, whose distribution varies by geography. The electricity generation system within a state will look very different depending on whether that state is endowed with abundant coal resources, or hydropower, or natural gas, or wind, or wood, or insolation. It is not EPA's task to choose among these technologies, but to set carbon intensity standards that are achievable for each category.

Some commentators have described this as a "fenceline" issue, arguing that EPA must set standards that are achievable inside the fenceline of the regulated sources. Another way to look at the same issue is to think of technology categories in terms of the inputs (including capital) and outputs that they use. Coal-fired generation is a lawful practice. EPA can make it more carbon-efficient, but is not authorized to ban it in favor of some other preferred technology. This is not because the other technology is somehow outside the fence; rather, it is because it uses different inputs and outputs that have an inherently different carbon footprint.

III. With Trading of Carbon Offsets

The value of emissions trading is well established, and there is no need to elaborate on it here. It is worth emphasizing, however, that emissions trading is not something that works only within a cap-and-trade program; trading is perfectly compatible with the two features listed above: standards that constrain the intensive margin, and that are tiered by category of technology.

EPA's most successful emissions trading program was for tetraethyl lead gasoline additive. That program, adopted in 1982, held all refiners and importers of leaded gasoline to a standard of 1.1 grams/leaded-gallon—a constraint on the intensive margin. Refiners traded lead offsets, denominated in grams. Trading allocated lead additive efficiently while it lasted, and it enabled EPA to phase out leaded fuel altogether by 1987.

An early example of a tiered standard with trading across technology tiers is the National Highway Traffic Safety Administration's (NHTSA) CAFE standard for light trucks for model years 1983-85. President Carter's Regulatory Analysis Review Group (RARG) provided comments[3] to NHTSA Administrator Joan Claybrook in March 1980, proposing a "composite standard" for two-wheel drive and four-wheel drive trucks. NHTSA adopted a variant of this composite standard in its final rule, which was later superseded when Congress decided to set CAFE standards directly by statute.

It is worth looking back at the details of RARG's composite standard, and how it might be adapted to apply to EGUs. The problem NHTSA faced was that different manufacturers specialized in different kinds of trucks. Some produced only 4x4 trucks; others produced a mix that was dominated by 4x2 trucks. CAFE averaging (the equivalent of emissions trading) took place only within companies, not across companies. How could NHTSA set a standard that recognized the inherently different fuel economy of 4x4 trucks, and accommodated the different fleet mixes produced by the various manufacturers, and that at the same time provided the right incentives to improve fuel economy both within technology categories and across them?

---

[3] The relevant portion of the RARG report is attached to this comment. The full text of the RARG report is available at: http://cowps.mercatus.org/1980/03/31/national-highway-traffic-safety-administrations-proposed-light-truck-average-fuel-economy-standards-for-model-years-1983-1985/.

The RARG composite standard calculated a fuel economy target for each manufacturer based on the mix of vehicles in its fleet. This composite standard did not mandate that manufacturers change their fleet mix, but it did allow them credit for doing so. It provided an incentive both to improve the fuel economy within each technology category, and to improve the fuel economy of the fleet by shifting the mix in the direction of more 4x2 vehicles.

EPA could do something similar in setting state-by-state targets for reduced CO_2 emissions. Begin by establishing achievable CO_2 intensity standards for each category of EGU technology. Then set a composite target for each state by weighting those standards according to the mix of technologies used to generate power within the state. A state could meet its composite target by improving the efficiency at each of the EGUs within the state, but it would also get credit for shifting the generation mix in the direction of lower carbon emissions.

How does that differ from what EPA did in its 2015 Clean Power Plan? Primarily it differs by basing the composite standard on technology-based standards that are achievable. This would not force states to change their fleet mix, but would allow them to take credit for changes in the fleet mix, as well as for improvements in the CO_2 intensiveness within technology categories.

IV. With a Price Target Tied to the Social Cost of Carbon

It is impossible to know in advance the precise cost of achieving CO_2 reduction targets, since many other economic factors may cause the cost to vary over time and space. Yet we know that an efficient system of emissions reduction will have a uniform cost per ton of CO_2, regardless of where it is emitted. And a price that is stable over time will best enable states and utilities to plan for cost-effective emissions reduction measures.

For these reasons, EPA should consider mechanisms that use market feedback to adjust targets, with the aim of getting the price of tradeable CO2 offsets to stabilize at close to the domestic social cost of carbon. This can be done with a "safety valve" feature that automatically adjusts the intensity standards to prevent unanticipated price spikes. Ideally, such a mechanism will stabilize allowance prices across time and space, and at the same time prevent large transfers across state lines and across technology categories. In effect, the agency would be adjusting the shadow price of the various emissions standards in order to maintain consistency with the SCC.

Conclusion: "The Last-Ton Tax"

The result of all of the above recommendations might be called a "last-ton tax." Every EGU will face a marginal cost of CO_2 emissions that is equal to the SCC. States will have an incentive (but not an obligation) to shift the mix of generation technologies in the direction of lower carbon emissions. At the same time, the last-ton tax avoids imposing an inframarginal tax on those CO_2 emissions that are unavoidable. It thereby avoids large transfers of income from consumers to pay a CO_2 scarcity rent, and avoids the wasteful effects of a rent-seeking contest to obtain those rents.

The safety valve mechanism allows a national market to develop for trading CO_2 offsets, but uses market feedback to adjust the emissions targets so that prices are stable and the rule does not inadvertently cause large transfers between states or between categories of EGU technology.

This implementation of the last-ton tax would encourage each technology to come into equilibrium with the SCC, but it would not necessarily cause the mix of technologies to come into equilibrium (at least not quickly), because the tiered standard compensates for the inherently higher carbon footprint of fossil-fueled generation. The last-ton tax goes as far as possible towards the desirable efficiency properties of a Pigovian tax on carbon, within the limits of EPA's current statutory authority.

Attachments:

> *Excerpt from the RARG report of March 31, 1980, explaining the structure of a composite standard.*
>
> *Public Interest Comment on the EPA's 2014 NPRM on the Clean Power Plan.*

EXECUTIVE OFFICE OF THE PRESIDENT
COUNCIL ON WAGE AND PRICE STABILITY
WINDER BUILDING, 600 - 17TH STREET, NW.
WASHINGTON, D.C. 20506

March 31, 1980

Honorable Joan Claybrook
Administrator
National Highway Traffic Safety
Administration
Washington, D.C. 20590

Dear Ms. Claybrook:

In my letter to you of February 25, 1980, I identified the concerns on which the Regulatory Analysis Review Group (RARG) was focusing in its review of the National Highway Traffic Safety Administration's proposed Light Truck Fuel Economy Standards for Model Years 1983-85. The RARG now has completed that review, and the outcome is the enclosed report which I request be placed in the public record for this proceeding.

Sincerely,

R. Robert Russell
Director

cc: Members of Regulatory Analysis Review Group

Docket Section
National Highway Traffic Safety
Administration
Docket No. FE78-01; Notice 1
Room 5108
400 Seventh Street, S.W.
Washington, DC 20590

National Highway Traffic Safety Administration's

Proposed Light Truck Average Fuel Economy

Standards for Model Years 1983-1985

Report of the

Regulatory Analysis Review Group

Council on Wage and Price Stability

March 31, 1980

B. AN ALTERNATIVE APPROACH

The Composite Standard

The alternative approach explored here involves setting fuel economy targets for different categories of trucks, and using a pre-determined fleet mix for each manufacturer to turn these targets into a composite standard. To illustrate how such a standard would be set, Table IV converts NHTSA's base case standards for 4X2 and 4X4 trucks into a single composite standard for each manufacturer.

Table IV

Calculating a Composite Standard for each manufacturer

		4X2		4X4		Composite
		mpg	Fraction	mpg	Fraction	Standard (mp
1983:	GM	18.0	.70	15.6	.30	17.21
	Ford	"	.74	"	.26	17.31
	Chrysler	"	.74	"	.26	17.31
	AM	"	0	"	1.00	15.60
	IH	"	0	"	1.00	15.60
1984:	GM	18.70	.70	16.1	.30	17.90
	Ford	"	.74	"	.26	18.01
	Chrysler	"	.74	"	.26	18.01
	AM	"	0	"	1.00	16.10
	IH	"	0	"	1.00	16.10
1985:	GM	19.70	.70	16.2	.30	18.50
	Ford	"	.74	"	.26	18.65
	Chrysler	"	.74	"	.26	18.65
	AM	"	0	"	1.00	16.20
	IH	"	0	"	1.00	16.20

Note: The fleet mix fractions are those estimated by NHTSA for the 1983 model year. The composite was calculated as an harmonic average, i.e., the inverse of the composite is the weighted average of the inverses of the two separate standards.

These composite standards should achieve about the same fuel savings as the separate base-case standards. However, they give manufacturers both the opportunity and the incentive to meet the standard by replacing 4X4's with 4X2's. Moreover, using the composite approach, NHTSA could establish a richer set of performance categories than simply these two, without removing incentives to shift the fleet mix. To implement the composite approach, NHTSA would use truck categories only as a tool for calculating the fuel economy standards from category by category fuel-economy targets. The composite standard itself would, in the usage of the statute, be classless. 1/

With appropriate legislation, the fuel economy standards for cars and trucks might be combined into an overall composite standard. This would add more flexibility for finding the least-cost method of saving fuel. Setting separate standards for cars and truc could conceivably reduce overall fuel economy, by inducing a shift in demand away from cars towards less efficient trucks, although market responses to higher fuel prices seem to have blunted this shift. In any event, a composite car/truck standard would give manufacturers proper incentives to avoid such shifts. NHTSA should explore this option, although it is not available under the current law.

1/ Nothing in the statute forbids this approach. The statute requires that passenger car standards be the same for all manufacturers. There is no similar requirement for the truck standards. Indeed, the statute explicitly authorizes separate standards for different classes of trucks, which would inevitably result in varying effects on the different manufacturers. Since this is explicitly permitted, it seems unlikely that composite standards, which would result in similarly varying effects are forbidden. NHTSA's treatment of this issue in the preamble to its final truck standards for model years 1980-81 suggests that it agree 43 FR 11997-8. There, NHTSA discussed a proposed fleet-average stan dard at some length - eventually rejecting it on policy grounds - without suggesting that it might be illegal.

RARG recognizes that proposing a composite standard for each truck manufacturer raises several potential difficulties, including possible effects on the competitive structure of the industry. We are not prepared to recommend that NHTSA adopt this modification since a full assessment will require considerable additional analysis. We do recommend that NHTSA consider the approach and, in its final regulatory analysis, address the strengths and weaknesses of the approach. In the remainder of this Section we summarize some of the advantages and potential difficulties of the scheme.

Advantages of a Composite Standard

The fundamental advantage of composite standards is that they give manufacturers greater flexibility to meet given fuel-economy objectives at lowest cost. Manufacturers have greater discretion to choose which classes of vehicles to improve, and whether to increase production of the more fuel-efficient classes. This flexibility may lower costs for three major reasons. First, the manufacturers may later acquire better information about the costs and effectiveness of fuel-saving options than is available to NHTSA at the time the standards are set. Second, the least cost mix of improvements may vary from manufacturer to manufacturer, making it impossible for NHTSA to set class-specific standards that are optimal for every manufacturer. Third, composite standards allow manufacturers to improve fuel economy by substituting more fuel-efficient classes of trucks for less fuel-efficient classes.

To illustrate this last possibility, suppose that the 4X2 and 4X4 standards are 18 and 16 mpg respectively. Further suppose that, at these levels, the resource cost for an additional 1 mpg improvement in either class is $50 per car. Under class specific standards, the manufacturer will spend up to $50 per mpg to improve the fuel economy of any particular vehicle, but will forego an opportunity to improve average fuel economy that may be much cheaper -- replacing some 4X4's with 4X2's. Each such replacement generates a 2 mile per gallon improvement. For small changes, this improvement is virtually costless because at equilibrium prices and quantities, the manufacturer and consumer both are indifferent to small shifts in the relative production of each type of vehicle. The cost -- measured by consumers' losses from inferior truck attributes -- will increase as the production shifts away from the prior equilibrium mix of 4X4 and 4X2 trucks. The composite standard will give the manufacturers the incentive to substitute among classes to the same degree that they exploit every other opportunity to improve fuel economy. 1/ Thus, the composite standard may better encourage achieving mandated fuel savings at least cost, where cost includes both resource cost and changes in consumer valuation of attributes.

1/ There is a possible complication in comparing the resource costs of technological modifications with the costs resulting from mix shifts toward truck attributes that are less desirable to consumers. Manufacturers will treat as costs, the reduction in profits from shifting from the optimum number of, say, 4X4 trucks. But profit change may not accurately reflect the welfare cost of the shift. The bias depends upon the circumstances of demand and costs facing the manufacturer, and on the manufacturer's market power; profit changes will overstate welfare losses under some circumstances and understate welfare losses in others. We have no reason to expect, however, that these distortions will af t the relative cost-minimization advantages of a composite standard cc pared to a class-by-class approach.

The composite standard may also give NHTSA the freedom to use a richer set of classes to deal with the problem of setting standards that are achievable by the manufacturer least capable of making fuel economy improvements. For example, in the 1980-81 rule-making, IH requested that NHTSA establish a separate class for trucks with a gross vehicle weight rating over 6000 pounds to accommodate its greater proportion of heavy trucks. The agency declined, noting that establishing such a separate class would remove any incentive for IH to improve fuel economy by reducing the GVWR of its trucks below 6000 pounds. 1/ Instead, NHTSA found other grounds for establishing a special IH class. But under a composite standard, the IH proposal would have had fewer drawbacks. A composite that included a class for heavier trucks would accommodate the IH fleet, without removing the incentive to downsize it. As discussed avove, the criteria for designating additional classes should be performance-related differences in potential fuel economy and different representation in different fleets.

1/ 43 FR 11997-8.

Difficulties with the composite standard

1. Establishing fleet mix assumptions

An essential feature of the proposed composite standards is that the assumed fleet mix for each manufacturer must be established before the model year to which the standards apply. 1/ Projecting fleet mix accurately can be quite difficult, however, and embedding such projections in composite standards may be inadvisable. For example, projecting fleet mixes might require NHTSA to make quite subjective judgments about manufacturers intended marketing plans. To avoid having to make projections, NHTSA might consider basing each year's composite standard on a lagged fleet mix; for example, a manufacturer's 1985 standard might be based on its 1983 fleet mix. A two or three year lag would allow the standards to be continually and automatically adjusted as manufacturers changed their fleet composit .

While using a lagged fleet mix assumption will give an incentive to produce more fuel efficient classes, the incentive may be partly diluted by the manufacturers' recognition that this year's change in fleet mix will tighten a future year's standard. There may also be other drawbacks to the use of lagged data, such as opportunities for dynamic "gaming" by manufacturers and limitations on the manufacturers' flexibility to respond to abrupt changes in market demand for truck types. We urge NHTSA to explore the advantages and disadvantages of alternative methods of specifying fleet mix.

1/ The fleet mix assumption must be established in advance so that manufacturers have a fixed benchmark against which improvements in fuel economy can be measured. In contrast, if the composite stan d were based on the actual fleet mix of the model year, the manuturer would not have an incentive to increase its mix of fuel-efficient trucks.

2. Competitive effects

A composite standard is likely to change the relative prices of truck types, compared to prices under a class-specific standard. Manufacturers will tend to increase the price of trucks with poor fuel economy, and lower the price of fuel-efficient trucks as they meet the composite standard.

While this change in the relative prices of different truck classes should result in consumer choices more consistent with overall national objectives, the change may also have unequal impacts on different manufacturers. For instance, the profitability of companies that specialize in 4X4s could increase compared to those who produce mostly 4X2s. On the other hand, the increased felxibility of the composite standard may benefit primarily those manufacturers with the most diversified fleets, or those with inadequate capital to improve all models simultaneously. These differential firm effects could influence the structure of the industry, and deserve careful examination.

Another difficulty under composite standards is that they may impose an unfair burden on a manufacturer which decides to increase its mix of a less fuel-efficient class of vehicle. While the market as a whole may be shifting toward the more fuel-efficient classes, there is no reason to force each manufacturer to do so. The composite standard would make it very difficult for a manufacturer to, say, terminate his production of 4X2s while continuing to make 4X4s. For this reason, NHTSA might preserve the option of meeting class specific standards as an alternative means of compliance.

Providing both options may, however, have its own drawbacks. For instance, if two manufacturers started to specialize in different classes, one of them could get credit for increasing the fuel efficiency of its fleet mix, while the other selected the class-standard option and avoided being charged for decreasing the efficiency of its mix.

THE GEORGE WASHINGTON UNIVERSITY

WASHINGTON, DC

Public Interest Comment[1] on

The Environmental Protection Agency's Proposed Rule:

Carbon Pollution Emission Guidelines for Existing Stationary Sources – Electric Utility Generating Units

Docket ID No. EPA-HQ-OAR-2013-0602

RIN: 2060-AR33

December 1, 2014

Brian F. Mannix

Visiting Scholar[2]

The George Washington University Regulatory Studies Center

The George Washington University Regulatory Studies Center works to improve regulatory policy through research, education, and outreach. As part of its mission, the Center conducts careful and independent analyses to assess rulemaking proposals from the perspective of the public interest. This comment on the Environmental Protection Agency's proposed rule setting carbon emissions guidelines for electric generating units (EGUs) does not represent the views of any particular affected party or special interest, but is designed to evaluate the efficiency consequences of EPA's choice of an intensity constraint for carbon emissions, and to give guidance to states on the relative advantages and disadvantages of this type of constraint.

[1] This comment reflects the views of the author, and does not represent an official position of the GW Regulatory Studies Center or the George Washington University. The Center's policy on research integrity is available at http://regulatorystudies.columbian.gwu.edu/policy-research-integrity.

[2] Brian F. Mannix served (2005 – 2009) as EPA's Associate Administrator for Policy, Economics, and Innovation. He is currently a visiting scholar with the George Washington University Regulatory Studies Center, 805 21st St. NW, Suite 609, Washington, DC. He can be reached at BMannix@gwu.edu or (202) 994-7543.

Introduction

EPA's proposed rule sets state-by-state carbon intensity targets for the production of electricity. States are expected to adopt some form of economic incentive regulatory system to achieve these targets, but there has been a great deal of confusion about how, exactly, such a system should work. Several states have asked EPA how to translate these "rate-based" intensity targets into equivalent "mass-based" targets that could form the basis of a cap-and-trade or similar system; EPA recently responded to these questions with a supplementary notice.[3]

This comment will argue that it would be a serious mistake for states to convert intensity goals to mass-based goals. Economic theory suggests that the costs of achieving emissions reduction using a mass-based control system will be an order of magnitude more expensive than achieving the same reduction with a rate-based system, when the costs of rent-seeking (which a rate-based system can better resist) are taken into account. Moreover, states that adopt a mass-based system place themselves at a severe competitive disadvantage, not only with respect to other states, but also with respect to foreign jurisdictions that adopt a rate-based target or no target at all. Finally, a mass-based system of emissions control would have regressive distributional effects that can easily be avoided with a rate-based system of control.

Historically, emissions trading has been most successful when designed to achieve an intensity goal. For example, between 1982 and 1987 EPA used an emissions trading system to phase out the use of tetraethyl lead as an octane booster in gasoline.[4] Other countries followed, and the United Nations Environment Program recently estimated that the global benefit of removing lead from gasoline now amounts to at $2.4 trillion per year.[5] A key factor leading to the success of this effort was EPA's deliberate use of an intensity target, rather than a mass-based target.

The economic literature on emissions trading contains confusing and contradictory discussions about the relationship between *constraints on the intensive margin* (intensity constraints) and *constraints on the extensive margin* (mass-based constraints). This comment, which will be most accessible to economists, is intended to provide a structure for thinking about these options and their advantages and disadvantages.

This comment will *not* address many other important questions raised by EPA's proposal: e.g., climate science, the estimation of the Social Cost of Carbon, EPA's legal authority for the proposed rule, or the calculation of state-by-state targets. Regardless of the disposition of this particular rulemaking, any attempt to regulate carbon emissions – particularly at the state level – will require that policy makers understand the dynamics of emissions trading and the implications of working under a constraint on the intensive or extensive margin.

---

[3] Notice; Additional Information Regarding The Translation Of Emission Rate Based Co2 Goals To Mass Based Equivalents, https://www.federalregister.gov/articles/2014/11/13/2014-26900/carbon-pollution-emission-guidelines-for-existing-stationary-sources-electric-utility-generating.

[4] The lead trading system was developed by the author at the Office and Management and Budget in 1981-82, and implemented by EPA at OMB's request.

[5] http://www.unep.org/newscentre/default.aspx?DocumentID=2656&ArticleID=8917

Options for Reducing Emissions:

In order to illuminate the critical policy choices in designing regulatory instruments for reducing carbon emissions, and to explore their economic implications, this section will review six distinct control options:

1. A Carbon Tax
2. An Auctioned Cap-and-Trade System
3. A Rebated Carbon Tax
4. A Rebated Cap-and-Trade System
5. An Output-Compensated Carbon Tax
6. An Output-Compensated Emissions Trading System

Economic theory tells us that the first two options, which at some level are equivalent ("dual"[6]) to each other, will raise the effective price of carbon emissions and will thereby reduce the level of emissions, which will follow an ordinary "Marshallian" demand curve. The second pair of options are also mutually dual, but will reduce emissions by following an income-compensated demand curve. The final pair of options are mutually dual and follow an output-compensated demand curve.

Option 6 is simply emissions trading with an intensity constraint, and it has multiple advantages over other options. To understand both its advantages and disadvantages, however, we will first have to explain where those different demand curves come from and what they mean.

1. A Carbon Tax

There is an extensive literature regarding the efficiency advantages of using a price instrument – a Pigovian emissions tax – to "internalize" the externalities associated with air pollution. In many ways CO_2 emissions, if they are indeed harmful, would be an ideal candidate for an emissions tax set equal to the external "Social Cost of Carbon" (SCC). As the author has argued elsewhere: "The SCC may appear to be a gross oversimplification of a complex underlying reality; but, in fact, it is the right simplification to undertake. This is because any damage that greenhouse gas emissions may inflict on global climate systems is independent of the source of the emissions. To the climate, all CO2 molecules look the same. . . [A]ny cost-effective portfolio of climate policies will have a single implicit marginal cost of carbon."[7]

[6] http://en.wikipedia.org/wiki/Duality_%28optimization%29

[7] Brian Mannix and Susan Dudley, *Public Interest Comment on The Interagency Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis under Executive Order No. 12866*., 26 February 2014.

https://regulatorystudies.columbian.gwu.edu/sites/regulatorystudies.columbian.gwu.edu/files/downloads/OMB_2013-0007_SCC.pdf

In addition to its theoretical advantages, however, a carbon tax has some major drawbacks. As illustrated in the drawing below, a carbon tax would impose two distinct costs on consumers[8] of electricity: the real resource costs of eliminating some fraction of emissions, plus the cost of paying the tax on those emissions that are *not* eliminated. Typically, the cost of paying the tax – that is, the tax revenues – will be several times larger than the cost of reducing emissions. Economists will usually treat this tax as a "transfer payment" to the government rather than a real resource cost to the economy. Nonetheless, from consumers' perspective, both of these costs will cause an increase in the price of electricity. A tax on carbon emissions from EGUs will effectively translate into a tax on electricity, which is likely to be very unpopular, in part because it is regressive compared to other sources of state revenue. Such a tax would also render a state much less competitive in attracting businesses that use electricity. Moreover, to the extent that an EGU carbon tax drove electricity-using businesses to other jurisdictions, the purpose of the tax will have been defeated. Carbon emissions can effectively flee the tax, causing a rebound effect. Thus a carbon tax is particularly ill-suited for use at the state level, because states have few options for avoiding that counterproductive outcome. In Figure 1, below, the downward sloping line is an ordinary Marshallian demand curve that shows how the level of carbon emissions will respond to a price, or Pigovian tax.

Fig. 1: Effect of a Tax on Emissions

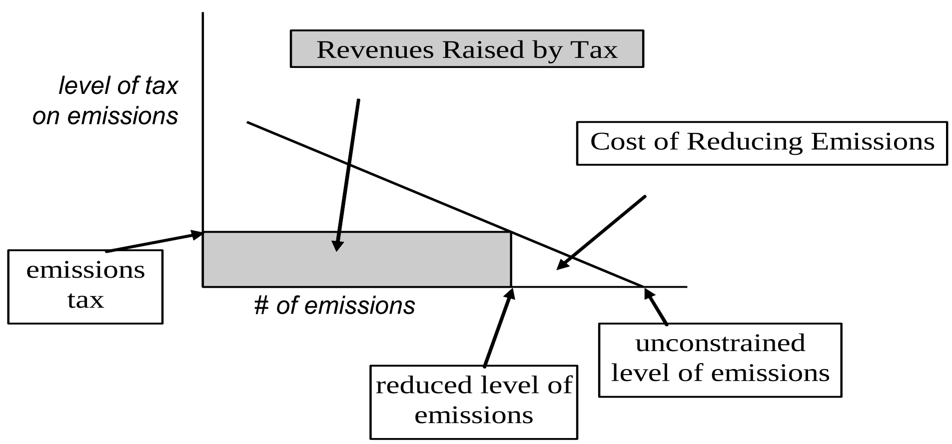

Before leaving the carbon tax, we need to make note of three distinct ways that it will cause emissions to be reduced below the level of unconstrained (untaxed) emissions. First, it causes a *technical substitution effect*, as electricity producers substitute lower carbon-intensity technologies for higher ones. Second, it causes a *consumer substitution effect* by raising the price of electricity, thereby inducing consumers to buy more of other goods and less of electricity.

[8] A carbon tax, or almost any other method of controlling carbon emissions, will impose costs whose incidence will be felt both downstream of the EGUs (by residential and commercial electricity consumers) and upstream (by owners of carbon-intensive infrastructure and fossil fuel deposits, especially coal-fired EGUs and coal mines). The majority of the costs, in the long run, are likely to fall downstream; hence our focus here will be on consumer impacts. This will help simplify the diagrams, which will show various demand curves for emissions allowances, but policy makers should understand that the actual incidence of costs is more complex.

Third, it causes a *consumer income effect*, as consumers find that higher electricity prices have left them less able to afford everything, including electricity. The latter two effects are sometimes referred to in combination as the *output effect*, since both of them result in lower output of electricity. In the illustration below, both components of the output effect are exaggerated for clarity. In practice, the technical substitution effect will dominate, and the income effect will likely be very small, in terms of its effect on carbon emissions. Of course, the effect of a carbon tax on consumer incomes is likely to be quite large, in terms of political acceptability, but that is another matter. The downward sloping line in Figure 2 is an ordinary factor-market demand curve (because carbon emissions are a factor of production for electricity), and these three effects are standard features of any such demand curve.

Figure 2: Three mechanisms by which a tax will reduce emissions

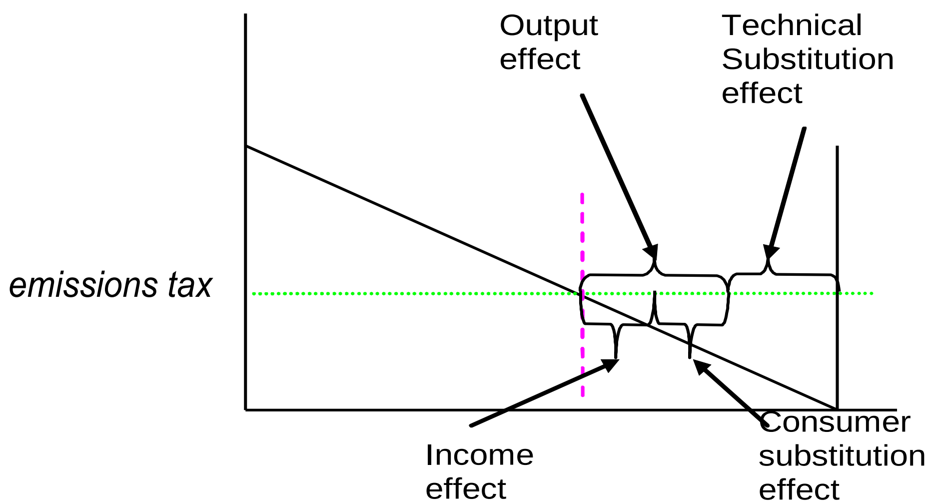

2. A Cap-and-Trade System

A cap-and-trade system, with the allowances auctioned by the government, is the dual equivalent of a Pigovian emissions tax. As such, its effect on emissions can be described by the same factor-market demand curve we saw above. In the diagram below, the shaded rectangle, which had represented the revenue raised by a carbon tax, now has been relabeled a "regulatory scarcity rent," and it represents the market value of the outstanding emissions allowances. This is sometimes called a "Tullock rectangle," named for Gordon Tullock, who died on November 4, 2014. Tullock pointed out that the creation, by regulation, of a scarcity rent, will inevitably touch off a rent-seeking contest to capture that value. It will take place in political, administrative, and judicial arenas, and – to a first approximation – will likely continue up to the point where the entire scarcity rent has been consumed. Thus in practice the Tullock rectangle represents, not a transfer that can be ignored, but a real resource cost that substantially reduces the efficiency of the regulatory system. Of course, taxes are not immune from rent-seeking, so a carbon tax would also likely induce some of these same costs. But, with a tax, there is some expectation that at least a portion of the revenue will accrue to the treasury. With a cap-and-trade system there may be no such expectation, making them particularly susceptible to rent-seeking.

Apart from that, a cap-and-trade system will suffer from many of the same drawbacks as a carbon tax: it will impose an unnecessarily large burden on consumers, and it will substantially impair the competitiveness of any jurisdiction that adopts it.

Figure 3: Demand for Emissions Allowances under Cap-and-Trade

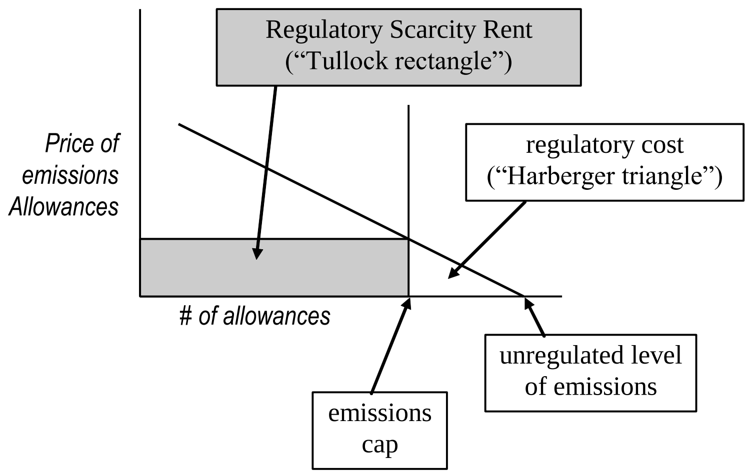

3: A Rebated Carbon Tax

In order to ameliorate the consumer impact of a carbon tax, some have proposed a rebated carbon tax, also known as an "income-compensated" or a "revenue-neutral" carbon tax. The rectangle of revenue that was collected by the tax would be returned to the public in the form of a cut in the income tax or sales tax, so that the net change in the public tax burden would be zero. One difficulty is that the rebate cannot easily be targeted on the same consumers who are bearing the incidence of the carbon tax. Another is that dramatic changes in tax policy have uncertain outcomes.

Note that, if successfully adopted, the effect of an income-compensated carbon tax would be to eliminate the (very small) consumer income effect on carbon emissions, so that emissions would be slightly higher than they would be under an uncompensated carbon tax. In effect, the level of carbon emissions would not follow the black line – what economists call an ordinary or "Marshallian" demand curve. Instead, it would follow the slightly higher red line – what economists call an income-compensated demand curve. It reflects the technical substitution effect and the consumer substitution effect, but the consumer income effect (on carbon emissions) has been eliminated.

Figure 4: Deriving an Income-Compensated Demand Curve

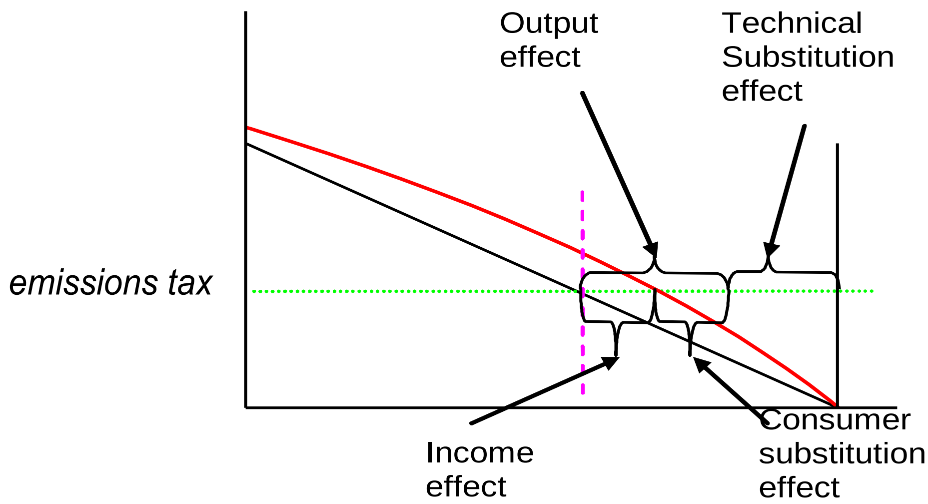

4: A Rebated Cap-and-Trade

We can imagine doing something similar with a cap-and-trade system. The emissions allowances would be auctioned, and the resulting revenues returned to consumers by cutting other taxes. In practice, it would not likely work out that way. Various interests would stake a claim on the revenues, or would lobby for a set-aside of free allowances to be allocated by some political or administrative process. Again, by creating a scarcity rent, cap-and-trade systems tend to stimulate costly rent-seeking. But, assuming that could be avoided, a rebated cap-and-trade system would work very much like a rebated carbon tax, and would follow the same income-compensated demand curve.

5: An Output-Compensated Carbon Tax

There is an alternative means of returning the revenues from a carbon tax to the consumers who paid it. The revenues can be placed in a fund which is used to subsidize electricity production. Every EGU would pay a tax on carbon emissions, and receive a subsidy on electricity output. Those with relatively high carbon emissions would be net payers into the fund; those with relatively low carbon emissions would be net recipients. Overall, the net revenues to the fund – and the net tax burden passed on to consumers in the price of electricity – would be zero.

This tax/subsidy system has been used to manage conventional NOx emissions from EGUs and other large point sources in Sweden. It has the virtue of targeting the benefit of the subsidy on exactly the same consumers who bear the burden of the tax.[9] In enacting its system of refunded emissions payments, Sweden was conscious of the fact that output compensation would

[9] Note that, in order to work as described, the tax must be on current carbon emissions and the subsidy must be on current electricity production. Various proposals to use "historical" emissions as the basis of allocating emissions allowances will not work the same way, since historical emissions are completely inelastic.

substantially mitigate the damage to the nation's competitiveness that otherwise would have been incurred as a result of an emissions tax. Sweden was also deliberately seeking to create incentives to produce a technical substitution effect, rather than a change in consumers' lifestyles or market baskets.[10]

As seen in Figure 5, the effect of an output-compensated carbon tax is limited to the technical substitution effect. It will not produce either a consumer income effect nor a consumer substitution effect. Emissions under an output-compensated carbon tax will follow the blue line – an output-compensated demand curve. Note that, because it uses only the technical substitution effect, an output compensated carbon tax will produce slightly fewer emissions reductions than an uncompensated tax at the same price per ton of carbon. Still, from the consumer's perspective, the costs of a compensated tax will be dramatically lower.

Figure 5: Deriving an Output-Compensated Demand Curve

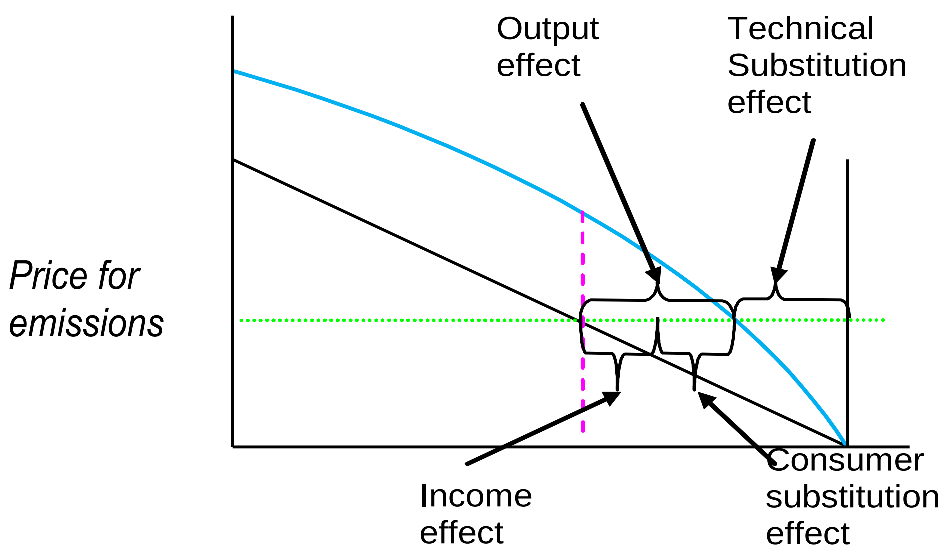

6. An Output-Compensated Emissions Trading System

With this last option, things get simpler rather than more complicated. It turns out that an output-compensated emissions trading system is simply emissions trading combined with an intensity constraint. States can simply require that all covered sources comply with the same carbon intensity constraint, denominated in tons per megawatt-hour, and can allow trading of carbon allowances among them. This is sometimes called an "offset market." More carbon-intensive sources will need to buy allowances from sources that are less carbon-intensive. As long as no one cheats, the overall system will meet the intensity goal, and trading will allow participants to

[10] For an excellent discussion of how this system works, see Thomas Sterner and Bruno Turnheim, "Innovation and Diffusion of Environmental Technology: Industrial NOx Abatement in Sweden under Refunded Emission Payments," http://www.rff.org/Documents/RFF-DP-08-02.pdf

find the least-cost means of doing so. There is no fixed pool of allowances to be allocated, and no central authority need collect any tax or pay any subsidy; as a result this system is very resistant to rent-seeking.

So, for example, when EPA set the intensity constraint on lead in gasoline at 1.1 grams of lead per gallon of leaded gasoline in 1982, refiners began trading lead allowances at a price that fluctuated around 2 cents per gram. This automatically translated into a subsidy on leaded gasoline of around 2.2 cents per gallon, because each gallon produced would earn 1.1 grams of lead allowances. The price at the pump incorporated both the effect of the lead "tax" and the effect of the gasoline "subsidy," which exactly offset each other. Other than enforcing the intensity constraint and monitoring the trading, there was little that EPA needed to do. There was no pool of allowances to be allocated, and no fund to collect revenues. Rent seeking in the lead phasedown program, which had been rampant prior to 1982, virtually vanished. Within five years, with little resistance, EPA was able to phase out lead use almost entirely.

Of course, emissions trading under an intensity constraint will produce only a technical substitution effect. It will not, for example, do much to encourage electricity conservation by consumers.[11] But its advantages, in terms of economic competitiveness, distributional effects, and especially resistance to rent-seeking, are substantial.

Putting it All Together

The chart below summarizes the relationships among the six options we have been discussing.

Constraint is applied to . . .	Extensive margin w/ No Compensation	Extensive Margin w/ Income Compensation	Intensive Margin w/ Output Compensation
Price instrument:	1. Carbon Tax	3. Income-compensated carbon tax	5. Output-compensated carbon tax (like Sweden)
Quantity instrument:	2. Cap & Trade	4. Cap & Trade w/ rebate	6. carbon trading under an intensity constraint (like lead phasedown)
Reduces emissions by . . .	Technical substitution effect; consumer substitution effect; income effect	Technical substitution effect; consumer substitution effect	Technical substitution effect only
Carbon emissions follow a . . .	Marshallian demand curve	Income-compensated demand curve	Output-compensated demand curve
Revenue/scarcity rents accrue to . . .	Revenue to the treasury or the distributees	Revenue may be returned to consumers via tax cuts or coupons	Revenue is retained by consumers

[11]There is an argument that, by producing only a technical substitution effect, an intensity based standard is more compatible with EPA's legal authority under Section 111(d). There is little evidence that Congress intended this section of the Clean Air Act to go beyond the use of technology to reduce emissions, to be used to penalize consumers, or to induce changes in consumer lifestyle. Such legal arguments are beyond the scope of this paper, however.

The one disadvantage of Option 6 (emissions trading under an intensity constraint) is that it is limited to the technical substitution effect, and therefore will not achieve emission reductions associated with reductions in electricity demand – reductions that theoretically could be achieved by a simple unrebated carbon tax and that might be economically justified. On the other hand, it has substantial advantages in that it achieves substantial emission reductions while simultaneously minimizing the increase in the price of electricity by avoiding any net tax burden. This avoids undesirable consumer impacts, and also minimizes the damage to the economic competitiveness of any jurisdiction using this system.

Finally, trading under a uniform intensity constraint has the advantage of providing a very small attack surface for rent-seekers. It is awkward for any government agency's economic analysis to provide a frank discussion of the costs of rent seeking, but they are very real and very large. Consumer advocates, especially, ought to be alert to the danger. If a state converts an intensity target into a mass-based target, with the resulting pool of carbon allowances subject to allocation via some political/administrative process, the effect on consumers will be dramatic. The costs to consumers will rise by an order of magnitude, generating revenues that will fund a massive redistribution of income to the various groups that will attempt to claim some share of the allocation.

Commercial consumers of electricity, too, should recognize that a state moving from an intensity-based constraint to a mass-based constraint will be unlikely to stay competitive with states (or with foreign jurisdictions) that choose not to embed a carbon tax in their electricity prices.

One other complication needs to be addressed. We know that an efficient Pigovian tax on carbon would be set equal to the Social Cost of Carbon, assuming the latter is correctly calculated. If a state is instead using emissions trading under an intensity constraint, how can it know the efficient level of the constraint? The short answer is, it cannot know, and neither can the EPA. Even if EPA has done a creditable job of calculating reasonable targets for the short term, they will quickly become obsolete as circumstances evolve. It is not possible to calculate the "right" targets for the states, regardless of whether the targets are expressed in terms of intensity or total emissions. The same is true in international negotiations for a global system of carbon controls.

What is possible, both for states and for nations, is to come to some agreement about the target price for carbon allowances, and to manage their emission control systems to converge on that target price. For states that use an intensity constraint, that means observing the price at which carbon allowances trade, and intervening in the market if it rises above a target level (the Social Cost of Carbon, properly calculated). There are various mechanisms to do this. A "safety valve" mechanism can, in the short term, sell additional carbon allowances into the market to prevent a price spike. In the longer term the state can adjust its intensity constraint so that the market price for carbon settles at an appropriate level, without excessive reliance on the safety valve mechanism. As markets evolve, states can compare the price of carbon across different jurisdictions, and seek adjustments in their targets as needed. Ultimately, a national (or international) system of emission control needs to coordinated around a common price for carbon, rather than on any central authority doling out quantitative targets.

Appendix for Economists[12]

For those who are interested, this Appendix gives a little more formal detail on the arguments outlined above, and on the nature of the duality between Options 5 and 6.

The Slutsky equation, & three income-compensated demand curves

Ordinarily we use the Slutsky equation to decompose the price elasticity of demand for a consumer good ($\varepsilon_{electricity}$) into a consumer substitution effect and a consumer income effect.

$$\varepsilon_{electricity} = \varepsilon_{cons\text{-}substitution} + (PQ/I)\ \varepsilon_{income}$$

ε_{income} is the income elasticity of demand for electricity, and the weighting factor (PQ/I) is simply the fraction of income (I) devoted to purchasing Q of this good at price P.

The ordinary (Marshallian) demand curve shows the relationship between quantity and price, and reflects both effects. In contrast, an "income compensated" demand curve is one that reflects only the substitution effect. Conceptually, this involves removing the income effect by compensating the consumer for the income loss that is implicit in a price increase – or, in our case, a tax increase.

But there is an ambiguity in defining how much the consumer should be compensated. Hicks proposed compensation that would exactly preserve the consumer's welfare – giving what is known as the Hicks compensated demand curve, which (for the diagrams in this paper, showing the effect of tax increases) lies above the Marshallian demand curve. Consumer welfare is not observable, however; so Slutsky instead proposed compensating the consumer to the point where the initial market basket could still have been purchased. This is compensation using a Laspeyre index of price changes, and it generates the Slutsky compensated demand curve, which lies above the Hicks compensated demand curve. An alternative (and less generous) measure of compensation uses a Paasche index of price changes, producing a compensated demand curve that lies below the Hicks compensated demand, but still above the Marshallian. The Laspeyre indexed and the Paasche indexed compensated demand curves effectively give an upper and lower bound on the locus of the Hicks curve; the region in between corresponds to Samuelson's "zone of darkness" in which we cannot really be certain of the sign of the consumer welfare effect.

When dealing with infinitesimal price changes (e.g., in the Slutsky equation above) it makes little difference which flavor of compensated demand curve we choose, since they converge for small price changes. When designing policies to reduce pollutants, however, we are not really interested in infinitesimal changes. Hence it is important to identify which curve we are dealing with. For reasons that will become clear in the next section, we are interested in the third type –

[12] For simplicity, throughout this Appendix I will continue to ignore supply side effects, and will assume that consumers bear the incidence of upstream taxes.

the Paasche indexed compensated demand curve, despite the fact that it is far more obscure in the literature than the more familiar Hicks and Slutsky variants.

The *extended* Slutsky equation, & three output-compensated demand curves

When a price increase or a tax applies to a factor of production, such as carbon, rather than a consumer good, we can use the factor-market Slutsky equation (analogous to the consumer-good Slutsky equation) to decompose the price elasticity of demand (ε_{carbon}) into a technical substitution effect and an output effect.

$$\varepsilon_{carbon} = \varepsilon_{tech\text{-}substitution} + (TE/PQ)\ \varepsilon_{electricity}$$

Here the weighting factor (TE/PQ) is simply the fraction of the firm's revenue (PQ) consumed by the tax T on emissions E. But we already know from the section above that $\varepsilon_{electricity}$ can be expanded further. Combining the two equations, we get:

$$\varepsilon_{carbon} = \varepsilon_{tech\text{-}substitution} + (TE/PQ)\ \varepsilon_{cons\text{-}substitution} + (TE/I)\ \varepsilon_{income}$$

This extended form of the Slutsky equation shows how the price elasticity of demand for a factor of production (in this case, carbon allowances) can be decomposed into a relatively large technical substitution effect, a smaller consumer substitution effect (weighted by tax revenues as a fraction of industry revenue), and a still smaller consumer income effect (weighted by tax revenues as a fraction of consumer income).

The Marshallian demand curve for carbon emissions will be a function of all three effects. The output-compensated demand curve will reflect only the technical substitution effect. Again, when we move away from infinitesimals, we can distinguish three different flavors of output-compensated demand curve, depending on which type of index we use for determining the level of compensation. The uppermost of the three output-compensated demand curves uses Laspeyre index compensation. The middle one, corresponding to the Hicks compensated demand curve, holds electricity output (rather than consumer welfare) constant. Note that electricity output is perfectly observable, so there is no reason to disparage the Hicks flavor of compensated demand curve in factor markets. Nonetheless, our interest lies with the third variant: the Paasche-indexed output-compensated demand curve. The reason is that this is the only one of the three curves which is revenue neutral. That is, with Paasche indexing, the output compensation (in the form of a subsidy for electricity output) is exactly equal to the revenues collected from the carbon tax.

We are interested in revenue neutrality, not because it sounds like an appealing political slogan, but because it is an inherent property of emissions trading with a constraint on the intensive margin. Just as there is a mathematical duality between the two instruments that operate on the extensive margin (an emissions tax, and a cap-and-trade system), there is a similar mathematical duality between the two instruments that operate on the intensive margin. A tax on emissions which is exactly offset by a subsidy on output, is dually equivalent to a system of emissions trading with an intensity constraint. Both of these instruments will cause the price and quantity of carbon emissions to trace a revenue-neutral (i.e., Paasche indexed) output-compensated demand curve.

The key to demonstrating this duality is to decompose the shadow price of the intensity constraint by adding one extra variable, and one extra constraint, to the maximization problem. If the intensity constraint is expressed in tons of carbon per megawatt-hour, then the shadow price of that constraint should be expressed in dollar-megawatt-hours per ton – an unfathomable dimension. That vector can be resolved into two components, however: a shadow price, or shadow tax, on carbon that is expressed in $/ton, and a negative shadow price, or shadow subsidy, on output that is expressed in $/megawatt-hour. Replacing one shadow price with two, however, adds an extra degree of freedom to the maximization problem. We need to add one more equation to remove that degree of freedom: the revenue from the carbon tax must exactly offset the cost of the electricity subsidy. This is the revenue-neutrality constraint.

We know that an offset trading market is revenue neutral, because there is no mechanism for any central authority either to collect revenue or to make payments. Every trade in the market has a buyer and a seller. Such offset trading – trading under a constraint on the intensive margin – is naturally revenue neutral, which is what causes it to follow a Paasche-indexed output compensated demand curve, just like the output-compensated tax that Sweden used in the illustration above.

This framework of different demand curves may appear excessively abstract, but it has important real-world consequences. Only by understanding the relationship between embedded shadow taxes and shadow subsidies, their incidence, and their incentive effects, can we hope to design emissions control systems that are effective, fair, and efficient.