THE GEORGE WASHINGTON UNIVERSITY

WASHINGTON, DC

Comparing a Rebated Carbon Tax with a Compensated Carbon Tax, And Revisiting the Distinction Between Economic and Social Regulation

Brian F. Mannix[1]

Draft Working Paper[2] – Do Not Cite

November 19, 2019

The GW University Regulatory Studies Center

Abstract:

Part I of this paper compares and contrasts a simple carbon tax, a rebated carbon tax, and an output-compensated carbon tax, using the electric power industry to illustrate the differences. While these "price instruments" are serious policy proposals, the goal of this paper is to explore the underlying microeconomic theory and, in particular, to highlight some important properties of the compensated demand curves that describe their effects. Part II of the paper explains the mathematical duality between those three price instruments and their corresponding quantity instruments: auctioned cap-and-trade and allocated cap-and-trade, both of which impose a constraint on the quantity of emissions, and emissions-*rate*-cap-and-trade (also called offset trading), which imposes a constraint on the emission-intensiveness of industrial output. Keeping in mind the properties of compensated demand curves, Part III of the paper argues that the mathematical form of their regulatory constraints can help explain many of the typical differences between economic and social regulation – including the tendency of economic regulatory agencies to be multi-headed "independent" agencies, their inclination to use more adjudication and formal rulemaking than is typical of social regulators, and their troubling susceptibility to agency capture. A mathematical appendix explains the six flavors of compensated demand curve and the shadow prices that shape them.

---

[1] Brian Mannix is a Research Professor at the GW Regulatory Studies Center.

[2] This working paper 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.

Introduction

Economists have long advocated either Pigovian taxes or another form of "market-based emissions control" as superior to rigid "command and control" regulations, and there have been some notable successes in applying them. But the literature often obscures the policy choices that need to be made when introducing a market-based system, and the consequences that follow from those choices. Most environmental regulation does not impose a fixed cap on emissions, but rather imposes a ratio constraint, limiting the amount of emissions per unit of output. Emissions trading is compatible with either a fixed cap (i.e., a constraint on the *extensive* margin) or a ratio- or rate-cap (a constraint on the *intensive* margin). 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.[3] Other countries followed, and the United Nations Environment Program has estimated that the global benefit of removing lead from gasoline now amounts to $2.4 trillion per year.[4] A key factor leading to the success of this effort was the regulators' use of an intensity target (grams per gallon), rather than a fixed quantity target for limiting lead use.

This paper is intended to develop a structure for thinking about the policy options available when designing a system of market-based emissions control, and the advantages and disadvantages of each. It will begin by looking at options for taxing carbon emissions from electricity generating units. It will *not* address many other important questions related to carbon emissions, such as the state of climate science or methods of estimating the social cost of carbon. Regardless of how those other questions are resolved, any attempt to regulate carbon emissions will require that policy makers understand the dynamics of emissions taxation and/or trading, and the implications of working under a constraint on the intensive or extensive margin.

---

[3] 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.

[4] Peter L. Tsai and Thomas H. Hatfield, "Global Benefits From the Phaseout of Leaded Fuel," Journal of Environmental Health Vol. 74, No. 5 (December 2011), pp. 8-15. https://www.jstor.org/stable/26329321.

I. Price Instruments 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 three distinct forms of taxation:

1. A Carbon Tax
2. A Rebated Carbon Tax
3. An Output-Compensated Carbon Tax

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."[5]

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[6] of

---

[5] 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

[6] 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 and will ignore any upstream incidence. 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. This paper makes another simplifying assumption: that markets are effectively competitive. We will ignore complications related to local monopolies and associated price controls or other economic regulation.

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 government revenue. Such a tax would also render a jurisdiction 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 emissions "leakage" in addition to economic losses. 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

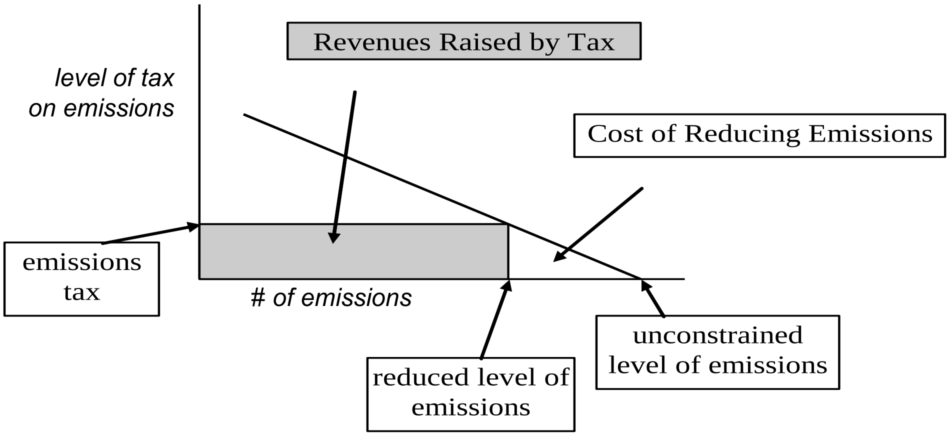

Before leaving the carbon tax, we need to make note of three distinct ways that it will cause emissions to be reduced below the level of unconstrained (untaxed) emissions. First, it causes a *technical substitution effect,* as electricity producers substitute lower carbon-intensity technologies for higher 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. 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

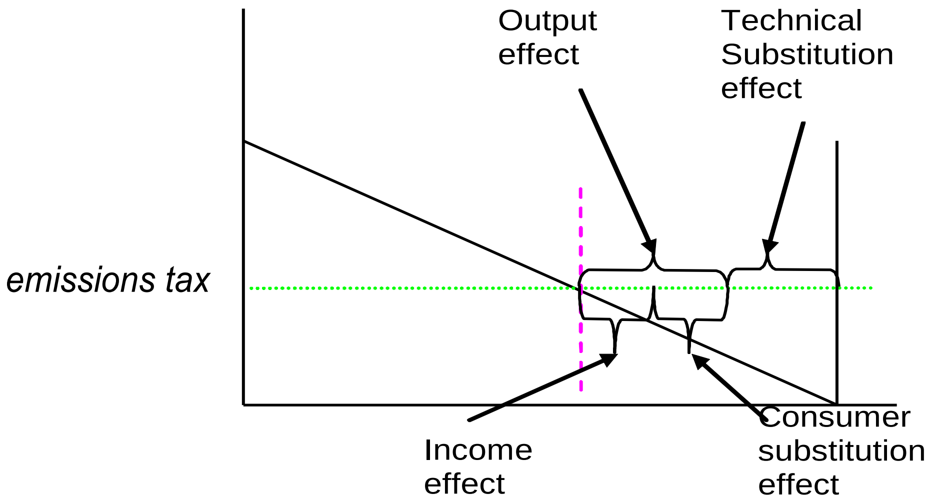

2: A Rebated Carbon Tax

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 a refundable tax credit, 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 shifts in tax policy can have uncertain outcomes, and will induce intense lobbying for favorable treatment – also known as rent-seeking.

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 3: Deriving an Income-Compensated Demand Curve

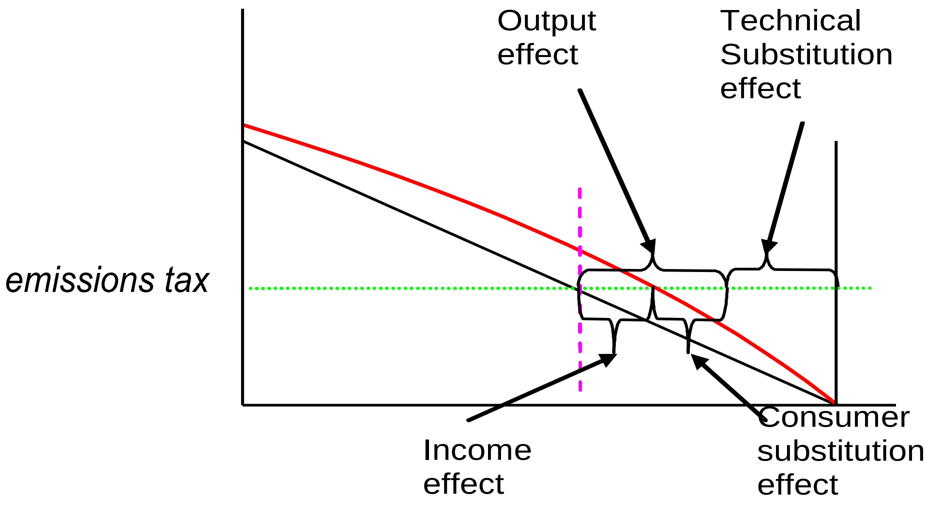

3: An Output-Compensated Carbon Tax

There is an alternative means of returning the revenues from a carbon tax to the consumers who pay 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.[7] In enacting its system of refunded emissions payments, Sweden was conscious of the fact that output compensation would

---

[7] Note that, in order to work as described, the tax must be on current emissions and the subsidy must be on current electricity production. Proposals to use “historical” emissions as the basis of the subsidy will not work the same way, since historical emissions are completely inelastic. The benefit of the subsidy would be captured by those sources of historical emissions, and (absent price controls) will not be passed through to consumers.

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.[8]

As seen in Figure 4, 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 below – an output-compensated demand curve. Note that, because it reflects only the technical substitution effect, an output compensated carbon tax will produce fewer emissions reductions than an uncompensated tax at the same price per ton of carbon. Still, from the electricity consumer's perspective, the cost of the output compensated tax will be dramatically lower.

Figure 4: Deriving an Output-Compensated Demand Curve

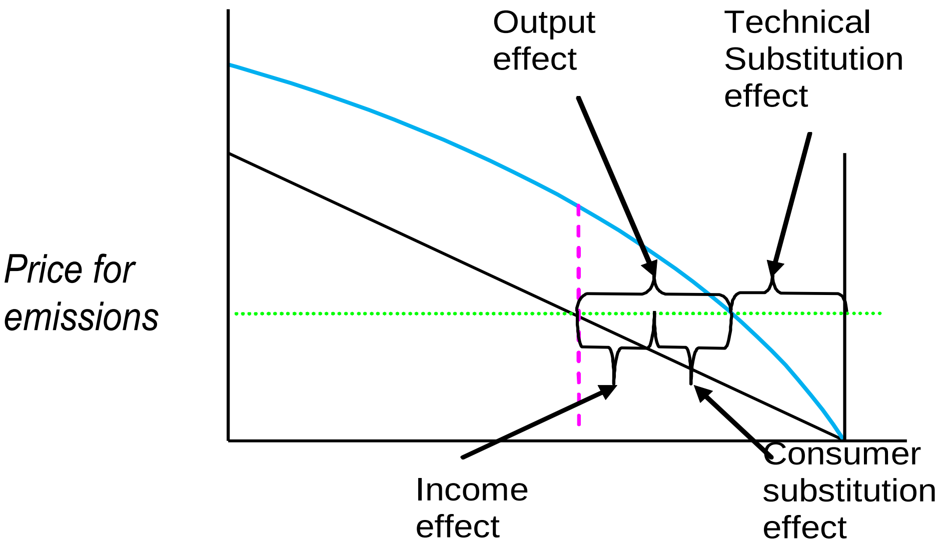

---

[8] 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

Note that, in contrast to the ordinary carbon tax and the rebated carbon tax, the compensated carbon tax does not create a need for "border adjustments" to avoid trade distortions with neighboring jurisdictions that are not taxing carbon. The rectangle of carbon-tax revenue in Figure 1 would cause a substantial increase in electricity prices, and rebating that revenue through other channels will not provide much relief. Under a compensated carbon tax, however, there is no net revenue. The price of electricity will reflect the burden of the carbon tax AND the offsetting benefit of the electricity subsidy. For this reason, a compensated carbon tax may be a much more practical option for a nation or state to pursue on its own, without the need for an interjurisdictional framework to deal with competitiveness concerns, emissions leakage, and border adjustment mechanisms.

II. Quantity Instruments for Reducing Emissions:

While emissions taxes have played an important role in the development of the theory of environmental economics, in practice they are quite rare. We have more experience with other forms of market-based emissions control. The three described in this section are the quantity instruments that correspond – in a mathematical "duality"[9] – to the three tax or price instruments described above.

4. A Cap-and-Trade System

A cap-and-trade system, with the allowances auctioned by the government, is the dual equivalent of a simple Pigovian emissions tax. As such, its effect on emissions can be described by the same factor-market demand curve we saw in Figure 1 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

---

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

up to the point where the entire scarcity rent has been exhausted. 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 it 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, it will substantially impair the competitiveness of any jurisdiction that adopts it, and it will encourage leakage of emissions to other jurisdictions.

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

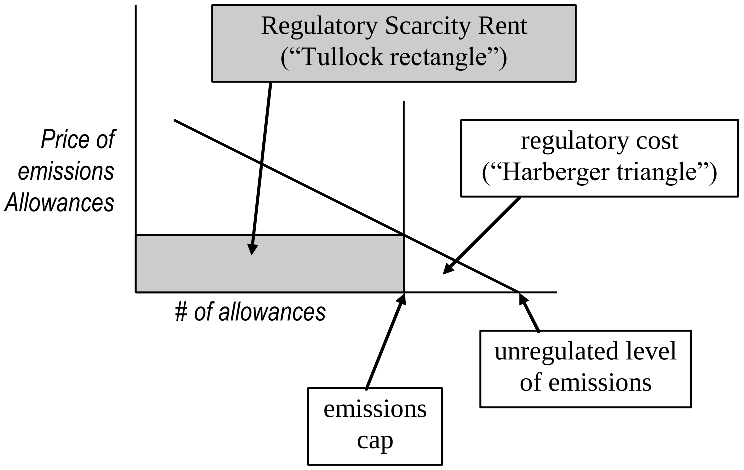

4: A Rebated Cap-and-Trade

We can imagine a rebated cap-and-trade system, as the dual equivalent of a rebated emissions tax. 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 (unrealistically) that rent-seeking could be avoided, a rebated cap-and-trade system would work very much like a rebated carbon tax, and would cause emissions to follow the same income-compensated demand curve illustrated earlier in Figure 3. And even a successful rebate would not do anything to mitigate the economic competitiveness and emissions leakage issues.

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 – often called "offset trading" in the literature. We can simply require that all covered sources comply with the same carbon intensity constraint, denominated in tons of carbon emissions per megawatt-hour of electricity output, and can allow trading of carbon allowances among them. 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 specified emissions-intensity goal, and trading will allow participants to 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 nor pay any subsidy; as a result this system is relatively 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. More than half of all refineries in the U.S. shut down in the next few years, because they were uneconomical and had been operating only to collect the federal subsidies: price controlled "oil entitlements" from DOE (ended in 1981) and extra lead allowances from EPA (ended in 1982). Within five years, with little resistance, EPA was able to phase out lead use almost entirely.

One disadvantage of emissions trading under an intensity constraint is that it will produce only a 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. In the case of lead in gasoline, technical substitution turned out to be sufficient to eliminate emissions entirely. But that will not always be the case. Applied to carbon emissions from electricity generation, such a system would not, for example, do much to encourage electricity conservation by consumers. But it would still be very effective in promoting technical change and it has other substantial advantages, in that it achieves emission reductions while simultaneously minimizing the increase in the price of electricity by avoiding any net tax burden. This avoids undesirable distributional impacts, and also minimizes the damage to economic competitiveness and the emissions leakage that arises from those competitiveness effects.

Finally, trading under a uniform intensity constraint has the advantage of providing a very small attack surface for rent-seekers. The market price of emissions allowances would follow an output-compensated demand curve, which does not have a Tullock rectangle. Just as an output-compensated carbon tax would have no net revenue, an output-compensated emissions trading system would have no fixed pool of allowances to be allocated. It would put a price on carbon emissions, but the price would be offset by a subsidy on electricity output,[10] so that the scarcity rents associated with carbon emissions are passed through to consumers and are difficult to divert to other potential recipients.

---

[10] A regulatory constraint on emissions intensity has two shadow prices: a "shadow tax" on the numerator, emissions, and a "shadow subsidy" on the denominator, output. See the Appendix for more detail.

No regulatory system is immune from rent-seekers; the ongoing contest between oil refiners and ethanol refiners (all at consumers' expense) under EPA's renewable fuel standard is instructive in that regard. Still, it is worth making the effort to minimize rent-seeking costs. International negotiations over climate change have repeatedly failed to make progress, in part because they have succumbed to rent-seeking at the national level. How much do we need to reduce carbon emissions, and which countries should accept what share of the burden?[11] The negotiations might have made more progress if instead they had focused on an output-compensated carbon tax, or its dual equivalent, an output-compensated emissions trading system – options which would not require any decisions about allocating the burden, and which would minimize concerns about international competitiveness and trade distortions.

III. Economic and Social Regulation Revisited

There are additional useful insights to be gained by looking at regulation through the lens of income-compensated and output-compensated demand curves. In 1974 economist Murray Weidenbaum, founder of the Center for the Study of American Business (now the Weidenbaum Center) at Washington University in St. Louis, divided the world of regulation into two broad categories. Economic regulation mostly meant the old regulated industries subject to rate-of-return regulation: railroads, trucks, buses, and pipelines by the Interstate Commerce Commission; airlines by the Civil Aeronautics Board; electromagnetic spectrum by the Federal Communications Commission, and so on. The second category was social regulation, including the health, safety, and environmental regulation of which EPA is the prime example. Just after he drew this distinction, the U.S. embarked on a decade of deregulation, as Presidents Ford, Carter, and Reagan begun to dismantle many of the old economic regulatory commissions. But that same decade was a period of rapid growth for EPA and the other social regulatory agencies. Clearly they are very different creatures.

---

[11] CEQ Chair Jim Connaughton, who led international negotiations for the U.S. 2007 – 2009, commented that: "We did spend an inordinate amount of time arguing about how to share the carbon space." Personal communication, January 2009.

I will argue that a fundamental difference between them is that social regulations are largely non-rivalrous, whereas economic regulatory agencies are engaged in – entangled in – an inherently rivalrous enterprise. What do I mean by that? Consider the nature of a license or permit issued by one of these agencies. Early in its history, the FCC awarded a certain slice of radio spectrum to Fetzer Communications. Shortly afterwards, a Mr. Ashbacker objected because, in granting the license to Mr. Fetzer, the Commission had effectively denied it to him. As the aggrieved party, Mr. Ashbacker was entitled to a hearing, and the Supreme Court agreed.[12] Ever since, the FCC has had to fashion its proceedings to protect the rights of the people to whom it is not giving a license. And the same is true of every other regulatory commission, whether it is handing out trucking routes or taxi medallions: on a day-to-day basis the main business of the commission is to settle rivalrous claims – granting a license to one is to deny to another. However much a commission might like to focus on the public interest, it inevitably spends more time and energy deciding among competing private interests.[13]

Contrast that with an environmental permit. There are some exceptions, but for the most part if Jane gets a permit to build something and I present a similar application, I'll get to build one too. In fact, the more I can show that my application is just like Jane's, the more likely to get a favorable treatment, rather than less. Environmental permits generally are not rivalrous.

The reason for this is that economic regulatory agencies generally place constraints on the extensive margin – the number of taxicabs, for example. Once they have done that, they are forced to tackle the much more contentious question of who gets to drive those taxicabs. Economic regulatory agencies work in the world of Marshallian demand curves or income-compensated demand curves, and this forces them to confront the task of divvying up the Tullock rectangle of rents that are created by the regulatory constraints they administer.

---

[12] *Ashbacker* v. *FCC*, 1944

[13] See Justice Frankfurter's dissent in *Ashbacker*, in which he complains of the FCC being distracted from its public mission by having to deal with private disputes.

In contrast, social regulatory agencies tend to place constraints on the intensive margin – grams per mile, miles per gallon, latrines per acre, ppm, and so on. Such compound constraints imply a compound shadow price with two components: a shadow tax on the targeted harm, and a shadow subsidy on the output with which it scales. If emissions trading is allowed, the market price will trace out an output compensated demand curve – one that lacks a Tullock rectangle. Regulatory constraints on the intensive margin simply do not create the same type of regulatory rents that are associated with exclusive permits. This does not make social regulatory agencies immune to rent-seeking; it just tends to take a different form. Theories of "agency capture," for example, seem to work well for economic regulatory agencies, but less well for social regulatory agencies.

This insight helps explain why economic regulation, with its "dividing up of spoils," so often uses distinct procedures, not only in the U.S. but around the world. It tends to be conducted at independent agencies, with some insulation from political interference. It tends to use more formal regulatory procedures and adjudication, in order to give due process to rivalrous contenders for its licenses. It is governed by case law and precedent to a far greater degree than social regulation, where policy analysis and benefit-cost analysis play a greater role.

We can also recognize that the differences between economic and social regulation has implications for what types of market failure each of them might be suitable for, as well as what types of regulatory failure each might be susceptible to. Economic regulation will be more appropriate when there is an underlying scarce resource that must be rationed, such as electromagnetic spectrum, or a depletable fishery. And it will work best when it can develop a regulatory system that employs, or at least mimics, property rights in that scarce resource. Social regulation be more appropriate when the goal is to encourage the use cleaner technologies, with fewer harmful side effects, to produce otherwise desirable outputs. In any case, regulators of all types must be alert to the various ways that rent-seeking can undermine their efforts.

Mathematical Appendix

For those who are interested, this Appendix gives a little more formal detail on the arguments outlined above, on the different flavors of compensated demand curves, and on the nature of the duality between a compensated emissions tax and emissions offset trading.[14]

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 ($\varepsilon_{\text{electricity}}$) into a consumer substitution effect and a consumer income effect.

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

$\varepsilon_{\text{income}}$ 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

---

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

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 all three types 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 – 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 ($\varepsilon_{\text{carbon}}$) into a technical substitution effect and an output effect.

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

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.