

Externalities and imperfect competition share the characteristic of diverting market outcomes from Pareto optimality. Though they differ in the source of market imperfection, they share some characteristics that make it useful for us to consider them together when thinking about potential government policy responses. We will consider externalities first.

We already know that a first-best response to externalities can in principle be achieved via Pigouvian taxes, which induce those causing the externality to internalize the external social costs (or benefits) of their actions. There are, of course, problems in implementing Pigouvian taxes; we may not know exactly what the marginal damage (or benefit) of an activity will be at a first-best equilibrium, since we are starting from a different equilibrium, and we don't have direct market valuations of the damage (or benefit). There may also be political reasons for tax-based solutions to be eschewed in favor of alternative approaches, if for example hidden taxes are more acceptable than explicit ones. There also may be administrative problems with taxing or otherwise controlling externalities directly, if it is hard to observe or measure their levels. For example, automobiles may cause negative externalities that are proportional to their tailpipe emissions of greenhouse gases, but the emissions from any particular automobile may be difficult for government to measure.


Alternatives to Corrective Taxes

Bovenberg and Goulder discuss alternatives to corrective taxation. The closest is probably *subsidies for pollution abatement*. If government pays polluters a subsidy for each unit that pollution falls below some benchmark level, then this is equivalent to providing them with a lump-sum transfer combined with a Pigouvian tax. That is, if the pollution benchmark is B units of emissions and the subsidy rate is s , the polluter receives a total subsidy $s(B - X)$ from the government, where X is the amount of pollution it emits. This is equivalent to getting a lump-sum transfer of sB and facing a tax at rate s per unit of pollution. In a world where the government can adjust lump-sum taxes and transfers, the subsidy-based policy and the tax-based policy are essentially identical. But, if the government must raise other, distortionary taxes to fund the lump-sum transfer implicit in the abatement subsidy, this policy may be less attractive, leaving aside the potentially adverse distributional consequences associated with transfers to polluters.

Quotas

Another common alternative to corrective taxation is *quotas* – using controls on quantities rather than prices to modify behavior. In the absence of uncertainty, quotas can be used to achieve the same outcome as corrective taxes. Consider the case where we know the marginal private costs of abatement and the marginal benefits of abatement. (Abatement can come from reducing output of a product that involves pollution, or from a switch in inputs – from coal to natural gas, for example; each method is costly to the producer, and we assume that producers efficiently choose among alternatives optimally, i.e., in a manner that minimizes their marginal abatement costs.) Representing costs, benefits and units of abatement in the following graph, we see that t^*

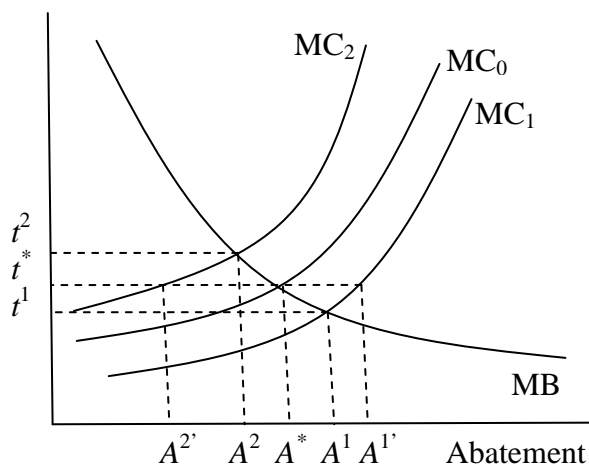
would be the optimal Pigouvian tax. If A^0 represents complete abatement, then producers would abate up to A^* and choose to pay the tax at that point on remaining emissions, as further abatement would cost them more than the tax.



How could we replicate this outcome using quotas? Suppose that polluters are required to obtain a permit for each unit of pollution, and that government auctions a total number of permits equal to $(A^0 - A^*)$. It follows that total pollution cannot exceed $(A^0 - A^*)$. Also, the market-clearing price of permits will be t^* , since at the equilibrium abatement level, A^* , the cost to producers of not buying a permit will be the marginal cost of further abatement. Hence, the level of abatement and payments by producers will be the same as under the tax regime.

Note that this equivalence requires that permits (1) are auctioned by the government; and (2) can be freely purchased by any agent in the market. A violation of (1), for example by giving away the permits, would leave marginal incentives unchanged but would reduce government revenue, just as in the above comparison of an abatement subsidy and a corrective tax. Even if the permits are given away, permit trading would still ensure that the same equilibrium prevailed. However, if condition (2) is violated, for example if permits are given away and must be used by the producers who receive them, then the equilibrium will be different, because the marginal abatement costs of each producer will not be equalized, and hence the overall costs of abatement will increase. Thus, *nontradeable* permits/quotas are inefficient and dominated by a policy that allows trading. It is worth noting that, in a setting in which trading is not allowed, the optimal level of abatement would be lower, since the marginal costs of abatement are higher.

Another difference between price and quantity approaches arises with uncertain abatement costs. Following Weitzman (*REStud* 1974), which approach is more effective depends on the relative slopes of the MB and MC curves in the above figure. Suppose the expected MC curve is MC_0 , but that costs might be higher or lower. If costs are lower, we would then wish to have a lower tax or fewer permits (more abatement); if costs are higher, we would want the opposite.



If we must fix either the abatement level A^* or the tax rate t^* , then either policy will induce deadweight loss. Under the tax policy, abatement will shift too much (to $A^{2'}$ or $A^{1'}$); under the abatement policy, it will not shift enough (or at all). Which distortion is greater depends on the relative slopes of the MB and MC curves. For example, if the MB curve is vertical, optimal abatement doesn't change and so the permit policy is best. If the MB curve is flat, then the optimal cost of abatement doesn't change, and so tax policy is best.

In reality, systems may resemble a mix of the two approaches; for example allowing more permits to be issued if the price of permits exceeds the expected price by more than a certain amount, or reducing the number of permits if the price is lower than expected, is a mixture of price and quantity schemes, effectively imposing a maximum and minimum permit price.

Performance Standards

A common approach to controlling externalities is a performance standard. An important example in the United States is the CAFE (Corporate Average Fuel Economy) standard requiring individual automobile producers to satisfy a certain miles-per-gallon rating for their annual sales. CAFE standards have some obvious flaws: (1) they cannot be traded among manufacturers; (2) they apply to average miles per gallon, whereas the logic of controlling gasoline usage would suggest using gallons per mile; (3) they are unrelated to the number of miles actually driven and the amount of gasoline actually used; and (4) they apply to some but not all categories of vehicles. While some or all of these important flaws need not apply to other performance standards, a fundamental problem with this class of regulations is that they implicitly combine a tax on externalities with a *subsidy* to production of the associated good.

Consider a firm that uses two inputs in producing output, labor (L) and energy (E), the first of which is “clean” (i.e., no externalities) and the second of which is “dirty” (negative externalities). We impose a performance standard that the ratio of energy to output cannot exceed a certain ratio, r . As a result, the firm’s optimization problem involves maximizing the Lagrangian, $pf(L, E) - wL - qE + \mu\left(r - \frac{E}{f(L, E)}\right)$, where output is $x=f(L, E)$, p is the price of output, w and q are input prices, and μ is the Lagrange multiplier of the regulation constraint. First-order conditions for L and E may be written:

$$pf_L - w + \mu \frac{E}{x^2} f_L = 0 \Rightarrow \left(p + \mu \frac{E}{x^2}\right) f_L = w; \text{ and}$$

$$pf_E - q + \mu \frac{E}{x^2} f_E - \mu \frac{1}{x} = 0 \Rightarrow \left(p + \mu \frac{E}{x^2}\right) f_E = q + \mu \frac{1}{x}$$

Now, suppose instead that we impose a subsidy at rate s on output and a tax at rate t on the use of energy, also requiring a balanced budget, e.g., $sx = tE$. The firm now seeks to maximize $(p + s)f(L, E) - wL - (q + t)E$, leading to first-order conditions:

$$(p + s)f_L = w \Rightarrow \left(p + \frac{tE}{x}\right) f_L = w$$

$$(p + s)f_E = q + t \Rightarrow \left(p + \frac{tE}{x}\right) f_E = q + t$$

Comparing the two sets of first-order conditions, we see that they are identical if we set $\mu = tx$. The intuition is that the firm relaxes the performance standard by producing more output, as doing so then allows the use of more energy. Indeed, given the conflicting incentives regarding usage of the dirty input (reducing its factor intensity but increasing its use through output expansion), it is possible that a performance standard may have the perverse effect of increasing overall use (see., e.g. Holland et al., *AEJ: Policy*, 2009). One could, of course, combine a performance standard with an output tax to mimic a corrective tax, as noted by Bovenberg and Goulder. In this simple example such a two-part policy would seem gratuitously complex, but in

realistic circumstances the two-part policy might be easier to implement, for example if input use or emissions were unobservable but output and the technology used could be verified.


Legal Remedies

One potential channel for forcing producers of negative externalities to internalize the damage they cause is the legal system. For example, if government established a legal liability rule, under which a producer that caused damages to consumers or to other producers could be sued for damages, then the producers of the externality would be induced to reduce emissions, at least to the extent that the marginal cost of abatement was lower than the marginal damage. Even without liability rules, one could imagine a broader use of the legal system to get to efficient outcomes. For example, victims of externalities could contract with polluters, agreeing to pay the polluters to abate. Assuming that the victims' marginal benefits of abatement exceed the marginal costs to the producers of abating, a welfare-increasing contract should be possible.

This reasoning dates from the classic paper by Coase (*J. Law & Econ.*, 1960), and is known as the Coase Theorem: well-defined property rights and a well-functioning legal system can correct externalities without additional government intervention. However, it is a realistic solution only under very restrictive conditions, breaking down where the individual causing the damage to a particular victim is hard to determine or where there is a large number of victims, in which case there would be a severe problem of coordination among those seeking a legal remedy.

Possible Complications due to Nonconvexities

As already discussed, one problem with correcting externalities is that we may not know what marginal damage is, precisely because the market imperfection inhibits the normal process by which valuations are revealed by transactions. But a deeper potential problem is that moving in a direction that we estimate is a *local* improvement in welfare may not move us toward a *global* optimum. By the nature of externalities, their costs depend not only on the level of activity, but also on the presence of victims. This interaction makes corner solutions more likely. For example, consider an economy with two goods, *X* and *Y*. Producing *X* generates no externalities, while *Y* can be produced using one of two technologies: technology *A* is more expensive but limits emissions, while technology *B* is less expensive but produces emissions with such intensity that any production of *Y* precludes production of *X*. Society's production possibilities frontier will then resemble the following graph, where the convex portion corresponds to technology *A* and the point on the horizontal axis to technology *B*:



The optimal point can be either on the convex portion of the PPF or at point B, depending on consumer preferences. But if we start at some initial point, say C, where production of *X* and *Y* both occur, we will observe the presence of an unpriced externality and be inclined to impose a tax on production of *Y*, which will shift the mix of production in the direction shown, increasing production of *X* and reducing production of *Y*. This will move toward a local welfare maximum, but away from the global maximum if that is at point B.