

THE USE OF STANDARDS AND PRICES FOR PROTECTION OF THE ENVIRONMENT

*William J. Baumol and Wallace E. Oates**

Princeton University, Princeton, N.J., USA

Summary

In the Pigouvian tradition, economists have frequently proposed the adoption of a system of unit taxes (or subsidies) to control externalities, where the tax on a particular activity is equal to the marginal social damage it generates. In practice, however, such an approach has rarely proved feasible because of our inability to measure marginal social damage.

This paper proposes that we establish a set of admittedly somewhat arbitrary standards of environmental quality (e.g., the dissolved oxygen content of a waterway will be above x per cent at least 99 per cent of the time) and then impose a set of charges on waste emissions sufficient to attain these standards. While such *resource-use prices* clearly will not in general produce a Pareto-efficient allocation of resources, it is shown that they nevertheless do possess some important optimality properties and other practical advantages. In particular, it is proved that, for any given vector of final outputs such prices can achieve a specified reduction in pollution levels at minimum cost to the economy, even in the presence of firms with objectives other than that of simple profit maximization.

In the technicalities of the theoretical discussion of the tax-subsidy approach to the regulation of externalities, one of the issues most critical for its application tends to get the short end of the discussion. Virtually every author points out that we do not know how to calculate the ideal Pigouvian tax or subsidy levels in practice, but because the point is rather obvious rarely is much made of it.

This paper reviews the nature of the difficulties and then proposes a substitute approach to the externalities problem. This alternative, which we shall call the environmental pricing and standards procedure, represents what we consider to be as close an approximation as one can generally achieve in practice to the spirit of the Pigouvian tradition. Moreover, while this method does not aspire to anything like an optimal allocation of resources, it will be shown to possess some important optimality properties.

* The authors are members of the faculty at Princeton University. They are grateful to the Ford Foundation whose support greatly facilitated the completion of this paper.

1. Difficulties in Determining the Optimal Structure of Taxes and Subsidies

The proper level of the Pigouvian tax (subsidy) upon the activities of the generator of an externality is equal to the marginal net damage (benefit) produced by that activity.¹ The difficulty is that it is usually not easy to obtain a reasonable estimate of the money value of this marginal damage. Kneese & Bower report some extremely promising work constituting a first step toward the estimation of the damage caused by pollution of waterways including even some quantitative evaluation of the loss in recreational benefits. However, it is hard to be sanguine about the availability in the foreseeable future of a comprehensive body of statistics reporting the marginal net damage of the various externality-generating activities in the economy. The number of activities involved and the number of persons affected by them are so great that on this score alone the task assumes Herculean proportions. Add to this the intangible nature of many of the most important consequences—the damage to health, the aesthetic costs—and the difficulty of determining a money equivalent for marginal net damage becomes even more apparent.

This, however, is not the end of the story. The optimal tax level on an externality generating activity is not equal to the marginal net damage it generates *initially*, but rather to the damage it would cause if the level of the activity had been adjusted to its *optimal* level. To make the point more specifically, suppose that each additional unit of output of a factory now causes 50 cents worth of damage, but that after the installation of the appropriate smoke-control devices and other optimal adjustments, the marginal social damage would be reduced to 20 cents. Then a little thought will confirm what the appropriate mathematics show: the correct value of the Pigouvian tax is 20 cents per unit of output, that is, the marginal cost of the smoke damage *corresponding to an optimal situation*. A tax of 50 cents per unit of output corresponding to the current smoke damage cost would lead to an excessive reduction in the smoke-producing activity, a reduction beyond the range over which the marginal benefit of decreasing smoke emission exceeds its marginal cost.

The relevance of this point for our present discussion is that it compounds enormously the difficulty of determining the optimal tax and benefit levels. If there is little hope of estimating the damage that is currently generated, how much less likely it is that we can evaluate the damage that would occur in an optimal world which we have never experienced or even described in quantitative terms.

There is an alternative possibility. Instead of trying to go directly to the optimal tax policy, one could instead, as a first approximation, base a set of

¹ We will use the term marginal *net* damage to mean the difference between marginal social and private damage (or cost).

taxes and subsidies on the current net damage (benefit) levels. Then as outputs and damage levels were modified in response to the present level of taxes, the taxes themselves would in turn be readjusted to correspond to the new damage levels. It can be hoped that this will constitute a convergent, iterative process with tax levels affecting outputs and damages, these in turn leading to modifications in taxes, and so on. It is not clear, however, even in theory, whether this sequence will in fact converge toward the optimal taxes and resource allocation patterns. An extension of the argument underlying some of Coase's illustrations can be used to show that convergence cannot always be expected. But even if the iterative process were stable and were in principle capable of yielding an optimal result, its practicality is clearly limited. The notion that tax and subsidy rates can be readjusted quickly and easily on the basis of a fairly esoteric marginal net damage calculation does not seem very plausible. The difficulty of these calculations has already been suggested, and it is not easy to look forward with equanimity to their periodic revision, as an iterative process would require.

In sum, the basic trouble with the Pigouvian cure for the externalities problem does not lie primarily in the technicalities that have been raised against it in the theoretical literature but in the fact that we do not know how to determine the dosages that it calls for. Though there may be some special cases in which one will be able to form reasonable estimates of the social damages, in general we simply do not know how to set the required levels of taxes and subsidies.

2. The Environmental Pricing and Standards Approach

The economist's predilection for the use of the price mechanism makes him reluctant to give up the Pigouvian solution without a struggle. The inefficiencies of a system of direct controls, including the high real enforcement costs that generally accompany it, have been discussed often enough; they require no repetition here.

There is a fairly obvious way, however, in which one can avoid recourse to direct controls and retain the use of the price system as a means to control externalities. Simply speaking, it involves the selection of a set of somewhat arbitrary standards for an acceptable environment. On the basis of evidence concerning the effects of unclean air on health or of polluted water on fish life, one may, for example, decide that the sulfur-dioxide content of the atmosphere in the city should not exceed x percent, or that the oxygen demand of the foreign matter contained in a waterway should not exceed level y , or that the decibel (noise) level in residential neighborhoods should not exceed z at least 99% of the time. These acceptability standards, x , y and z , then amount to a set of constraints that society places on its activities. They represent the decision-maker's subjective evaluation of the minimum

Swed. J. of Economics 1971

standards that must be met in order to achieve what may be described in persuasive terms as “a reasonable quality of life”. The defects of the concept will immediately be clear to the reader, and, since we do not want to minimize them, we shall examine this problem explicitly in a later section of the paper.

For the moment, however, we want to emphasize the role of the price system in the implementation of these standards. The point here is simply that the public authority can levy a uniform set of taxes which would in effect constitute a set of prices for the private use of social resources such as air and water. The taxes (or prices) would be selected so as to achieve specific acceptability standards rather than attempting to base them on the unknown value of marginal net damages. Thus, one might tax all installations emitting wastes into a river at a rate of $t(b)$ cents per gallon, where the tax rate, t , paid by a particular polluter, would, for example, depend on b , the BOD value of the effluent, according to some fixed schedule.¹ Each polluter would then be given a financial incentive to reduce the amount of effluent he discharges and to improve the quality of the discharge (i.e., reduce its BOD value). By setting the tax rates sufficiently high, the community would presumably be able to achieve whatever level of purification of the river it desired. It might even be able to eliminate at least some types of industrial pollution altogether.²

Here, if necessary, the information needed for iterative adjustments in tax rates would be easy to obtain: if the initial taxes did not reduce the pollution of the river sufficiently to satisfy the preset acceptability standards, one would simply raise the tax rates. Experience would soon permit the authorities to estimate the tax levels appropriate for the achievement of a target reduction in pollution.

One might even be able to extend such adjustments beyond the setting of the tax rates to the determination of the acceptability standards themselves. If, for example, attainment of the initial targets were to prove unexpectedly inexpensive, the community might well wish to consider making the standards stricter.³ Of course, such an iterative process is not costless. It means that at least some of the polluting firms and municipalities will have to adapt their

¹ BOD, biochemical oxygen demand, is a measure of the organic waste load of an emission. It measures the amount of oxygen used during decomposition of the waste materials. BOD is used widely as an index of the quality of effluents. However, it is only an approximation at best. Discharges whose BOD value is low may nevertheless be considered serious pollutants because they contain inorganic chemical poisons whose oxygen requirement is nil because the poisons do not decompose. See Kneese and Bower on this matter.

² Here it is appropriate to recall the words of Chief Justice Marshall, when he wrote that “The power to tax involves the power to destroy” (*McCulloch vs. Maryland*, 1819). In terms of reversing the process of environmental decay, we can see, however, that the power to tax can also be the power to restore.

³ In this way the pricing and standards approach might be adapted to approximate the Pigouvian ideal. If the standards were revised upward whenever there was reason to believe that the marginal benefits exceeded the marginal costs, and if these judgments were reasonably accurate, the two would arrive at the same end product, at least if the optimal solution were unique.

operations as tax rates are readjusted. At the very least they should be warned in advance of the likelihood of such changes so that they can build flexibility into their plant design, something which is not costless (See Hart). But, at any rate, it is clear that, through the adjustment of tax rates, the public authority can realize whatever standards of environmental quality it has selected.

3. Optimality Properties of the Pricing and Standards Technique

While the pricing and standards procedure will not, in general, lead to Pareto-efficient levels of the relevant activities, it is nevertheless true that the use of unit taxes (or subsidies) to achieve the specified quality standards does possess one important optimality property: it is the least-cost method to realize these targets.¹ A simple example may serve to clarify this point. Suppose that it is decided in some metropolitan area that the sulfur-dioxide content of the atmosphere should be reduced by 50%. An obvious approach to this matter, and the one that often recommends itself to the regulator, is to require each smoke-producer in the area to reduce his emissions of sulfur dioxide by the same 50%. However, a moment's thought suggests that this may constitute a very expensive way to achieve the desired result. If, at existing levels of output, the marginal cost of reducing sulfur-dioxide emissions for Factory *A* is only one-tenth of the marginal cost for Factory *B*, we would expect that it would be much cheaper for the economy as a whole to assign *A* a much greater decrease in smoke emissions than *B*. Just how the least-cost set of relative quotas could be arrived at in practice by the regulator is not clear, since this obviously would require calculations involving simultaneous relationships and extensive information on each polluter's marginal-cost function.

It is easy to see, however, that the unit-tax approach can *automatically* produce the least-cost assignment of smoke-reduction quotas without the need for any complicated calculations by the enforcement authority. In terms of our preceding example, suppose that the public authority placed a unit tax on smoke emissions and raised the level of the tax until sulfur-dioxide emissions were in fact reduced by 50%. In response to a tax on its smoke emissions, a cost-minimizing firm will cut back on such emissions until the marginal cost of further reductions in smoke output is equal to the tax. But, since all economic units in the area are subject to the same tax, it follows that the marginal cost of reducing smoke output will be equalized across all activities. This implies that it is impossible to reduce the aggregate cost of the specified decrease in smoke emissions by re-arranging smoke-reduction quotas: any alteration in this pattern of smoke emissions would involve an increase in

¹ This proposition is not new. While we have been unable to find an explicit statement of this result anywhere in the literature, it or a very similar proposition has been suggested in a number of places. See, for example, Kneese & Bower, Chapter 6, and Ruff, p. 79.

smoke output by one firm the value of which to the firm would be less than the cost of the corresponding reduction in smoke emissions by some other firm. For the interested reader, a formal proof of this least-cost property of unit taxes for the realization of a specified target level of environmental quality is provided in an appendix to this paper. We might point out that the validity of this least-cost theorem does not require the assumption that firms are profit-maximizers. All that is necessary is that they minimize costs for whatever output levels they should select, as would be done, for example, by a firm that seeks to maximize its growth or its sales.

The cost saving that can be achieved through the use of taxes and subsidies in the attainment of acceptability standards may by no means be negligible. In one case for which comparable cost figures have been calculated, Kneese & Bower (p. 162) report that, with a system of uniform unit taxes, the cost of achieving a specified level of water quality would have been only about half as high as that resulting from a system of direct controls. If these figures are at all representative, then the potential waste of resources in the choice between tax measures and direct controls may obviously be of a large order. Unit taxes thus appear to represent a very attractive method for the realization of specified standards of environmental quality. Not only do they require relatively little in the way of detailed information on the cost structures of different industries, but they lead automatically to the least-cost pattern of modification of externality-generating activities.

4. Where the Pricing and Standards Approach is Appropriate

As we have emphasized, the most disturbing aspect of the pricing and standards procedure is the somewhat arbitrary character of the criteria selected. There does presumably exist some optimal level of pollution (i.e., quality of the air or a waterway), but in the absence of a pricing mechanism to indicate the value of the damages generated by polluting activities, one knows no way to determine accurately the set of taxes necessary to induce the optimal activity levels.

While this difficulty certainly should not be minimized, it is important at the outset to recognize that the problem is by no means unique to the selection of acceptability standards. In fact, as is well known, it is a difficulty common to the provision of nearly all public goods. In general, the market will not generate appropriate levels of outputs where market prices fail to reflect the social damages (or benefits) associated with particular activities. As a result, in the absence of the proper set of signals from the market, it is typically necessary to utilize a political process (i.e., a method of collective choice) to determine the level of the activity.¹ From this perspective, the selec-

¹ As Coase and others have argued, voluntary bargains struck among the interested parties may in some instances yield an efficient set of activity levels in the presence of externalities. However, such coordinated, voluntary action is typically possible only in small groups. One can hardly imagine, for example, a voluntary bargaining process involving all the persons in a metropolitan area and resulting in a set of payments that would generate efficient levels of activities affecting the smog content of the atmosphere.

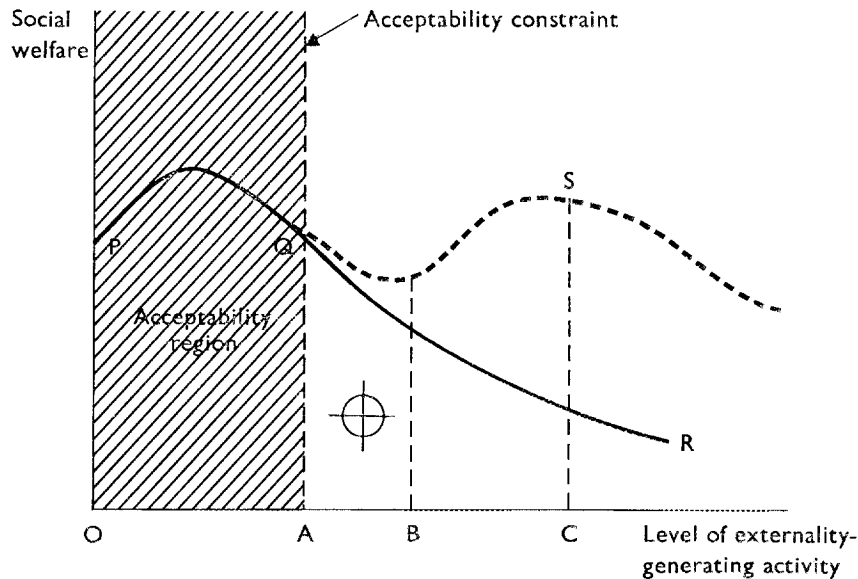


Fig. 1

tion of environmental standards can be viewed as a particular device utilized in a process of collective decision-making to determine the appropriate level of an activity involving external effects.

Since methods of collective choice, such as simple-majority rule or decisions by an elected representative, can at best be expected to provide only very rough approximations to optimal results, the general problem becomes one of deciding whether or not the malfunction of the market in a certain case is sufficiently serious to warrant public intervention. In particular, it would seem to us that such a blunt instrument as acceptability standards should be used only sparingly, because the very ignorance that serves as the rationale for the adoption of such standards implies that we can hardly be sure of their consequences.

In general, it would seem that intervention in the form of acceptability standards can be utilized with any degree of confidence only where there is clear reason to believe that the existing situation imposes a high level of social costs *and* that these costs can be significantly reduced by feasible decreases in the levels of certain externality-generating activities. If, for example, we were to examine the functional relationship between the level of social welfare and the levels of particular activities which impose marginal net damages, the argument would be that the use of acceptability standards is justified only in those cases where the curve, over the bulk of the relevant range, is both decreasing and steep. Such a case is illustrated in Fig. 1 by the curve *PQR*. In a case of this kind, although we obviously will not have an accurate knowledge

Swed. J. of Economics 1971

of the relevant position of the curve, we can at least have some assurance that the selection of an acceptability standard and the imposition of a unit tax sufficient to realize that standard will lead to an increase in social welfare. For example, in terms of the curve PQR in Fig. 1, the levying of a tax sufficient to reduce smoke outputs from level OC to OA to ensure that the quality of the air meets the specified environmental standards would obviously increase social welfare.¹

On the other hand, if the relationship between social welfare and the level of the externality-generating activity is not monotonically decreasing, the changes resulting from the imposition of an acceptability standard (e.g., a move from S to Q in Fig. 1) clearly may lead to a reduction in welfare. Moreover, even if the function were monotonic but fairly flat, the benefits achieved might not be worth the cost of additional intervention machinery that new legislation requires, and it would almost certainly not be worth the risk of acting with highly imperfect, inconclusive information.

In some cases, notably in the field of public utility regulation, some economists have criticized the employment of acceptability standards on both these grounds; they have asserted that the social costs of monopolistic misallocation of resources are probably not very high (i.e., the relevant portion of the social-welfare curve in Fig. 1 is not steep) and that the regulation can itself introduce inefficiencies in the operations of the regulated industries.

Advocacy of environmental pricing and standards procedures for the control of externalities must therefore rest on the belief that in this area we do have a clear notion of the general shape of the social welfare curve. This will presumably hold true where the evidence indicates, first that a particular externality really does have a substantial and unambiguous effect on the quality of life, if, for example, it makes existence very unpleasant for everyone or constitutes a serious hazard to health; and second that reductions in the levels of these activities do not themselves entail huge resource costs. On the first point, there

¹ The relationship depicted in Fig. 1 is to be regarded as an intuitive device employed for pedagogical purposes, not in any sense as a rigorous analysis. However, some further explanation may be helpful. The curve itself is not a social-welfare function in the usual sense; rather it measures in terms of a numeraire (kronor or dollars) the value, summed over all individuals, of the benefits from the output of the activity minus the private and net social costs. Thus, for each level of the activity, the height of the curve indicates the net benefits (possibly negative) that the activity confers on society. The acceptability constraint indicates that level of the activity which is consistent with the specified minimum standard of environmental quality (e.g., that level of smoke emissions from factories which is sufficiently low to maintain the quality of the air in a particular metropolitan area). There is an ambiguity here in that the levels of several different activities may jointly determine a particular dimension of environmental quality, e.g., the smoke emissions of a number of different industries will determine the quality of the air. In this case, the acceptable level of pollutive emissions for the firm or industry will clearly depend on the levels of emissions of others. If, as we discussed earlier, unit taxes are used to realize the acceptability standards, there will result a least-cost pattern of levels of the relevant externality-generating activities. If we understand the constraint in Fig. 1 to refer to the activity level indicated by this particular solution, then this ambiguity disappears.

is growing evidence that various types of pollutants do in fact have such unfortunate consequences, particularly in areas where they are highly concentrated. [On this see, for instance, Lave & Seskin]. Second, what experience we have had with, for example, the reduction of waste discharges into waterways suggests that processes involving the recycling and reuse of waste materials can frequently be achieved at surprisingly modest cost.¹ In such cases the rationale for the imposition of environmental standards is clear, and it seems to us that the rejection of such crude measures on the grounds that they will probably violate the requirements of optimality may well be considered a kind of perverse perfectionism.

It is interesting in this connection that the pricing and standards approach is not too different in spirit from a number of economic policy measures that are already in operation in other areas. This is significant for our discussion, because it suggests that regulators know how to work with this sort of approach and have managed to live with it elsewhere. Probably the most noteworthy example is the use of fiscal and monetary policy for the realization of macroeconomic objectives. Here, the regulation of the stock of money and the availability of credit along with adjustments in public expenditures and tax rates are often aimed at the achievement of a selected target level of employment or rate of inflation. Wherever prices rise too rapidly or unemployment exceeds an "acceptable" level, monetary and fiscal variables are readjusted in an attempt to "correct" the difficulty. It is noteworthy that this procedure is also similar to the pricing and standards approach in its avoidance of direct controls.

Other examples of this general approach to policy are not hard to find. Policies for the regulation of public-utilities, for instance, typically utilize a variety of standards such as profit-rate ceilings (i.e., "fair rates of return") to judge the acceptability of the behavior of the regulated firm. In the area of public education, one frequently encounters state-imposed standards (e.g., subjects to be taught) for local school districts which are often accompanied by grants of funds to the localities to help insure that public-school programs meet the designated standards. What this suggests is that public administrators are familiar with this general approach to policy and that the implementation of the pricing and standards technique should not involve insurmountable administrative difficulties. For these reasons, the achievement of specified environmental standards through the use of unit taxes (or subsidies) seems to us to possess great promise as a workable method for the control of the quality of the environment.

¹ Some interesting discussions of the feasibility of the control of waste emissions into waterways often at low cost are contained in Kneese & Bower. In particular, see their description of the control of water quality in the Ruhr River in Germany.

5. Concluding Remarks

It may be useful in concluding our discussion simply to review the ways in which the pricing and standards approach differs from the standard Pigouvian-prescription for the control of externalities.

(1) Under the Pigouvian technique, unit taxes (or subsidies) are placed on externality-generating activities, with the level of the tax on a particular activity being set equal to the marginal net damage it generates. Such taxes (if they could be determined) would, it is presumed, lead to Pareto-efficient levels of the activities.

(2) In contrast, the pricing and standards approach begins with a predetermined set of standards for environmental quality and then imposes unit taxes (or subsidies) sufficient to achieve these standards. This will not, in general, result in an optimal allocation of resources, but (as is proved formally in the appendix) the procedure does at least represent the least-cost method of realizing the specified standards.

(3) The basic appeal of the pricing and standards approach relative to the Pigouvian prescription lies in its workability. We simply do not, in general, have the information needed to determine the appropriate set of Pigouvian taxes and subsidies. Such information is not, however, necessary for our suggested procedure.

(4) While it makes no pretense of promising anything like an optimal allocation of resources, the pricing and standards technique can, in cases where external effects impose high costs (or benefits), at least offer some assurance of reducing the level of these damages. Moreover, the administrative procedures—the selection of standards and the use of fiscal incentives to realize these standards—implied by this approach are in many ways quite similar to those used in a number of current public programs. This, we think, offers some grounds for optimism as to the practicality of the pricing and standards technique for the control of the quality of the environment.

References

1. Bohm, P.: Pollution, Purification, and the Theory of External Effects. *Swedish Journal of Economics* 72, no. 2, 153–66, 1970.
2. Coase, R.: The Problem of Social Cost. *Journal of Law and Economics* 3, 1–44, 1960.
3. Hart, A.: Anticipations, Business Planning, and the Cycle. *Quarterly Journal of Economics* 51, 273–97, Feb. 1937.
4. Kneese, A. & Bower, B.: *Managing Water Quality: Economics, Technology, Institutions*. Baltimore, 1968.
5. Lave, L. & Seskin, E.: Air Pollution and Human Health. *Science* 21, 723–33 Aug. 1970.
6. Portes, R.: The Search for Efficiency in the Presence of Externalities. *Unfashionable Economics: Essays in Honor of Lord Balogh* (ed. P. Streeten), pp. 348–61. London, 1970.
7. Ruff, L.: The Economic Common Sense of Pollution. *The Public Interest*, Spring 1970, 69–85.

Swed. J. of Economics 1971

APPENDIX

In the text, we argued on a somewhat intuitive level that the appropriate use of unit taxes and subsidies represents the least-cost method of achieving a set of specified standards for environmental quality. In the case of smoke-abatement, for instance, the tax-subsidy approach will automatically generate the cost-minimizing assignment of "reduction quotas" without recourse to involved calculations or enforcement.

The purpose of this appendix is to provide a formal proof of this proposition. More precisely, we will show that, to achieve *any* given vector of final outputs along with the attainment of the specified quality of the environment, the use of unit taxes (or, where appropriate, subsidies) to induce the necessary modification in the market-determined pattern of output will permit the realization of the specified output vector at minimum cost to society.

While this theorem may seem rather obvious (as the intuitive discussion in the text suggests), its proof does point up several interesting properties which are noteworthy. In particular, unlike many of the propositions about prices in welfare analysis, the theorem does not require a world of perfect competition. It applies to pure competitors, monopolists, or oligopolists alike so long as each of the firms involved seeks to minimize the private cost of producing whatever vector of outputs it selects and has no monopsony power (i.e., no influence on the prices of inputs). The firms need not be simple profit-maximizers; they may choose to maximize growth, sales (total revenues), their share of the market, or any combination of these goals (or a variety of other objectives). Since the effective pursuit of these goals typically entails minimizing the cost of whatever outputs are produced, the theorem is still applicable. Finally, we want simply to emphasize that the theorem applies to whatever set of final outputs society should select (either by direction or through the operation of the market). It does not judge the desirability of that particular vector of outputs; it only tells us how to make the necessary adjustments at minimum cost.

We shall proceed initially to derive the first-order conditions for the minimization of the cost of a specified overall reduction in the emission of wastes. We will then show that the independent decisions of cost-minimizing firms subject to the appropriate unit tax on waste emissions will, in fact, satisfy the first-order conditions for overall cost minimization.

Let

x_{iv} represent the quantity of input i used by plant v ($i = 1, \dots, n$), ($v = 1, \dots, m$),

z_v be the quantities of waste it discharges,

y_v be its output level,

$f_v(x_{1v}, \dots, x_{nv}, z_v, y_v) = 0$ be its production function,

p_i be the price of input i , and

k the desired level of $\sum_v z_v$, the maximum permitted daily discharge of waste.

Swed. J. of Economics 1971

In this formulation, the value of k is determined by the administrative authority in a manner designed to hold waste emissions in the aggregate to a level consistent with the specified environmental standard (e.g., the sulphuric content of the atmosphere). Note that the level of the firm's waste emissions is treated here as an argument in its production function; to reduce waste discharges while maintaining its level of output, the firm will presumably require the use of additional units of some other inputs (e.g., more labor or capital to recycle the wastes or to dispose of them in an alternative manner).

The problem now becomes that of determining the value of the x 's and z 's that minimize input cost

$$c = \sum_i \sum_v p_i x_{iv}$$

subject to the output constraints

$$y_v = y_v^* = \text{constant} \quad (v = 1, \dots, m)$$

and the constraint on the total output of pollutants

$$\sum_v z_v = k.$$

It may appear odd to include as a constraint a vector of given outputs of the firms, since the firms will presumably adjust output levels as well as the pattern of inputs in response to taxes or other restrictions on waste discharges. This vector, however, can be *any* vector of outputs (including that which emerges as a result of independent decisions by the firms). What we determine are first-order conditions for cost-minimization which apply to *any* given vector of outputs no matter how they are arrived at. Using $\lambda_v (v=1, \dots, m)$ and λ as our $(m+1)$ Lagrange multipliers, we obtain the first-order conditions:

$$\left. \begin{aligned} \lambda_v f_{vz} + \lambda &= 0 & (v = 1, \dots, m) \\ p_i + \lambda_v f_{vi} &= 0 & (v = 1, \dots, m) (i = 1, \dots, n) \\ y_v &= y_v^* & (v = 1, \dots, m) \end{aligned} \right\} \quad (1)$$

where we use the notation $f_{vz} = \partial f_v / \partial z_v$, $f_{vi} = \partial f_v / \partial x_{iv}$.

Now let us see what will happen if the m plants are run by independent managements whose objective is to minimize the cost of whatever outputs their firm produces, and if, instead of the imposition of a fixed ceiling on the emission of pollutants, this emission is taxed at a fixed rate per unit, t . So long as its input prices are fixed, firm v will wish to minimize

$$c = tz_v + \sum_i p_i x_{iv}$$

subject to

$$y_v = y_v^*.$$

Direct differentiation of the m Lagrangian functions for our m firms immediately yields the first-order conditions (1)—the same conditions as before, provided t is set equal to $\hat{\lambda}$. Thus, if we impose a tax rate that achieves the desired reduction in the total emission of pollutants, we have proved that this reduction will satisfy the necessary conditions for the minimization of the program's cost to society.¹

¹ In this case, λ (and hence t) is the shadow price of the pollution constraint. In addition to satisfying these necessary first-order conditions, cost-minimization requires that the production functions possess the usual second-order properties. An interesting treatment of this issue is available in Portes. We should also point out that our proof assumes that the firm takes t as given and beyond its control. Bohm discusses some of the problems that can arise where the firm takes into account the effects of its behavior on the value of t .