loss to the farmer is measured unambiguously by the loss of output or the cost of obtaining clean water, whichever is less. But when the damaged party is a consumer (and this, we have argued, is the more typical case), willingness to pay may differ from required compensation because the former is constrained by the consumer's income. The result is that the assignment of property rights will affect resource use.[18]

In summary, it appears that the Coase theorem fails as a challenge to pollution-control policy involving some form of public intervention. It does offer an insight into the virtues of the market in dealing with certain kinds of externalities, but generally not those associated with pollution or other environmental disruption.

The cost-effectiveness of a tax

Another kind of challenge to a pollution tax comes not from a school of academic economists, as in the case of the Coase theorem, but from noneconomists. The contention is that the information required to implement a tax (the marginal damage, at the optimal point, to all pollution receivers) is not available and is not likely to be. One implication is that neither a tax nor the economic theory on which it rests is particularly relevant to practical attempts to deal with pollution. Many economists accept, at least provisionally, the claim that we do not know enough about damage functions to design a tax to achieve full Pareto optimality.[19] But these same economists have shown how a tax can be used to achieve the more modest, but still important, objective of cost-effective control.[20]

[18] Income effects were analyzed by Dolbear (1967) and Mishan (1967). For an amusing critique of the Coase theorem and extensions as applied to pollution, see Mishan's "Pangloss on Pollution" (1971*b*).

[19] Methods of estimating damages will be discussed in detail in the next section.

[20] Versions of this result have been obtained and discussed by many people. See, for example, the work of Kneese (1964), Ruff (1970),

Suppose we view the problem as one of choosing, through the political process, a standard of environmental quality (much as we choose amounts of other public goods, such as national defense) and then seeking a method to achieve it at least cost. In what follows we shall show that a tax will do this and that direct controls on emissions, a method favored by many noneconomists, probably will not. However, there are circumstances in which controls may be superior to a tax or can usefully supplement a tax. Our approach in proving the cost-minimization theorem is similar to the one adopted in the preceding section. We shall first derive necessary conditions for achieving a preselected level of pollution at minimum cost and then show that the same conditions are satisfied by the decentralized decisions of polluting firms subject to an appropriate tax.

Formally, the planner's problem is to minimize the sum of expenditures on two kinds of inputs (those used to produce conventional goods and services and those used to control pollution), subject to restrictions on production, on pollution, and on the relationship between production and pollution. Previously, we considered pollution as just another factor of production. This, of course, implied some expenditure on control, because less pollution meant more of other costly inputs. Here, however, the expenditure is made explicit in order to obtain an expression for the indicated pollution tax in terms of the cost of control. This has some advantages in interpretation and in comparing the costs of a tax with those of other methods, such as direct controls, but it sacrifices some

Baumol and Oates (1971, 1975*a*), Baumol (1972), and Mishan (1974). The clear nontechnical discussion by Ruff can be particularly recommended to noneconomists. A detailed empirical study of the comparative costs of taxes or effluent charges as opposed to uniform controls (discussed in the following text) to achieve a desired level of water quality in the Delaware estuary was discussed by Kneese (1977). The conclusion of the study was that the desired quality could be achieved for about half the cost with taxes.

detail in modeling the role of pollution within the firm, as we shall see.[21]

In symbols, the problem is as follows:

$$\text{minimize} \sum_i \sum_k p_i r_{ik} + \sum_k p_v v_k \tag{6.16}$$

subject to

$$f^k(r_{1k}, \ldots, r_{nk}) = y_k^* \qquad (k = 1, \ldots, h) \tag{6.17}$$

$$g^k(y_k^*, v_k) = s_k \qquad (k = 1, \ldots, h) \tag{6.18}$$

and

$$\sum_k s_k \leq s^* \tag{6.19}$$

where r_{ik} is the amount of input i and v_k is the amount of control input v employed by firm k, p_v is the price of v, y_k^* is the output of firm k, $g^k(\cdot)$ is a function that relates smoke emissions to levels of output and control for each firm, s^* is the environmental-quality standard, and the other symbols are as previously defined.

There are at least two features of this model that deserve further explanation. As indicated in equation (6.18), smoke emissions are determined by two things: the level of output, y^*, and the input, v, devoted to abatement or control. This formulation is not as rigid as it may seem, because the control input can be understood rather broadly as a method or technique for reducing emissions in combination with physical factors such as labor and capital. Just one such input is specified for simplicity without loss of generality.

A vector of outputs, the y_k^*, is specified, because otherwise the problem would be trivial. By having the firms produce nothing or very little, the planner obviously can minimize costs and satisfy the smoke constraint. What we are interested in are

[21] For an approach that treats pollution as an input, but is similar in other respects to ours, see the work of Baumol and Oates (1975*a*).

the conditions for minimizing costs associated with any given output, just as in the ordinary theory of the firm. The output actually selected presumably will depend on demand and on the planner's objective or the firm's objective. We assume only that it is desired to produce the chosen output at least cost, and we seek the conditions that will assure this. As before, we do not suppose that a planner can really determine input use at the firm level. We simply pose the problem in order to show how a much less ambitious approach, the setting of a (uniform) tax, can achieve the same results.

Proceeding with the solution, the Lagrangian expression can be written, first substituting $g^k(\cdot)$ directly for s_k, as

$$L = \sum_i \sum_k p_i r_{ik} + \sum_k p_v v_k + \sum_k \lambda_k [y_k^* - f^k(\cdot)] + \lambda \left[\sum_k g^k(\cdot) - s^* \right] \quad (6.20)$$

Differentiating with respect to the r_{ik} and v_k, and assuming no corner solutions, we obtain the necessary conditions for a minimum:

$$p_i - \lambda_k f_i^k = 0 \qquad (\text{all } i,k) \qquad (6.21)$$

$$p_v + \lambda g_v^k = 0 \qquad (\text{all } k) \qquad (6.22)$$

Now suppose the decisions on input levels will be made by the individual firms. The problem facing each is to minimize the sum of expenditures on inputs and a pollution tax, subject to the same restrictions on production and on the relationship between production and pollution. Note that our results will apply to imperfectly competitive firms as well, because we may assume that they are interested in keeping costs down, however much they choose to produce.[22]

[22] We must also assume that the firms are price takers in factor markets, importantly including the market for pollution. That is, the tax rate is not influenced by firm activities. This issue was further discussed by Bohm (1970) and Baumol and Oates (1975*a*). A potential difficulty with the factor price assumption is that after

Thus the firm's problem is as follows:

$$\text{minimize} \sum_i p_i r_{ik} + p_v v_k + t_k s_k \tag{6.23}$$

subject to equations (6.17) and (6.18). The Lagrangian expression, again subsituting $g^k(\cdot)$ for s_k, is

$$L^k = \sum_i p_i r_{ik} + p_v v_k + t_k g^k(\cdot) + \alpha_k[y_k^* - f^k(\cdot)] \tag{6.24}$$

where t_k is the pollution tax. Differentiating with respect to the r_{ik} and v_k, we obtain

$$p_i - \alpha_k f_i^k = 0 \qquad \text{(all } i\text{)} \tag{6.25}$$

and

$$p_v + t_k g_v^k = 0 \tag{6.26}$$

Comparing these conditions to equations (6.21) and (6.22), it is clear that they are the same, provided the tax t_k is set equal to λ, the shadow price of the pollution constraint, for all k.[23] It is clear that λ depends on the standard, s^*. For full efficiency, s^* would be set where the marginal damage from pollution just equals the marginal benefit, but this would bring us back to the preceding section's approach, which we have suggested is impaired by lack of information about damages.

Still, we have shown a great deal. Let us take stock. We have shown that a uniform tax on polluters ($t_k = \lambda$, all k) will achieve a preselected standard for environmental quality at minimum cost, provided the tax is set appropriately. It is important that the result emerges from the decentralized

imposition of a tax, the prices either may be changed or may no longer reflect real factor scarcities (assuming they did so in the original problem of social-cost minimization). My guess is that this difficulty is likely to be of little empirical importance.

[23] It must also be true that $\alpha_k = \lambda_k$. Because the equations and parameters are the same in both cases (provided $t_k = \lambda$), the solution values of the parameters, including α_k, must be the same. Away from equilibrium, α_k is, in general, not equal to λ_k.

decisions of the polluting firms. The central authority need know nothing about the control options facing each firm in setting the tax, and it need do nothing beyond setting the tax. On the other hand, to set the tax appropriately, the authority must solve for λ, the change in the minimum expenditure on production and control associated with a small change in the pollution constraint. This is a kind of aggregate marginal cost of control, and in practice it might be estimated from knowledge of the costs of an "average" polluter.[24] Even when this is not feasible, however, a uniform tax has the desirable property of minimizing the cost of achieving some quality standard, and doing so in a decentralized fashion.

To see this, consider the expression for the tax implicit in equation (6.26). Rewriting to make the tax explicit, we have

$$t = -p_v/g_v^k \tag{6.27}$$

The RHS is the price of the control input divided by its marginal product, or the marginal cost of control. Now, suppose the tax required to achieve a given quality standard, call it q^*, where q^* represents units of pollution abated and is inversely related to s^*, is not known. Instead, a tax is set that will result in a different quality, q^{**}. The marginal cost of control will still be equated across sources of pollution, because each will push control to the point that the marginal cost equals the common tax. This is shown for two sources with different control costs in Figure 6.3. A tax t^* will achieve the desired quality level q^* at the least cost, but a tax t^{**} will achieve q^{**} at the least cost.

The advantage of a tax over direct controls on emissions is easily demonstrated in this format as well. Suppose the two sources in Figure 6.3 are producing the same amount of

[24] For more on this suggestion, see the work of Kneese and Bower (1968). The reader seeking a discussion of some of the theoretical efficiency issues treated in this chapter, especially taxes versus direct controls, in a detailed and realistic setting might wish to consult the Kneese-Bower volume on the economics, technology, and institutions of water-quality management.

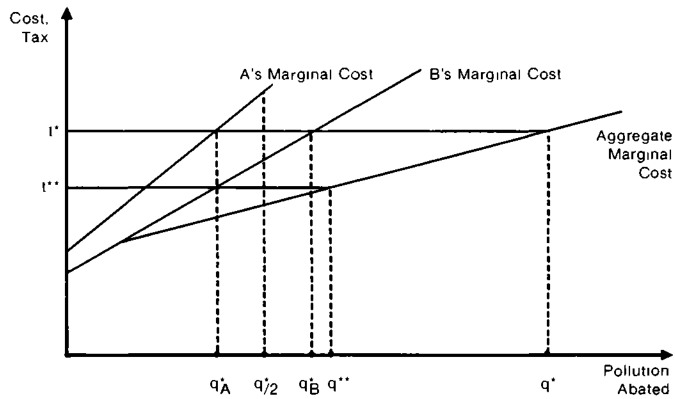

Figure 6.3. Minimum-cost tax.

pollution, before the tax or other control. Now it is desired to achieve a reduction to q^*. One obvious way to do this is to impose a uniform control on each source: a reduction of $q^*/2$. The difficulty is that in general this will result in violation of the cost-minimizing equimarginal outcome assured by the tax. As long as marginal costs differ, the cost of achieving q^* can be reduced by shifting a unit of abatement from the high-cost source to the low-cost source. Of course, a uniform reduction (which can also be stated in percentage terms for sources of different sizes) may have some appeal on grounds of equity, but it certainly will not be cost-effective.

Alternatively, the control can be tailored to the individual source to achieve the standard at least cost, as under the tax. In Figure 6.3 this will involve setting a standard of q_A^* for source *A* and q_B^* for *B*. The difficulty here is that the central authority will need to know the control cost function for each individual source. When there are just two sources, the difficulty may not be serious, although even in this case the incentive to misrepresent is clear. And when there are a great many sources, it is not realistic to imagine that the central authority can be informed about the types and costs of options available to each for controlling pollution.

Another advantage that has been claimed for a tax, as opposed to direct controls, is that the tax provides a continuing incentive to the polluter to cut back on emissions. No matter how low they already are, cutting back further will reduce tax payments. This may be especially important in a dynamic setting, where polluters are encouraged to seek new low-cost ways of cutting back.[25]

A disadvantage of a tax is that extensive monitoring of emissions is required. Thus far, we have tended to ignore the administrative costs of the policy alternatives. Yet it is clear (as noneconomists, especially, have argued in their attacks on the feasibility of a tax) that the real resource costs of monitoring can be substantial.

A first response to this criticism is that it appears to apply to direct controls as well as to other alternatives such as a subsidy or a permit system. Certainly this is true for controls on emissions, whether uniform or individually tailored. However, monitoring costs may be considerably lower for another form of control: a requirement that the polluter use a particular type of control technology. This is a popular approach in the management of both air quality and water quality in the United States.

My impression is that there is no reason to believe that mandated technology will be cost-effective, any more than other controls. Horror stories of almost perverse inefficiency in

[25] Kneese and Schultze (1975), in a nontechnical discussion of the history of air- and water-pollution policies in the United States and desirable changes in these policies, argued that the incentive for technical change in pollution control may be the most important criterion for judging a policy. Discussions of tax effects on control technology have been provided by V. L. Smith (1972), Orr (1976), and, most rigorously and comprehensively, Magat (1978). For a comparison of technical change under a subsidy for pollution control, as opposed to a tax, see the work of Wenders (unpublished). The conclusion is that a tax provides superior incentives.

specific instances are common knowledge among students of environmental economics.[26] And as the history of mandated emission-control devices on automobiles suggests, continuing inspection may be required to ensure that the devices are functioning properly (indeed, that they are in place and functioning at all). Prospects are perhaps better in other areas, but it is difficult to imagine a technology that does not require some monitoring. A fair conclusion might be that the question of which approach to pollution control accomplishes a desired degree of control at the least cost, including monitoring cost, is an empirical one. It is conceivable that there are cases in which a mandated technology will represent the least-cost alternative.

There are a couple of other situations in which direct controls may be an improvement over a tax. One occurs when the desired emission level is zero, as, for example, with a highly toxic substance. In this situation, a simple ban on use may be indicated.[27] A second situation occurs when rapid or temporary variation in emission levels is desired, as a consequence, for example, of changing weather patterns. Taxes, subsidies, and the number of pollution permits sold can all be varied to meet changing emission targets, but this may be impractical over short periods. Changing prices can be costly, and this is presumably one reason that peak or time-of-day prices are not more widely employed. An in-place tax system, on air pollu-

[26] For example, recycling, considered by many to be the ideal control technology, is not among the mandated technologies that qualify for water-pollution-control subsidies (Kneese and Schultze, 1975), with the result that the choice of technology is biased away from recycling. Similarly, low-sulfur western U.S. coal is discriminated against by the proposed New Source Performance Standards (NSPS) for coal-burning plants that mandate scrubbers. The advantage of the low-sulfur coal is that it does not need scrubbers to meet almost any reasonable ambient air-quality standard; yet this natural advantage is impaired by the mandate.

[27] For a detailed discussion of the alternatives for dealing with toxic substances, see the work of Portney (1978), which appears in an RFF volume containing articles on several aspects of U.S. environmental policy.