as circumstances under which such measures can, in fact, aggravate the country's short-run economic problems. Next, we consider what the rest of the world can do about a country whose economic activity generates externalities that are harmful to people outside its borders. In such a case, the usual free-trade argument may no longer apply, and an appropriate set of tariffs against the offending products may serve as a partial substitute for Pigouvian taxes; we find that nonzero tariffs are generally required for international Pareto optimality, taking into account the interests of the externality-generating country as well as those of the rest of the world.

In Chapter 17, the concluding chapter, we return to environmental decision-making in the domestic context. Here we examine an issue in regulatory federalism: Which level of government should determine standards for environmental quality? Should the central government set uniform national standards, or should "local" governments determine standards appropriate to their own jurisdictions? A purely economic view suggests an unambiguous answer to this question: Standards for environmental quality should balance marginal gains against marginal control costs, jurisdiction by jurisdiction, for pollutants that do not travel across the boundaries of the jurisdiction - standards for such pollutants should thus be local in nature. However, the political economy of local decision-making complicates matters. What if local agencies, in their eagerness to attract new business investment and jobs, reduce environmental standards to attract new firms? Will not the result be destructive interjurisdictional competition leading to excessive environmental degradation? Chapter 17 explores these issues and finds that for the basic case, local competition need not lead to inefficiently low levels of local environmental quality. Local-standard setting can lead to desirable outcomes. However, we find that there are various sorts of circumstances where fears that excessive pollution will result from local decisions are justified. Here, some constraints on local choice may well be justified.

CHAPTER 11

Efficiency without optimality: the charges and standards approach

The results arrived at in Chapter 8 may seem to constitute insuperable barriers to a rational environmental policy. The very presence of externalities is likely to produce a large number of local maxima among which, in practice, it seems impossible to choose with any degree of confidence; we may not even know in which direction to modify the level of an externality-generating activity if we want to move toward an optimum. It should be emphasized that these problems beset equally all attempts to achieve optimality by any of the means usually proposed - direct controls and centralized decision-making at one extreme and pricing schemes, such as the Pigouvian taxes and subsidies, at the other.

Nevertheless, we believe that it is possible to design policies for the control of externalities that are reasonably efficient. The approach that we will propose in this and the next chapter consists of the use of a set of standards that serve as targets for environmental quality coupled with fiscal measures and other complementary instruments used as means to attain these standards. The standards, while admittedly somewhat arbitrary, are, in principle, not unlike the growth or employment goals that have guided governmental macroeconomic policies. In both cases, employment and environmental policy, the approach is, in practice, basically of the "satisficing" variety, with acceptability standards based on individual judgments and, often, compromise. Yet, in both cases, the choice of effective *means* to achieve the established goals has been facilitated by a substantial body of economic theory. This theory suggests that fiscal measures can contribute to the efficiency of a program to control externalities. Moreover, the use of these fiscal measures in combination with standards for acceptable environmental quality, avoids, at least in part, the policy problems that have been raised in Chapters 7 and 8.

Much of the material in this chapter is taken from W. J. Baumol and W. E. Oates, "The Use of Standards and Prices for Protection of the Environment," which originally appeared in *Swedish Journal of Economics* LXXIII (March, 1971), 42–54 and was reprinted in P. Bohm and A. Kneese, Eds., *The Economics of the Environment: Papers from Four Nations,* London & Basingstoke: Macmillan, 1971.

Although in this chapter we emphasize the efficiency properties of effluent fees, we should not be taken to argue that this is always the best or the only way to deal with externalities. In the following chapters, we expand the analytic framework to allow the introduction of other policy tools and show that, under certain circumstances, an optimal environmental policy requires the use of several such measures.

1 Information requirements for optimization policy

The use of predetermined standards as an instrument of environmental policy recommends itself primarily because of the vast information required by the alternative approaches. Economists have long been aware of the enormous amount of information necessary to achieve anything that can even pretend to approximate optimality by means of centralized calculation. This is a major component of the Mises-Hayek argument against the potential effectiveness of full-scale central planning and direction. For the case of externalities, the argument is, if anything, strengthened by the analysis of Chapters 7 and 8, which emphasizes that data relating only to the neighborhood of an economy's initial position are particularly likely, in the presence of externalities, to lead the planner in the wrong direction.

Prohibitive information requirements not only plague centrally directed environmental programs, they raise similar difficulties for the calculation of optimal Pigouvian 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, and it is usually not easy to obtain a reasonable estimate of the money value of this marginal damage. There is a promising body of work applying a variety of techniques to the valuation of the damages from a polluted environment.[1] 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 difficulties in quantifying many of the most important consequences – the damage to health, the aesthetic costs – and the problems in determining a money equivalent for marginal net damage become quite 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

[1] For a useful survey of the techniques for estimation of the value of environmental amenities, see A. Myrick Freeman, *The Benefits of Environmental Improvement* (Baltimore: Johns Hopkins University Press, 1979).

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 fifty 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 twenty cents. As our results in Part I indicate, the correct value of the Pigouvian tax is twenty cents per unit of output, that is, the marginal cost of the smoke damage *corresponding to an optimal situation.* A tax of fifty cents per unit of output corresponding to the current smoke damage 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 that we have never experienced or even described in quantitative terms.

One alternative route toward optimality may seem to be more practical. Instead of trying to go directly to the optimal tax policy, as a first approximation, one could base a set of taxes and subsidies on the current net damage (benefit) levels. In turn, as outputs and damage levels were modified in response to the present level of taxes, the taxes themselves would be readjusted to correspond to the new damage levels. It might be hoped that this would constitute a convergent iterative process with tax levels affecting outputs and damages, these, in turn, leading to modifications in taxes, and so on.

Unfortunately, such an iterative process also requires information that is very difficult to acquire. At each point in the sequence of learning steps, one must be able to evaluate what the preceding step has achieved and to determine the directions to further improvement. But, knowing neither the relevant costs nor the incremental damages corresponding to each conceivable step, that is precisely what we cannot calculate. Because we are unable to measure social welfare, and because we do not know the vector of inputs and outputs that characterize "the optimum," we simply do not know whether a given change in the tax rate has moved us toward that optimum or has even been able to improve matters. There seems to be no general way in which we can get the information necessary to implement the Pigouvian tax-subsidy approach to the control of externalities.[2]

[2] There may be particular instances where careful analyses can produce some rough estimates of benefits and costs that can serve as the basis for a Pigouvian tax. For an interesting

2 The environmental charges 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. There is a fairly obvious way to avoid recourse to direct controls and retain the use of the price system as a means to control externalities: it involves the selection of a set of standards for an acceptable environment.[3] 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 sulfurdioxide content of the atmosphere in the city should not exceed x percent, that the oxygen demand of the foreign matter contained in a waterway should not exceed y, or that the decibel (noise) level in residential neighborhoods should not exceed z, at least 99 percent 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 standards that must be met in order to achieve what may be described as "a reasonable quality of life." The defects of this procedure are obvious, and, because we do not want to minimize them, we shall examine the problem of the choice of standards in a later section.

For the moment, however, we want to emphasize the role of the price system in the realization of these standards. The point here is simply that the public authority can impose a system of charges that would, in effect, constitute a set of prices for the private use of social resources, such as air and water. The charges (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. For example, one might tax all installations emitting wastes into a river at a rate $t(b)$ cents per gallon, where the tax rate, t, paid by a particular polluter, would, for example, depend on b, the BOD[4] value of the effluent, according to some fixed schedule. Each polluter would then be given a financial incentive to reduce

Footnote 2 *(cont.)*
application to the control of airport noise levels, see D. Harrison, "The Regulation of Aircraft Noise," in T. Schelling, Ed., *Incentives for Environmental Protection* (Cambridge, Mass.: M.I.T. Press, 1983), 41–143.

3 This proposal is not new. Most attempts to write a system of effluent charges into law are based on a set of prescribed standards.

4 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, but 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.

the amount of effluent he discharges and to improve the quality of the discharge (that is, 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.[5]

In marked contrast to an attempt at optimization, should iterative adjustments in tax rates prove desirable in a charges and standards approach, the necessary information would be easy to obtain. They require no data on costs or damages – only figures on current pollution levels. 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 might soon permit the authorities to estimate the tax levels appropriate for the achievement of a target reduction in pollution.[6]

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.[7] Of course, such an iterative process is not costless. It means that some of the polluting firms and municipalities will have to modify their 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 that may itself not be cheap.[8] But at any rate it is clear that, through the adjustment of tax rates, the public authorities can usually realize whatever standards of environmental quality have been selected.

3 Optimality property of the pricing and standards technique: cost minimization

Although the pricing and standards procedure will not, in general, lead to Pareto-optimal levels of the relevant activities, it is nevertheless true

5 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 v. Maryland,* 1819). In terms of reversing the process of environmental decay, the power to tax can be also the power to restore.

6 Of course, the political problems likely to beset either iterative process must not be minimized.

7 In this way, the charges and standards approach might be adapted to approximate the Pigouvian outcome. 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 might well arrive at the same end product, at least if the optimal solution were unique.

8 See A. G. Hart, "Anticipation, Business Planning and the Cycle," *Quarterly Journal of Economics* LI (February, 1937), 273–97.

that the use of unit taxes (or subsidies) to achieve specified quality standards does possess one important property: under appropriate conditions,[9] it is the least-cost method for the achievement of these targets.[10]

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 percent. 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 percent. However, a moment's thought suggests that this may constitute a very expensive way to achieve the desired result. If, at current 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 would be arrived at in practice by the regulator is not clear, because 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 percent. 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, because 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 rearranging smoke-reduction quotas: any alteration in this pattern of smoke emissions would involve an increase in smoke output by one firm

9 These conditions are spelled out later in this and the next chapters. Specifically, we will see in Chapter 13 that the presence of stochastic influences can sometimes make other instruments of control more efficient than taxes.

10 This proposition is not new. For some early discussions, see, for example, Kneese and Bower, *Managing Water Quality,* Chapter 6; and L. Ruff, "The Economic Common Sense of Pollution," *The Public Interest* XIX (Spring, 1970), 69–85. There is a similar proof by Charles Upton in "Optimal Taxing of Water Pollution," *Water Resources Research* IV (October, 1968), 865–75. The theorem takes no explicit account of metering costs which can, of course, be substantial. However, there seems to be little reason to expect these to be out of line with the enforcement costs associated with other environmental protection methods.

the value of which to the firm would be less than the cost of the corresponding reduction in smoke emissions by some other firm. 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 the next section.

It is significant that the validity of this least-cost theorem does not require the assumption that the firms generating the externalities are profit maximizers or perfect competitors. All that is necessary is that they minimize costs for whatever output levels they select, as would be done, for example, by an oligopolistic firm that seeks to maximize its growth or its sales, and that the market prices of the inputs reflect reasonably well the opportunity costs of their utilization.[11]

4 Derivation of the cost-minimization theorem

Let us turn now to a formal derivation of the optimality property of the charges approach that was described in the preceding section. 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.

Although this theorem may seem rather obvious (as the intuitive discussion in the last section suggests), its proof does point up several interesting properties. As already emphasized, unlike many of the propositions about prices in welfare analysis, the theorem does not require a world of perfect competition. It applies alike to generators of externalities who are pure competitors, monopolists, or oligopolists, 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 (that is, no influence on the prices of inputs) and so long as input prices approximate their opportunity costs. 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). Because the effective pursuit of these goals typically entails minimization of the cost of whatever outputs are produced, the theorem applies to whatever set of final outputs society should select (whether by central direction or the operation of the market).[12]

11 A similar argument suggests that the rationing of pollution by the sale of pollution licenses (rights) at a market-clearing price offers the same advantages in cost minimization. We shall demonstrate the validity of this argument in the next chapter.

12 The theorem may even be extended to certain agencies that are not cost-minimizers overall, but have incentives to minimize expenditures on pollution control. See W. Oates and

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

- r_{ik} represent the quantity of input i used by plant k $(i=1,\ldots,n)$, $(k=1,\ldots,m)$;
- s_k be the quantities of waste it discharges;
- y_k be its output level;
- $y_k=f^k(r_{1k},\ldots,r_{nk},s_k)$ be its production function;
- p_i be the price of input i; and
- s^* the desired level of $\sum s_k$, the maximum permitted discharge of waste per unit of time.

In this formulation, the value s^* 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 (for example, 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 (for example, 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 r_{ik} and s_k that minimize input cost for all firms together:

$$\min c=\sum_i \sum_k p_i r_{ik}, \tag{1}$$

subject to the output constraints

$$f^k(r_{1k},\ldots,r_{nk},s_k)=y_k \geq y_k^*=\text{constant} \quad (k=1,\ldots,m)$$

and the constraint on the total output of pollutants

$$\sum_k s_k \leq s^*.$$

It may appear odd to include, as a constraint, a vector of given outputs for the firms, because the firms will presumably adjust output levels as

Footnote 12 *(cont.)*
D. Strassmann, "Effluent Fees and Market Structure," *Journal of Public Economics* XXIV (June, 1984), 29–46.

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 that apply to *any* given vector of outputs no matter how it is reached.[13]

Using $\lambda_1,\ldots,\lambda_m$, and λ as our $m+1$ Lagrange multipliers, we obtain as Kuhn-Tucker conditions

$$\begin{array}{ll} \lambda-\lambda_k f_s^k \geq 0 & s_k(\lambda-\lambda_k f_s^k)=0 \\ p_i-\lambda_k f_i^k \geq 0 & r_{ik}(p_i-\lambda_k f_i^k)=0 \\ y_k^*-f^k(r_{1k},\ldots,r_{nk},s_k)\leq 0 & \lambda_k[y_k^*-f^k(r_{1k},\ldots,r_{nk},s_k)]=0 \\ \sum s_k-s^*\leq 0 & \lambda(\sum s_k-s^*)=0 \end{array} \tag{2}$$

for all i, k, where we have written f_s^k for $\partial f^k/\partial s_k$ and f_i^k for $\partial f^k/\partial r_{ik}$.

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_s. So long as its input prices are fixed, firm k will wish to minimize the cost of whatever output level it produces; that is, it will minimize[14]

$$c=t_s s_k+\sum_i p_i r_{ik} \tag{3}$$

subject to

$$f^k(r_{1k},\ldots,r_{nk},s_k)\geq y_k^*.$$

Direct differentiation of the m Lagrangian functions for our m firms immediately yields the first-order conditions (2); these are the same conditions

[13] The reason for prespecification of the vector of output has its analogue in the elementary theory of the firm. Where we use a cost-minimization premise in the analysis of the firm's input choices, it is obviously not correct to assume that it seeks to operate at as low a cost per unit as possible, without specifying its output level. For the firm's output level is determined by demand relationships as well as costs, and the output it decides to produce may be far from that which minimizes average costs. It is, however, reasonable to posit that whatever the output level it selects for itself, the firm will seek to produce it at as low a cost as possible. Our premise here is the analogue of this last assumption.

[14] Note again that this assumes identity between the prices in (1) and (3), that is, that input prices to the private firm correspond to the cost of their use to society. Thus, although our result does not require pure competition in the regulated firm, it does call for input prices that are not too far from their competitive values.

as before,[15] provided t_s is set equal to λ where λ (and hence t_s) is the shadow price of the pollution constraint – the marginal social cost of an increase in the stringency of the pollution standard.[16]

We have thus proved

Proposition One. A tax rate set at a level that achieves the desired reduction in the total emission of pollutants will satisfy the necessary conditions for the minimization of the program's cost to society.[17]

The preceding discussion indicates, incidentally, that pricing can play an effective role as a substitute for part of the information that is pertinent in the presence of externalities. In an illuminating remark, S. C. Kolm reminds us that the choice of efficient measures for the control of externalities requires, in principle, detailed information both about the benefits these measures offer the various members of the economy and the costs they impose on each of them.[18] The pricing mechanism offers no help with respect to the first of these because the very presence of externalities means that an individual decision maker's behavior does not reflect all of the relevant social benefits.

However, pricing does serve to eliminate the need for detailed cost information.[19] Under a system of central direction, a planner who wants to calculate the least-cost allocation of pollution quotas among the firms under his control must, as is shown in (2), have at his disposal data giving

15 The last of the Kuhn-Tucker conditions, $\Sigma s_k \leq s^*$, obviously has no counterpart in the calculation of the individual firm. However, it will clearly be satisfied if the s_k corresponding to a given set of prices is unique.

16 Clearly, the value of λ is an important datum and would be helpful in selecting a standard if that figure were available. Unfortunately, this information is lost in the standards and charges approach because no optimality calculation is carried out in the process. There are, indeed, no free lunches.

17 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, "The Search for Efficiency in the Presence of Externalities," in *Unfashionable Economics.* We should point out also that our proof assumes that the firm takes t_s as given and beyond its control. Peter Bohm in "Pollution, Purification, and the Theory of External Effects," *Swedish Journal of Economics* LXXII, No. 2 (1970), 153–66, discusses some of the problems that can arise where the firm takes into account the effects of its behavior on the value of t_s. See also our discussion in Chapter 6.

18 S. C. Kolm, "Économie de l'Environment" (unpublished manuscript), Chapter 2.

19 "This advantage, not needing to know the value of the right to pollute, is one of the great points of interest of the method of regulation by taxation (or subsidy). It is a property of *decentralization* of decisions: by requiring everyone to pay a financial charge equal to the damage he causes, one leaves the necessity of knowing the value of the right to pollute entirely in the hands of the person who knows it best – the polluter himself." Kolm, *ibid.*, Chapter 2, p. 4.

all of the f_s^k and f_i^k (that is, the marginal product figure for every input, *i*, and for every polluting plant, *k*). The herculean proportions of the task of collecting this mass of information and then carrying out the requisite calculations is clear.[20] A pricing approach dispenses with the need for all these data and computations because it gives that portion of the optimization calculation over to an automatic process. This suggests that the charges and standards approach may be looked upon as a procedure that frankly abandons any attempt to obtain extensive information on benefits but which uses the pricing system where it is at its best, in the allocation of damage-reducing tasks in a manner that approximates minimization of costs, even though detailed data on the costs of these tasks are unavailable.

5 Geographical and other appropriate variations in tax rate[21]

Even the cost-minimization claims for the standards and pricing approach must be qualified carefully. The theorem as stated runs into several problems in practice that may complicate its applicability.

One relevant assumption implicit in the preceding analysis asserts that there is a direct and additive relationship between the emission of pollutants and the degree of welfare loss suffered by the community. However, that is not always the case. A firm that emits waste into the upper parts of a river may do more or less damage to the community than one that discharges the same amount of effluent downstream. The upstream emissions may be less damaging than those downstream if the upper part of the river is sufficiently unpolluted to permit natural processes to disperse or degrade a considerable portion of the wastes before anyone is affected by them. On the other hand, if there is little natural cleansing of the upstream discharges, they may well be more costly to society than discharges into the lower parts of the river because people and activities along the entire length of the river may be affected primarily by upstream emissions.

Because the social damage caused by upstream and downstream discharges obviously differs, it is not appropriate to tax them at the same rate. In such circumstances, an equal tax per unit of effluent in the two

20 Although the calculation has ignored the costs of surveillance, obviously such outlays would be required under any system of environmental regulation. There seems to be reason to believe that, in many applications, the routine metering costs that would be needed will be considerably smaller than the costs of surveillance *and* judicial enforcement that are the instruments of direct controls.

21 This section is based on comments by Elizabeth Bailey and on two illuminating papers: Thomas H. Teitenberg, "Taxation and the Control of Externalities: Comment," *American Economic Review* LXIV (June, 1974); and Susan Rose-Ackerman, "Effluent Charges: A Critique," *Canadian Journal of Economics* VI (November, 1973), 512–28.

regions will generally *not* minimize the cost of a specified reduction in pollution as a simple counterexample demonstrates. Suppose that only the area near the mouth of the river is polluted so that the objective of the program is to reduce the level of pollution in that portion of the waterway. Suppose, moreover, that treatment of emissions will cost fifteen cents per gallon in a typical downstream plant but only ten cents per gallon upstream. Finally, assume that although all of the downstream firms' discharges add directly to the filth in the polluted part of the river, half of the upriver plants' discharges are eliminated automatically by natural processes. In that case, a tax of twelve cents per gallon of effluent will induce only the upstream plants to cleanse or reduce their emissions, because only their private costs of treatment per gallon are smaller than the tax rate. But to society this is an inefficient outcome, for ten cents nets it only a *half* gallon reduction in filth downstream, whereas treatment by a downstream plant would reduce pollutant discharge by a full gallon for only fifteen cents.[22]

Not only geographic accidents of location can lead to this problem. It may arise out of the range of decisions available to the firm itself, with the result that a uniform tax on discharges can induce management to make the wrong decisions. Turvey cites the case of a firm that has the option of building a high or a low chimney for its smoke.[23] If the high chimney can disperse pollutants sufficiently to render them harmless, it may yield the same contribution to human welfare as the suppression of smoke emissions and do so at a lower cost in resources. However, a tax based on emissions will clearly always favor smoke suppression rather than dispersion via higher chimneys, whatever their relative social costs.[24]

The upshot of all this is that, for the minimum-cost theorem to hold, it is necessary for the tax to be based on the *effect* of an emission on the community, and not necessarily on the amount generated. In practice this can sometimes be done in a rough-and-ready way (for example, by basing effluent charges on, say, two parameters – the quantity emitted and the quality of the receiving waters, or the amount of smoke emitted and on chimney height). Another device that may sometimes work reasonably well involves the establishment of different zones, based on concentration of population and current pollution levels, with different tax rates imposed in different zones.[25] Where some such simple provision will do the trick, the issues raised in this section create no insuperable difficulties for the charges and standards procedure. However, where delicate differentiations are essential, the attractive simplicity of the proposal can dwindle rapidly.[26]

[22] This is obviously a highly simplified illustration. Engineering models of waterways describing the differential impact on water quality of emissions at different locations use relationships that are much more complex. See Rose-Ackerman's discussion of the Delaware Estuary Model, "Efficient Charges: A Critique."

[23] Ralph Turvey, "On Divergences Between Social Cost and Private Cost," *Economica* New Series XXX (August, 1963), 309–13.

[24] Higher chimneys can, of course, lead to other sorts of problems, such as distant acid-rain.

One instructive way of looking at the matter is that differences in the effects of equal quantities of emission upon the effective level of pollution require the policy maker to retreat part way toward explicit evaluation of the social damage resulting from an emission. He must determine the extent to which various emissions influence the level of pollution. Note, however, that the charges and standards procedure still does not require the calculation of the effects of pollution on health, recreation, and psychic pleasure, and the translation of each of these into common (money) units.

What all this suggests is that, although the charges and standards procedure should never be as difficult to implement as the ideal Pigouvian tax,[27] it may still be quite complicated to take advantage of all the cost savings it offers in theory, in applications where the level of pollution damage responds differently to emissions from alternative sources or locations. The importance of this qualification obviously depends upon the circumstances at hand. As was just noted, where such differential effects of emissions are unimportant or where some simple device, such as variations in the charge by zone can deal with them (at least roughly), the charges and standards procedure retains its appeal.

The magnitude of the cost savings promised by more efficient systems of pollution control is quite large. There is now a substantial empirical literature encompassing a variety of air and water pollutants that provides estimates of the potential cost-savings from the use of pricing measures instead of direct controls.[28] These studies typically make use of simulation models for particular pollutants. Such a model has two basic components: a dispersion model that traces emissions from each source to the resulting pollutant concentrations at each receptor (or measurement) point in the air shed or waterway, and a set of control cost functions for the

[25] For further discussion of this proposal, see T. Tietenberg, "Spatially Differentiated Air Pollutant Emission Charges: An Economic and Legal Analysis," *Land Economics* LIV (August, 1978), 265–77.

[26] We shall explore this issue more systematically in the next chapter.

[27] Obviously, the ideal Pigouvian tax would also have to be adjusted for any differential effects of emissions from different sources.

[28] For a useful survey of these studies, see T. Tietenberg, *Emissions Trading: An Exercise in Reforming Pollution Policy* (Washington, D.C.: Resources for the Future, 1985), Chapter 3.

sources. With such a model, it is possible to simulate the outcomes under different systems of environmental management. The typical procedure is to calculate the cost of attaining some predetermined level of environmental quality under the current direct control system and then to compute the least-cost solution. A comparison of the two provides a measure of the excess costs under the current system. Such studies have generally found that the least-cost solution entails costs that are only a modest fraction of those under the current direct control system: The estimates range from a high of roughly 50 percent of current costs to less than 10 percent. A system that can approach the least-cost solution thus can typically promise very large cost-savings. As we have discussed in this section, it probably involves something of an overestimate of the potential savings to assume that a system of effluent fees can realize the least-cost solution, for any system in practice will involve administrative compromises that will prevent the attainment of the least-cost outcome. Nevertheless, existing studies indicate that the costs of current programs involve inordinately excessive costs; if fee systems in practice could at least go some distance in the direction of the least-cost solution (which surely must be true), the cost-savings would be very large.

6 The charges and standards approach and multiple local optima

In one important respect, the charges and standards approach avoids completely the problem posed for the Pigouvian solution and for central planning by nonconvexities and the resulting presence of a multiplicity of local optima. Because it is a satisficing procedure, it makes no attempt to search for an optimum, and so there is no occasion for the decision maker to aim mistakenly for what is in fact a local optimum instead of the global one.

So long as the emission of a pollutant is a monotonically decreasing function of the magnitude of the charge imposed on it, a function that is not bounded away from zero, one can choose a set of tax levels sufficient to guarantee attainment of whatever standards happen to have been selected. If the quantity of pollutant S still exceeds the level called for by the adopted standards, one need merely increase the charge upon the emission of S until its quantity has been reduced to the "acceptable" level, and that is all there is to the matter.

The presence of a multiplicity of maxima does, however, require one significant qualification of the cost-minimization theorem. For although, at least in principle, the use of charges guarantees that a given set of standards will be achieved at some sort of minimum cost, this may, in fact,

be a local rather than a global minimum. Suppose, for example, that there are two ways of avoiding the pollution produced by some commodity X, an increase in the output of smoke suppressors, or the substitution of another commodity, Y, which emits little pollution. Assume, moreover, that there are decreasing average costs both in the production of smoke suppressors and in the manufacture of Y. In that case, there will be two cost-minimizing ways of getting the pollution down to the desired level, the elimination of a sufficient amount of X and its replacement by a suitable amount of Y, or through the production of a sufficient quantity of pollution-suppression equipment. Toward which of these minima the market process will converge depends on the initial position, for that will determine the relative initial costs of Y and suppressors. There certainly is no guarantee that the process will converge toward the less costly of the two minima.

However, the likelihood that this problem will be encountered is apparently unrelated to the presence or absence of externalities. Unlike the issues discussed in Chapter 8, the multiplicity of equilibria that is relevant for the cost calculation does not seem to be made more likely by the presence of externalities. For the nonconvexities induced by externalities stemming from X arise both in the social production possibility set for X and the activity Z, that is damaged by it. But the externality caused by X need not affect activities W and V whose purpose is to offset the pollution produced by X. Thus, it need not introduce nonconvexities into the XW or the WV production sets, which are the production sets pertinent for the determination of the cost-minimizing program of pollution control corresponding to a given output vector. Consequently, although it is true that the cost-minimization property of the charges and standards approach can run into multiple maximum problems, it seems no more likely to encounter these difficulties than a decision process in some other economic area. There seems to be no special reason to expect it to run afoul of the nonconvexities that are built into the economy by the presence of externalities and which serve as booby traps that threaten the effectiveness of any attempt to design an *optimal* externalities policy.

7 Where the charges and standards approach is appropriate

As we have emphasized, the most disturbing aspect of the charges and standards procedure is the somewhat arbitrary character of the criteria selected. There does presumably exist some optimal level of pollution (that is, 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 easy way to determine accurately the set of taxes necessary to induce the optimal activity levels.

Although this difficulty should not be minimized, it is important to recognize that the problem is not 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 output where market prices fail to reflect the social damages (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 (that is, a method of collective choice) to determine the level of the activity. From this perspective, the selection 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.

Because methods of collective choice, such as simple majority rule or decisions by an elected representative, can, at best, be expected to provide only rough approximations to optimal results, the general problem becomes one of deciding whether 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, intervention in the form of acceptability standards can be utilized with a degree of confidence only where there is 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 that 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 Figure 11.1 by the curve *PQR*. In a case of this kind, although we obviously will not have an accurate knowledge 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 achieve that standard will lead to an increase in social welfare. For example, in terms of the curve *PQR* in Figure 11.1, the levying of a tax sufficient to reduce smoke outputs from level *OC* to *OA* to insure that the quality of the air meets the specified environmental standards would obviously increase social welfare.[29]

[29] The relationship depicted in Figure 11.1 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

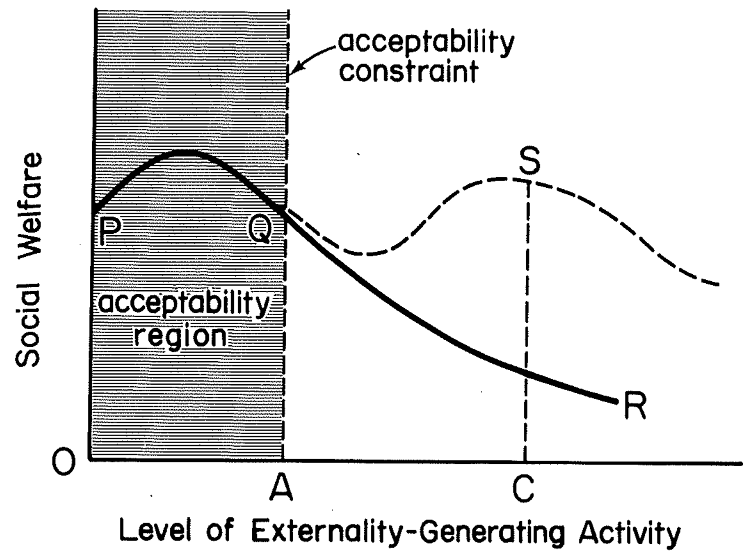

Figure 11.1

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 (for example, a move from *S* to *Q* in Figure 11.1) 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

usual sense; rather it measures, in terms of a numeraire (for example, 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 that is consistent with the specified minimum standard of environmental quality (for example, that level of smoke emissions from factories that 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 (for example, the smoke emissions of a number of different industries will determine the quality of the air). In this case, the acceptable level of polluting 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 implement the acceptability standards, there will result a least-cost pattern of levels of the relevant externality-generating activities. If we understand the constraint in Figure 11.1 to refer to the activity level indicated by this particular solution, then this ambiguity disappears.

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 (that is, the relevant portion of the social welfare curve in Figure 11.1 is not steep), and that the regulation can itself introduce inefficiencies into 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 is growing evidence that various types of pollutants do in fact have such unfortunate consequences, particularly in areas where they are highly concentrated. 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. 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.

CHAPTER 12

Marketable emission permits for protection of the environment

In the preceding chapter, we examined the case for a system of effluent fees for the attainment of a set of predetermined environmental standards. We found, in particular, that such a fee system has the capacity to achieve the standards at the least cost to society. There is, however, an interesting, and in certain circumstances, an appealing alternative to fees that also possesses the least-cost property: a system of marketable emission permits.

While economists were making the case for effluent fees in the 1960s, a political scientist at the University of Toronto, J. H. Dales, published a small volume in which he proposed, as an alternative to fees, a system of tradable property rights for the management of environmental quality.[1] Basically, Dales proposed that property rights be defined for environmental resources and then offered for sale to the highest bidder. For example, the environmental authority might create a limited number of permits for the discharge of a specified air or water pollutant; these permits would then be sold through some kind of auction.

It is easy to see that such a system can be used to achieve a specified environmental target and can do so, like fees, at minimum cost. In brief, the environmental authority can directly limit waste discharges to their target level by restricting the quantity of permits. As a market for these permits develops, a market-clearing price would emerge that (like a fee) will indicate to polluters the opportunity cost of waste emissions. Since all sources would face the same price for a permit, cost-minimizing behavior would result because marginal abatement cost would be equalized among these sources. This, as we have seen, is the first-order condition for the least-cost allocation of "pollution quotas" among sources.

In this chapter, we shall first discuss some of the attractive characteristics of the permit approach in a policy setting. We then turn to the issues of the design and operation of such a system.

[1] J. H. Dales, *Pollution, Property and Prices* (Toronto: University of Toronto Press, 1968).