4 Environmental Control under Uncertainty

4.1 Decision criteria under uncertainty

A great deal of discussion on environmental issues proceeds as though the consequences of actions are perfectly predictable. It is not uncommon, for example, to read an economic analysis of a fishery which contains no explicit mention of the fact that its growth function is only imperfectly known, or that there is possible disagreement about the size of the existing stock. This is often a necessary simplification. Stochastic models can rapidly become unmanageable, and an appeal to a computer in such cases often results only in the announcement of some numbers, with no accompanying insight about why the numbers have come out the way they have.

It should be recognized that the fact that an analysis contains no explicit mention of uncertainty does not necessarily mean that the analyst has pretended that there is no uncertainty. It could be that some kinds of averages of the various possibilities have implicitly been used in the discussion. The language of probabilities is the natural one to use in dealing with uncertainty, even although in the case of environmental problems the probabilities will often be *subjective* ones.[1] However, being subjective estimates, even

[1] For our purposes here I shall regard a probability distribution over various possibilities to be an *objective* one if there have been so many instances in the past that the probabilities can be estimated from the frequency distribution without too much difficulty (e.g. rainfall at a particular location). We use the term *subjective* probabilities for all others. The terms 'risk' and 'uncertainty' are often used to distinguish between these two cases. In what follows I shall use them interchangeably. For a good elementary account of statistical decision theory see Raiffa (1968).

experts disagree – often sharply – about the probable environmental effects of economic activities. Indeed, if there is a hallmark of environmental debates it is probably this.[2] But the fact that such probabilities are subjective is not an argument for not using the language of probabilities. Just as there is no single person whose estimates must always be relied upon, not every person's estimates of these probabilities are worth taking into account. For every new phenomenon there is some related phenomenon about which information is already available. Such evidence, in conjunction with pilot studies allow one to narrow down the family of distributions that might be used. As is invariably the case more information becomes available with the passage of time, so that the family is narrowed even further. The discussion in chapter 3, section 3.4, is relevant here. The best that can be achieved under such circumstances is to offer a range of policies that are optimal under the family of distributions. The social ranking of options will typically be only a partial ordering.

It is customary in welfare economics to encourage the government to accept the tastes *and* beliefs of individuals and then aggregate them in a suitable way. The fundamental theorem of welfare economics, referred to in chapter 2, is addressed to this kind of political environment. To be sure, it is recognized that it is desirable to make public various expert opinions and to enable individuals to base their beliefs on better information. The public provision of certain kinds of information, as we noted above, is one such implication. Nevertheless, the approach is to aggregate the *ex ante* 'preferences' of individuals, that is, preferences that incorporate individuals' tastes as well as their beliefs about various possibilities. Not surprisingly, a welfare optimum based on such an aggregation is called an *ex ante optimum*. Now, it may be asked why a government is required to respect individual *beliefs* about future possibilities in the same way as their 'tastes'. Tastes may be

[2] Aumann (1976) has shown that if two Bayesians hold the same prior beliefs about the occurrence of possible events and if their posterior beliefs about an event are common knowledge then these posteriors must be equal. To explain why environmentalists disagree one must therefore suppose either that they did not hold common priors or that their posterior beliefs are not common knowledge. There is a third possibility, of course, which is that they are not Bayesians!

refined or vulgar; nevertheless it may be held that a person's tastes must be respected. But beliefs can be *wrong*, and it is hard to see why it is undemocratic to disregard a person's beliefs if in fact they are wrong. It can no doubt be argued that we all have the right to take decisions on the basis of our own beliefs and to ignore evidence if it is psychologically convenient for us to do so (e.g. the hazards of smoking), and indeed that we have the right to make mistakes. But even if we accept this it is hard to see why a government ought to base *its* decisions on mistaken beliefs.

An alternative, therefore, is for the government to aggregate individual preferences over allocations at each state of nature – that is, their *ex post* preferences – and then aggregate these by the use of probability weights based on public information. A welfare optimum based on such an aggregation procedure is called an *ex post optimum* (also on occasion an Allais optimum; see Malinvaud, 1972).[3] This alternative therefore conveniently separates a person's *ex post* preferences from his beliefs and has the government respect the former but not necessarily the latter.

These foregoing arguments, and some further considerations that I shall develop in section 4.4, suggest that in many cases (most especially where public health is at issue) environmental protection is rather like a *merit good*, and so there is a case for the government to base its policies on only the most informed opinions.[4]

In what follows I shall, for expositional ease, assume that the government follows statistical decision theory and ranks options on the basis of their *expected net social benefits*. Now, even casual thinking on environmental problems alerts one to the fact that they sometimes involve a small chance of large-scale damage to society (or some large group) as a whole. However, these are precisely the kinds of problem statistical decision theory finds awkward to handle. It is possible to cast doubt on the plausibility of the 'expected utility hypothesis' in the case of risks that are

[3] In the *ex ante* case the government conducts an aggregation exercise only once, and in the *ex post* case twice. It is only in some restricted circumstances that *ex ante* and *ex post* optima are identical. See Broome (1981) and Hammond (1981a, 1981b) for deep explorations of these issues.

[4] 'The satisfaction of merit wants, by its very nature, involves interference with consumer preferences' (Musgrave, 1959, p. 13).

characterized both by 'low' probability and 'high' damage. On the other hand it cannot be claimed that there is anything as systematic and persuasive that can replace it. Seemingly appealing decision criteria, like 'maxi-min', display seriously unsatisfactory features when scrutinized.[5] 'Maxi-min' appears appealing precisely because it focuses uncompromisingly on the worst outcome associated with options, and in the field of environmental resources the worst may well be simply disastrous. Equally obviously, this feature of 'maxi-min' is its great weakness. A compromise, often resorted to, is to retain 'maxi-min's' distinguishing feature for worst outcomes by imposing constraints in the planning exercise by way of standards so that options that have the slightest chance of violating them are immediately ruled out – but otherwise to rank options by expected net social benefits. Likewise, really bad outcomes – even those with low probability – will be avoided by the expected social benefit criterion if the net benefit function is very steep at these points (see section 4.6 below).

4.2 Dependent versus independent uncertainties

It is convenient to distinguish between risks that are correlated across persons and those that are not. An increase in the emissions from automobiles in a region increases the chance that individuals will suffer from bronchial disorders. But to a reasonable approximation individuals face independent risks here depending on such personal factors as age and state of health. In contrast, the possible effects of massive deforestation on the global climate are jointly faced by all. Such risks are perfectly correlated, and the most extreme of these generate apocalyptic visions.

Environmental risks that are borne by individuals more or less independently of one another are somewhat easier to handle analytically, for one can appeal directly to the traditional theory of externalities. Moreover, economic theory tells us something

[5] The maxi-min criterion ranks options solely on the basis of their worst possible outcomes, no matter how low the probabilities of their occurrence (so long as they are positive) and no matter what the other possible outcomes are. A good account of different decision criteria under uncertainty is in Luce and Raiffa (1957).

about the relation between the ideal price of insurance against such risks and the risks themselves (see Malinvaud, 1972). It will be convenient first to draw out the distinction formally and see why it is the dependent case that usually generates the most acrimonious of debates, and ask whether one should expect this to be the case. We do this by means of a simple example. It will also clarify several other points that we have raised earlier.

Consider a group of N identical individuals. The representative person's valuation of his own income, we assume, can be represented by a function $U(Y)$, where Y is his income and where $U(Y)$ is an increasing and strictly concave function, as in figure 4.1; that is, $U'(Y) > 0$ and $U''(Y) < 0$. Strict concavity of $U(Y)$ means that he prefers a sure income to a lottery whose expected (or mean) outcome equals this sure income. Suppose that each person faces one of two possibilities: no damage, in which case his income level is $\bar{Y}$, or a damage, which is equivalent to an income loss L so that net income is $(\bar{Y} - L)$. If P is the concentration of pollutants, let $\pi(P)$ be the probability that the damage occurs. We naturally assume that $\pi'(P) > 0$.

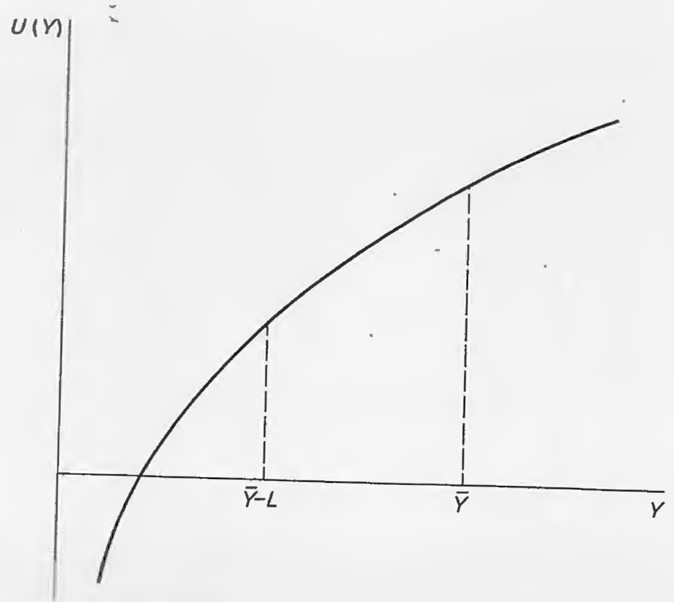

Figure 4.1

Consider first the case where individual risks are all independent of one another. If N is large then the fraction of people who will suffer the damage is pretty close to $\pi(P)$. In this case individuals can mutually insure themselves against the risk perfectly. If a person pays a premium of $\pi(P)\,L$ (his expected loss in income) and is guaranteed a net insurance payment of $(1-\pi(P))\,L$ in case of loss, he is guaranteed an income amounting to $\bar{Y} - \pi(P)\,L$ in either event. He is perfectly insured. The fact that N is large and that the risks are independent mean that such an insurance policy for each person is viable. Now suppose that we measure social benefits by the expected sum of the individual valuation functions. Then, under perfect insurance, expected social benefits is $E(B(P))_I$, where

$$E(B(P))_I = NU(\bar{Y} - \pi(P)\,L). \tag{4.1}$$

Suppose next that some activity increases the concentration level by a small amount, say, ΔP. Then the social damage caused by this can, on using equation (4.1), be expressed as

$$-\Delta E(B(P))_I \simeq N\pi'(P)\,LU'(\bar{Y} - \pi(P)\,L)\,\Delta P. \tag{4.2}$$

We come now to the other extreme case, where the risks are perfectly and positively correlated. In this case either *all* suffer the damage (with probability $\pi(P)$) or *none* suffers (the probability of this by definition $(1-\pi(P))$). In this case individuals cannot mutually insure themselves.[6] If P is the concentration level the expected social benefits, which we denote by $E(B(P))_D$, is

$$E(B(P))_D = N(1-\pi(P))\,U(\bar{Y}) + N\pi(P)\,U(\bar{Y} - L). \tag{4.3}$$

Suppose once again that some activity increases the concentration level by a small amount ΔP. On differentiating equation (4.3) with respect to P it is simple to check that the social damage caused in this case is:

$$-\Delta E(B(P))_D \simeq N\pi'(P)\,(U(\bar{Y}) - U(\bar{Y} - L))\,\Delta P. \tag{4.4}$$

[6] They may of course be able to insure themselves with some outside agency as, for example, would be the case with flood relief if financed through benefit taxation. We are considering the case where this is not possible.

Now notice that because $U''(Y) < 0$, $E(B(P))_I > E(B(P))_D$. This is precisely what one would expect since individuals are fully insured in the first case and not at all when the risks are fully correlated. But in estimating the social damage due to an *increase* in pollution one must compare the *declines* in expected social benefits – not the expected social benefits themselves. As regards this the matter is ambiguous, because without knowing what $\pi(P)$ is one cannot tell how $\Delta E(B(P))_I$ compares with $\Delta E(B(P))_D$. If $\pi(P)$ is 'small' (the chance that any given individual will suffer the loss is small), then $-\Delta E(B(P))_D > -\Delta E(B(P))_I$ and so an increase in 'collective risk' is the case to fear. Not so if $\pi(P)$ is nearly unity. In this case the social damage due to a marginal increase in pollution is greater when the risk is *not* collective – in the sense that the risks are independent.

While absurdly simple, the foregoing analysis suggests that there is no obvious reason why we ought to fear an increase in collective risks more than non-collective ones.[7] It also shows that it makes great sense to fear them more if $\pi(P)$ is small and the loss is 'large' – i.e. $U(\bar{Y}) - U(\bar{Y} - L)$ is large. But this is the case of a small collective risk of a mammoth social loss – precisely the kind of example over which people express their greatest anxieties.

4.3 Environmental research

One way to reduce risks to to spread them by choosing one's actions appropriately. Indeed, a good part of the early literature on the economics of uncertainty was concerned with exploring circumstances in which diversification pays, and with analysing the related question of how mutual insurance schemes enable a society to achieve this diversification.

A second way to reduce risks is to obtain further information on the uncertain areas. Pilot studies designed to investigate the environmental effects of pollutants (e.g. their effects on fisheries)

[7] I am using the terms 'collective' and 'non-collective' to describe the fully dependent and independent cases. In either case pollution is a public 'bad', in the sense of influencing everyone's chance of loss. Niehaus (1980) makes the observation that beyond a point the occupational and public risk of producing safety equipment exceeds the reduction of an existing risk.

and research designed toward discovering cheaper pollution abatement technologies are examples. This is not the place to discuss at length the strong *a priori* reasons for supposing that a market economy is unlikely to sustain the right amount of expenditure in obtaining such information and directing it along the right route (on this see for example Arrow, 1971b; Dasgupta and Heal, 1979). But it ought to be evident that 'knowledge' (or 'information') has the attributes of a public good. Thus there is a presumption that to the extent the producer of additional information cannot enforce property rights on the product there is a tendency towards insufficient production in a market economy. At the same time, however, there is a force operating in the opposite direction. Since knowledge is like a common pool (i.e. a common property resource), it is likely to be excessively *used* in a market economy.[8] The discussions in chapter 2 suggest therefore that knowledge, as an intangible good, ought perhaps to be subsidized in its production, and taxed in its use. Both these are on occasion observed in market economies. But there are obvious difficulties in implementing subsidies and taxes on such an intangible commodity as information. Indeed, arguments for the public provision of certain kinds of information, most especially the fruits of basic research, have been based on the non-appropriability of such commodities (see Arrow, 1971b). It is not an accident that government funds are usually involved in environmental research.

In this section we take it that the government can construct an (expected) damage function – a function that relates the (expected) social damage to the pollutant level. What we wish to argue here is

(1) that environmental research and development (R&D) projects often carry with them an insurance value, so that the social costs of risks associated with such projects are often *negative*
(2) that a government should not attempt a complete diversification among R&D strategies, even if they are uncorrelated, but instead would be well advised to specialize in only a few avenues of research.

[8] We are drawing attention only to the simplest points here. The matter is a great deal more complex. For some preliminary accounts of ~~the relation~~ between the structure of markets and the generation and use of information, see Dasgupta and Stiglitz (1980a, b).

Suppose society is uncertain about the precise effects of environmental pollution associated with specific economic activities. For example, it may not be known whether polychlorinated biphenyls (PCBs) have a large or small effect on marine food chains. Since research on this question is under way, only time will tell which is the case. The issue to be decided now is whether to undertake, for example, a research project to discover methods for breaking PCB molecules into harmless constituents. Even if such a project were successful, the realized social benefit from it will be high only if it is found that PCBs have a large detrimental effect on marine food chains; not otherwise. This is another way of saying that the social return on such a project is inversely associated with society's social income (i.e. national income corrected for environmental effects). This means that such a project provides society with insurance against adverse environmental effects. In this case, provided the variance of the project is not too great, a risk-averse society would prefer such a project to a sure project with the same expected return that was not environmental. In other words, society would prefer such an uncertain project to a sure project even if its expected social return is slightly less than the return from the sure project – the social cost of risk associated with the project is negative.

The second point is best illustrated by the observation that research activity in general is concerned with the acquisition of information. This acquisition requires the expenditure of resources, but not all information is worth this expenditure. Nor in general does the acquisition of information eliminate uncertainty, but this is no reason for not seeking it. Even when uncertainty is not eliminated, the information obtained may alter plans, and therein lies its value. Instead of taking an action in the absence of further information, one may wish to wait until more information is available. Of course, payment has to be made for this information. At the time one pays for the information (e.g. R&D expenditure) one does not know precisely what will be required (the outcome of the R&D project). However, one knows in advance that the optimal course of action will be based on the information acquired. The value of a research project is the expected social net benefit to be obtained from it. While we have provided a verbal account, the value of information can be repre-

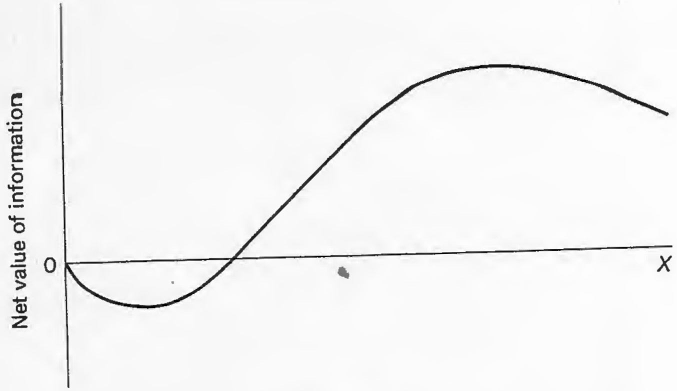

Figure 4.2

sented in a precise mathematical manner (see for example Marschak and Radner, 1972; or Dasgupta and Heal, 1979).

Suppose we were to represent environmental R&D projects of a certain kind by the degree of refinement in the experiments that define the projects. For example, all projects in the class so defined may be concerned with the effects of PCBs on marine food chains, but they may differ in the precision with which the investigator determines the effects. Suppose we were to denote the degree of precision (and therefore a research project) by x ($\geqslant 0$); thus a higher value of x denotes a more detailed experiment. We may represent the quantity of information by x; therefore, $x = 0$ is the most crude experiment of all – namely, no experiment! Suppose that the marginal social cost of information is positive. Then it can be shown (see Wilson, 1975) that under fairly weak conditions the net value of information declines in the neighbourhood of zero information (see figure 4.2). Thus if it is worth acquiring information of a certain type (as in figure 4.2), the amount of information ought to be no less than a certain positive level. But this means presumably that the decision-maker (here the government) ought to specialize in certain types of research.[9]

These arguments suggest the possibility of increasing returns in the value of information (which, incidentally, has nothing to do

[9] The possible non-concavity in the value of information was also demonstrated in an unpublished paper by Radner and Stiglitz (1975). For a simplified proof of the theorem, see Dasgupta and Heal (1979, chapter 13).

with possible lumpiness in R&D expenditure). Accordingly, a society ought to channel its R&D expenditure, not minutely into every possible avenue, but substantially into a few.

4.4 Planning mechanisms with dispersed information

In the preceding sections I have been concerned with certain aspects of what one might call 'games' against nature. In such situations the decision-maker is required to choose while uncertain about the true state of nature. What makes such 'games' relatively simple is that 'Mother Nature', it is generally thought, does not respond strategically to the decision-maker's choice. In the remainder of this chapter and in the appendix I study more complex games – those that arise in environmental management problems when there is an information gap between the decision-maker and those affected by his actions. This gap could, for example, be between the regulator and the firms engaged in environmental pollution (see sections 4.5–4.7 and the appendix). Firms typically will know more about abatement and clean-up costs than will the regulator, and may balk at providing correct information if it is not in their interest to do so. Information gaps presumably also exist between the regulator and the persons who are affected by pollutants. In fact it seems plausible that this latter information gap is even more difficult to close. While it is possible, at least in principle, for a regulator to discover firms' technological possibilities, such as abatement and clean-up costs (e.g. by engaging independent experts), determining the extent of a citizen's aversion to pollutants (i.e. learning about his mind) is an entirely different matter. Indeed, it is the potential impossibility of closing this latter gap – since individuals may well wish to give misleading information if asked – that has been central in discussions of the 'free-rider' problem in the theory of public goods (see for example Musgrave, 1959; Atkinson and Stiglitz, 1980).

The problem of devising appropriate incentive schemes in the face of information gaps has been a major concern of economists during recent years, and it is probably too early to attempt a rounded view of the findings. This is particularly so because

practical applications of some of the theoretical results have so far been rare.[10] In the remainder of this section I discuss the motivation and logic underlying such incentive schemes. I shall then illustrate some of the central issues by means of an extended example running through sections 4.5–4.7 and the appendix.

In any social organization there are certain pieces of information that are known (or which will become known) only by the individuals in question; that is, they are costly (or in the extreme, impossible) to monitor publicly. These *private* pieces of information include, (a) an individual's personal characteristics (e.g. his preferences, needs, or the pollution abatement costs for a firm, etc.); that is, what kind of an entity the agent *is*, (b) the actions that he takes (e.g. how hard a person works at his task), that is, what a person *does*, and (c) localized pieces of information about the state of the world – or certain aspects of specialized technological possibilities.[11] At the same time there are certain types of information that are publicly known, or which *can* be publicly observed at relatively little cost. These may be precise pieces of information (e.g. the amount of pollution emitted by a firm) or they may be statistical information (e.g. the age distribution in a given society at a given moment of time). Thus a planning mechanism essentially selects an outcome (i.e. an allocation of goods and services) which is a function of private decisions that are based on private information and public decisions that are based on publicly known information. The idea then is to choose among planning mechanisms on the basis of their outcomes as judged by the social welfare criterion that has been adopted.[12] The planning mechanisms I have looked at so far in this book are very simple examples of this.

[10] *The Review of Economic Studies* (Symposium on Incentive Compatibility), April 1979, presents a collection of theoretical essays, as does Laffont (1979). Experimental results are reported in Bohm (1972); Randall *et al.* (1974); Barnett and Yandle (1973); and Scherr and Babb (1975).

[11] In the insurance literature the terms *adverse selection* and *moral hazard* are used to characterize the problems raised by the first and second categories of private information respectively.

[12] For a general discussion of planning mechanisms with dispersed information see Dasgupta (1980).

In the voluminous recent literature on planning mechanisms with dispersed information the aim has been to devise schemes in which the government (or planner) invites individuals and firms to send in messages (suitably chosen) to the centre and at the same time publicly announces how the totality of received messages and public observations will be translated into public decisions (e.g. the rates of taxation on individuals and firms). If the set of admissible messages is identical to the set of possible types of private information the centre wants to know, the planning mechanism is called a *direct* one. Otherwise, it is called an *indirect* mechanism (see Dasgupta, Hammond and Maskin, 1979). In the appendix (section 4A.2) we shall provide an example of a direct mechanism in which firms will find it in their interest to tell the truth about their pollution abatement costs. An outstanding example of an indirect mechanism is the one provided by Groves and Ledyard (1977) in which individuals transmit quite abstract messages to the centre and the mechanism is so devised that equilibrium outcomes have the property of sustaining efficient allocations of public goods even although individuals' true underlying preferences for these public goods remain private information.

The *relevance* of such mechanisms for environmental planning can hardly be over-emphasized. To take only one example, they have an immediate bearing on pollution management problems when private damages are private information. But the question can be asked whether social-damage functions must invariably be based exclusively on individuals' perceptions of their private damages. One can argue that this is not how people inevitably view the matter. Policies that are guided by considerations of minimum food requirements of citizens are not necessarily based on individuals' perceptions of *their* personal needs; nor should they be so. Unquestionably, needs vary across persons. But it is surely right and proper for governments to aim at ensuring some standard of food intake for all citizens even if some can live on less. A society can take the view that all citizens have a basic right to enjoying a command over certain bundles of goods and services irrespective of what individual preferences are. Thus too with environmental issues bearing on health and the risk of death. In addition, as I have argued earlier it is difficult to see why the government must respect individual beliefs about various possibilities if they happen to be wrong. In other words, what I am trying to argue here is that the economist's desire for estimating individual 'willingness to pay' to reduce the risk of environmental damages may have become obsessive. Reliance on a social damage function which is not based exclusively on individual preferences is not an 'undemocratic' act. I shall take it in what follows that the social damage function (whether or not based on individual preferences) has been estimated and that the government is concerned with the choice of policies to influence the emission of pollutants by firms.

4.5 Taxation versus regulations under uncertainty

It is intuitively clear that the effects of optimum taxation and optimum regulation are unlikely to be the same when the planner faces uncertainty about matters that are relevant to the problem at hand. In what follows I use a simple formulation to see what the relevant effects are and how they need to be assessed. In particular I argue that if the resource in question displays threshold effects, then the optimal form of 'taxation' is more like a pure regulation than a pure tax; that is, the regulator ought to impose optimal effluent standards, and not optimal effluent charges.[13]

These points can be discussed most tellingly in the context of environmental pollution. To begin with, note that environmental effects of pollution usually take time to make their presence felt. Therefore, the uncertainty about the extent of social damage resulting from pollution will not be resolved until sometime in the future. The policy chosen today then must be independent of the resolution of this uncertainty.

For example, suppose were are considering a policy to restrict the discharge of PCBs by industries into the seas. To date we are still uncertain about the capacity of zooplanktons to absorb this effluent without undue damage. It is possible that this knowledge will be gained in the future, but today's decision about how much PCB ought to be discharged must be independent of this knowledge.

[13] By a 'pure' tax I mean a marginal tax rate which is independent of the rate of emission.

In fact the planner is usually uncertain as well about the cost of pollution control. It requires specialized knowledge. Moreover, abatement programmes take time to implement, and firms may not know today precisely what the costs will be. Furthermore, as I have argued earlier, even if firms know their true abatement costs, the regulator may not, and may not be able to elicit the truth from firms unless it is in their interest to tell the truth. Thus, today's decision on the amount of pollution permitted must be independent of the resolution of this uncertainty as well.

It is now clear why regulations (i.e. effluent standards) and taxes (i.e. effluent charges) are not identical in their effects. Recall that in the pure regulation scheme the planner selects the total quantity of pollution to be emitted. Firms are prohibited from polluting in excess of this. In the pure tax scheme, the planner imposes a constant tax rate for marginal units of pollution and individual firms then decide how much to pollute. Thus, for any given realization of the social-damage function (e.g. realization of the true threshold level of the resource being damaged by the pollutant) taxes encourage too little abatement if abatement costs are in fact higher than expected and they encourage too much abatement if they are lower. The problem is reversed for the

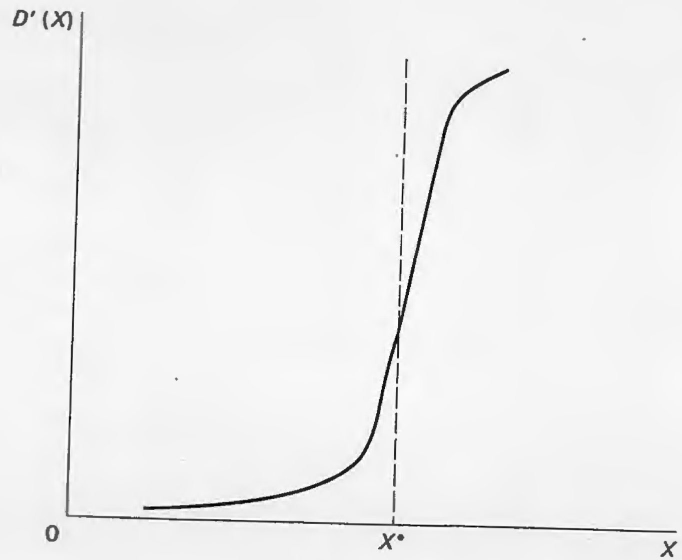

Figure 4.3

regulation scheme. Since the total quantity of pollution is decided by the planner in advance, it will be too little if costs are lower than expected and too much if they are higher.

Given that they are different in their impact, it is important to ask which is superior. As one would expect, the answer depends on the curvatures of the benefit and cost functions, and presently this will be confirmed by an example. But first it will be useful to obtain an intuitive feel for the proposition that it may well be desirable to rely on quantity restrictions rather than effluent charges (i.e. taxes) when the resource displays threshold effects.

As earlier, let X denote the total emission of a particular pollutant and let $D(X)$ be the social loss, in the sense of environmental damage caused by this emission. For the moment I am supposing that this loss function is known with certainty. Now suppose that the marginal loss function $[dD(X)/dX]$ takes the shape described in figure 4.3. Such a form seems plausible for a number of environmental problems, where X^* denotes the threshold level of pollution. That is, within a small neighbourhood of X^*, marginal damage due to the pollutant increases dramatically.

Now suppose that firms' abatement costs as a function of the quantity of pollution are unknown by the planner and are therefore functions of random variables as well. Regulation (i.e. the issue of a fixed quantity of licences to pollute) seems the better of the two schemes because the planner can ensure that the total level of pollution will be less than X^* – the level at which disaster strikes. Since cost functions are unknown to the planner, the only way to ensure against firms polluting beyond the level X^* via a pollution tax is to set a 'high' tax rate. However, a 'high' tax rate may be undesirable if there is a good chance that costs are lower than expected, because in such circumstances the amount of clean-up will exceed the amount desirable. There will be too little pollution!

The argument is still true if it is thought that there is a threshold level, and if the actual level is not known, only that it is within a range. Thus, for example, prolonged downgrading of the assimilative capacity of a medium caused by excessive levels of pollution can result in the instability of eco-systems. But it may not be known at what point the instability occurs. Here too the planner will wish to ensure against the possibility of disaster and guarantee that the level of pollution does not go much beyond the bottom value of the range. Once again, regulations can guarantee this, but pure taxes cannot, unless they are set at prohibitively high levels.

Regulations and pure taxes are polar types of plan instruments. They are an extreme case of tax functions defined on the level of effluent emission. A pure tax is a special case because the marginal rate is independent of the quantity discharged. A regulation is a zero tax rate up to the quota and an infinite tax rate for amounts in excess of the quota. It is seen that in the face of uncertainty the optimal tax scheme is one for which the tax rate is a function of the quantity of pollutants discharged and that in general it is neither of the two limiting forms just discussed. Nevertheless, for administrative reasons the planner may be forced to consider only the two limiting forms, and indeed, much of the debate on the appropriate form of intervention in economic activities involving environmental resources has centred on the relative advantages of pure taxation versus regulation.[14] The following section analyses these issues more formally in the context of a single firm emitting pollutants as a by-product of its production activity and looks briefly at the problem when more than one firm is so engaged. The appendix presents a formal analysis of this problem.

4.6 The case of a single firm

Consider once again a single firm whose net profit level as a function of the level of pollution X it emits is $B(X)$, and that the social damage from X is $D(X)$. Net social benefit $N(X)$ is taken to be $B(X)-D(X)$ and it is supposed that this is maximized at the pollution level X^* (see figures 3.1 and 3.2). It is clear that the optimum tax rate to impose on the firm is $t^* = D'(X^*)$.

Now suppose the regulator faces uncertainty about both the firm's net profit function $B(X)$ and the environmental damage function, $D(X)$. Regarding the former, it is natural to suppose that the firm knows its technology but that it is the government (regulator) which is uncertain about matters (e.g. abatement costs).

[14] Kneese and Schultze (1975) contains a good discussion of the relative merits of these polar modes which is based on institutional considerations.

Thus let $B(X, \phi)$ denote the net profit accruing to the firm when X is the level of emission and ϕ is the random variable reflecting the planner's uncertainty about the firm's technological possibilities. By hypothesis, the firm knows the true value of ϕ at the time it chooses X. However, I suppose that the planner does not know the true value of ϕ when it announces its policy.

Turning now to the social damage function D, I take it that the environmental consequences of the given pollutant are uncertain. Thus let $D(X, \theta)$ denote the social value of the damage sustained when X is the level of pollution and θ is the value of the radom variable reflecting the planner's uncertainty.

Suppose that the planner desires to choose that policy which will maximize the *expected* value of net social benefit, that is,

$$E[B(X, \phi) - D(X, \theta)] \tag{4.5}$$

where E is the expectation operator.[15] It should be noted that expression (4.5) is perfectly consistent with the regulator displaying an aversion toward risk, and we shall see this presently.

It will be supposed that the regulator can monitor the level of pollution X that the firm chooses to emit. For simplicity of exposition, assume that the random variables θ and ϕ are independent of each other. This is reasonable, since ϕ reflects uncertainties regarding the firm's technology, and θ reflects uncertainties regarding the effect of the firm's pollution on the environment. The regulator is interested in maximizing expected net social benefit (4.5). Moreover, he is aware that the firm, knowing the true value of ϕ, is interested solely in its net profits.[16] It is

[15] If $f(\bar{\theta}, \bar{\phi})$ is the probability that the two random variables realize the values $\bar{\theta}$ and $\bar{\phi}$ then by definition

$$E(B(X, \phi) - D(X, \theta)) = \sum_{\theta}\sum_{\phi} (B(X, \phi) - D(X, \theta))\, f(\theta, \phi).$$

I should emphasize that when I speak of the firm knowing its own technology I mean that it knows *more* than the regulator. Whatever uncertainty the firm faces about its own technology is assumed to be expected out in its own calculations.

[16] Thus the regulator and the firm pursue different goals. It is this feature that distinguishes the problem pursued here from the theory of teams. For a thorough discussion of the latter, see Marschak and Radner (1972).

then immediately apparent that the optimum policy consists in the planner imposing on the firm a pollution tax schedule $T(X)$ which is of the form:

$$T(X) = E[D(X, \theta)] \pm \text{constant}. \tag{4.6}$$

That is, the optimum pollution tax schedule is the expected social-damage function plus or minus a constant.[17] Faced with tax schedule $T(X)$, the firm will choose X so as to maximize

$$B(X, \phi) - T(X) = B(X, \phi) - E[D(X, \theta)] \pm \text{constant}, \tag{4.7}$$

that it, its profit, net of tax payment. Moreover, the regulator knows in advance that the firm will maximize (4.7).

Notice at once that when the regulator announces the tax schedule (4.6) he cannot predict precisely what the resulting level of pollution is going to be. This is because by assumption the regulator does not know the true value of ϕ, and as equation (4.7) makes clear, the firm's profit-maximizing choice of X depends on ϕ. Thus the imposition of the optimum tax schedule (4.6) results in an uncertainty about the amount of pollution that will eventually be emitted, thereby compounding the uncertainty about the final environmental damage. In order to maximize (4.7) the firm will choose that value of X at which marginal profit (excluding tax payment) equals the marginal tax rate.[18] That is,

$$\partial B(X, \phi)/\partial X = E[\partial D(X, \theta)/\partial X]. \tag{4.8}$$

From equation (4.8) it is clear that the profit-maximizing X, say X^*, is a function of the realized value of ϕ. That is, $X^* = X^*(\phi)$. The regulator by hypothesis does not know the true value of ϕ, but from equation (4.8) he can calculate the response function $X^*(\phi)$. Note as well that except for the limiting case where $D(X, \theta)$ is of the multiplicative form $Xg(\theta)$ (a case we have ruled out because we have supposed $D(X, \theta)$ to be strictly convex in X), the optimum pollution tax $E[D(X,\theta)]$ is not proportional to X. It follows that, except for this limiting case, the optimum tax rate on incremental pollution, namely, $\mathrm{d}E[D(X,\theta)]/\mathrm{d}X$, is not independent of the level of pollution.

[17] The 'constant' in expression (4.6) being by definition, independent of X, is essentially a lump-sum tax or subsidy, depending on its sign. Since by assumption there are no income effects its magnitude will not affect the outcome if we ignore distributional issues. In what follows the reader may wish to ignore the constant and suppose it to be nil.

[18] We are assuming that expression (4.7) is strictly concave in X for every possible ϕ. Moreover, the random variable has been so labelled that, without loss of generality, it is supposed that for each admissible value of ϕ, there is a unique solution of equation (4.8).

We conclude that in the presence of uncertainty about abatement costs, the optimal tax schedule, $E[D(X,\theta)]$, except for special cases, is neither a quota, nor a linear tax schedule (i.e. a marginal tax rate that is independent of the level of pollution emitted). That is, in the presence of uncertainty the control of environmental pollution is best conducted with the help of tax rates that vary with the quantity of pollutants discharged by a firm. Linear tax schedules and quantity regulations are merely suboptimal limiting forms of such policies.[19]

Let us conduct an exercise with this apparatus. Suppose that $\hat{X}$ is the level of pollution at which expected marginal social profit equals expected marginal social damage (i.e., $\hat{X}$ is the solution of the equation

$$E[\partial B(X, \phi)/\partial X] = E[\partial D(X, \theta)/\partial X]).$$

Now suppose that environmental uncertainties are small so that the social damage function $D(X, \theta)$ can be approximated around the level $\hat{X}$, in the form

$$D(X, \theta) = a_1(\theta) + D_1[X - \hat{X} - a_2(\theta)] + D_2[X - \hat{X} - a_3(\theta)]^2 \tag{4.9}$$

where D_1 and D_2 are positive constants, and where, without loss of generality, it is supposed that $a_1(\theta)$, $a_2(\theta)$ and $a_3(\theta)$ are random variables with zero expected values. From equation (4.9) it is then

[19] These two limiting forms were compared and contrasted in a seminal contribution by Weitzman (1974). Notice that for this example the imposition of $E(D(X, \theta))$ as the tax schedule is equivalent in its effect to that of the imposition of the optimum state-contingent pollution tax. Thus, let $X^*(\phi)$ be the solution of equation (4.8). Then define $t^*(\phi) = \partial B(X, \phi)/\partial X$, evaluated at $X^*(\phi)$. Then $t^*(\phi)$ is the optimum state-contingent tax rate. On state-contingent prices, see e.g. Dasgupta and Heal (1979, chapter 14). The example that follows in the text is that of Weitzman (1974).

apparent that

$$E[D(X,\theta)] = D_1(X-\hat{X}) + D_2(X-\hat{X})^2 + D_2E\{[a_3(\theta)]^2\}. \tag{4.10}$$

Thus, if the environmental damage function, $D(X, \theta)$, is of the form of equation (4.9), the optimal pollution tax schedule is of the form of equation (4.10). This is, of course, neither a linear tax nor a quota. As the third term on the right-hand side of equation (4.10) is a constant, we may as well ignore it (see footnote 12). The relevant terms depend on X, the first being linear in X and the second quadratic. Notice now that if D_2 is 'small', then in the neighbourhood of $\hat{X}$ the first term on the right-hand side of equation (4.10) dominates the second term, and so

$$E[D(X,\theta)] \simeq D_1(X-\hat{X}). \tag{4.11}$$

In this case the optimal tax schedule is approximately linear, with a constant marginal tax rate D_1. On the other hand, if D_2 is 'large', even a mild departure from $\hat{X}$ results in the firm being taxed heavily, as equation (4.10) confirms readily. Confronted with such a schedule, the firm will not wish to deviate unduly from the pollution level $\hat{X}$. The effect is then not dissimilar to the case in which the planner imposes $\hat{X}$ as a quota. Thus, if D_2 is large, the optimum pollution tax schedule resembles a quantity regulation.

The intuition behind these results ought to be clear enough. If the social damage function is of the form of equation (4.9), the marginal damage function is

$$\partial D(X,\theta)/\partial X = D_1 + 2D_2[X - \hat{X} - a_2(\theta)]. \tag{4.12}$$

We have already supposed by way of simplification that the uncertainty is 'small'. Thus the range of values $a_2(\theta)$ is permitted to take is small. From equation (4.12) it is apparent that the slope of the marginal damage function (i.e., $\partial^2 D(X,\theta)/\partial X^2$) is equal to $2D_2$. If D_2 is large, what equation (4.12) tells us is that marginal social damage increases dramatically with increasing pollution in the neighbourhood of an uncertain level of pollution, $\hat{X} + a_2(\theta)$. In other words, a large value of D_2 captures the fact that the pollution in question has a threshold effect. However, the threshold level of pollution, $\hat{X} + a_2(\theta)$, is unknown, with an expected value of $\hat{X}$. Therefore, if D_2 is large, the polluting firm, faced with a tax schedule of the form of equation (4.11) will choose not to pollute in excess of $\hat{X}$ and in fact will pollute at a level slightly short of $\hat{X}$ (this last, so as to pick up a small subsidy, as given by the first term of equation (4.10)).

One can also see from this example why the planner would go way off the mark in such circumstances if he were to rely on a linear tax schedule (i.e. a constant marginal tax rate). The point is that if the regulator is uncertain about the firm's technology – for example, if he is uncertain about abatement costs – the only way to ensure that the firm does not pollute beyond $\hat{X}$ is to set a high tax rate. But a high tax rate would be undesirable if abatement cost turned out to be lower than expected, because in such circumstances the amount of clean-up will exceed the amount desirable. There will be too little pollution! On the other hand, if D_2 is small, in the neighbourhood of $\hat{X}$ the marginal damage function is approximately constant (equal to D_1), as equation (4.12) makes clear. However, if marginal social damage is known and constant, it is obviously best to allow the firm full flexibility in finding the optimum level of pollution, since the firm by hypothesis knows the true value of ϕ and the regulator, by hypothesis, does not. A constant marginal tax rate (i.e. a linear tax schedule) allows the greatest amount of such flexibility, and so it is not surprising that if D_2 is small, the optimal tax schedule resembles a linear schedule. These considerations suggest that to the extent that environmental resources display threshold effects, the optimal tax schedule designed to limit their use resembles regulations governing amount of emission, and if for administrative reasons a choice has to be made solely between the optimum linear tax schedule and the optimum quota, the latter should be chosen.

4.7 The case of multiple firms

Where a single firm is engaged in causing environmental damage, it is a simple matter to compute the form of the optimum tax schedule. In the preceding section we noted in expression (4.6) that the regulator ideally should impose a pollution tax schedule which, up to an additive constant, is the expected value of the social damage function. The intuition behind this is rather obvious. For every possible realization of the random variable ϕ, such a tax schedule, if imposed on the firm, results in the firm's net profit function, equation (4.7), being identical with the social objective function. In other words, the tax schedule is so designed that the firm's objective (net of tax payment) coincides with society's objective.[20] In such a situation, the firm's profit-maximizing response cannot help but be the response that the regulator would ideally like the firm to make. Matters are more complicated when more than one firm is involved in damaging the environment. The problem is that the damage that any one firm imposes on the environment by marginally increasing its level of pollution discharge now depends, not only on its own level of discharge, but also on the levels discharged by others. This interaction, as we shall see, causes difficulties in the computation of the optimum pollution tax schedules, though the analysis is simple enough in general terms.

Suppose then that there are N firms engaged in emitting a specific type of pollutant. Firms are indexed by i or j ($i, j = 1, 2, \ldots, N$). Let X_i be the level emitted by firm i and let us suppose that the net profit function of firm i is $B(X_i, \phi_i)$, where ϕ_i is a random variable reflecting the regulator's uncertainty about abatement costs encountered by i. As in the previous section, we shall take it that at the time i chooses X_i it knows the true value of ϕ_i but that when the regulator announces his policy he is innocent of the actual value of ϕ_i. For ease of exposition suppose that the damage suffered by the environment due to pollution depends on the sum of the levels of pollution emitted by each of the firms and a random variable θ. Thus

$$D = D\left(\sum_{i=1}^{N} X_i, \theta\right).$$

The pollutant is, therefore, a 'public bad'. Expected net social benefit, we take it, is a direct generalization of expression (4.5), being the sum of the expected profits minus the expected social

[20] That is, the imposition of the tax schedule reduces the problem to a simple example in the theory of teams. For a pioneering discussion of incentives in teams, see Groves (1973).

damage, that is,

$$E\left[\sum_{i=1}^{N} B(X_i, \phi_i)\right] - E\left[D\left(\sum_{i=1}^{N} X_i, \theta\right)\right], \tag{4.13}$$

and it is this which the regulator is determined to maximize. For the remainder of this section we shall suppose, as before, that $B(X_i, \phi_i)$ is strictly concave in X_i for all ϕ_i, and that $D(X, \theta)$ is strictly convex in X for all θ, where

$$X = \left(\sum_{i=1}^{N} X_i\right)$$

is the total emission of the pollution in question.

Let us begin by analysing what the full optimum looks like. Suppose that at the time the firms choose their levels of emission, everyone (i.e. the N firms and the regulator) knows the realized values of the N random variables ϕ_i ($i = 1, \ldots, N$). Then, clearly, in order that expression (4.13) be maximized, firm i ($i = 1, \ldots, N$) should be made to choose that level of pollution X_i which maximizes

$$\sum_{i=1}^{N} B(X_i, \phi_i) - E_\theta\left[D\left(\sum_{i=1}^{N} X_i, \theta\right)\right]. \tag{4.14}$$

Suppose, without further ado, that it does. Then firm i should be made to choose that X_i at which

$$\partial B(X_i, \phi_i)/\partial X_i = E_\theta[\partial D(X, \theta)/\partial X] \qquad i = 1, \ldots, N \tag{4.15}$$

where

$$X = \sum_{i=1}^{N} X_i.$$

Equations (4.15) are N in number, and there are N unknowns X_i to be solved for. The ith equation in (4.15) says that the ith firm should pollute up to the level at which its actual marginal profit (left-hand side of equation (4.15)) equals the expected marginal damage due to aggregate pollution, where the expectation is carried out over the remaining random variable θ. Let $\tilde{X}_i$ ($i = 1, \ldots, N$) denote the solution of the system of equations

(4.15). It will be noticed that the realized values of the random variables ϕ_i ($i = 1, \ldots, N$) are parameters for this system of equations. Thus $\tilde{X}_i$ depends on the realized values of *all* the N random variables ϕ_j ($j = 1, \ldots, N$) and not merely on ϕ_i. Thus $\tilde{X}_i = \tilde{X}_i(\phi_1, \ldots, \phi_i, \ldots, \phi_N)$. This is precisely what intuition suggests. For, if it emerges that firm j's abatement cost at the margin is less than that of firm i, the former, in the interest of social welfare, should be forced to pollute less than the latter; the reverse if it emerges that i's abatement cost at the margin is less than that of j. It is therefore clear that the (full) optimum level of emission of firm i is of the form $\tilde{X}_i = \tilde{X}_i(\phi_1, \ldots, \phi_i, \ldots, \phi_N)$.

The question arises whether the full optimum can be enforced. Two distinct issues are involved here, and we have alluded to them earlier. First, it should be recognized that a knowledge of true abatement costs involves specialized technical knowledge, and while it is reasonable in many circumstances to assume that firms know their own abatement costs at the time they make their decisions, it is at least equally reasonable to suppose that the regulator does not. In this event it seems natural to allow the regulator to ask the firms to report their true abatement costs. But then recall that firms are interested only in their own private profits and not expected social benefits (equation (4.13)). If firms know in advance that their answers to the regulator's query will result in the enforcement of the optimum levels of pollution $\tilde{X}_i$, each firm will have a strong incentive to lie. Each firm would like to pretend that its marginal abatement cost is very high, reasoning that it will be allowed to pollute more than it would be allowed to were the truth known to the regulator. Its reasoning would be correct. Moreover, the regulator would know that this is how firms will reason. Therefore, he will know that the full optimum cannot be reached merely by calculating the functions $\tilde{X}_i$ and asking firms to divulge their private information. If the regulator wants the truth from firms in the environment we are considering, he must provide them with an incentive to tell the truth. The appendix to this chapter presents incentive schemes that will elicit the truth from firms.

This brings us to the second point, namely, that even if in principle the regulator can elicit the truth from firms, the cost of the transmission of this information from the individual firms to

the regulator typically will not be negligible: ϕ_i may be a large set of numbers. If such transmission costs are taken into account, it may not be sensible to try and reach what we have called the full optimum, $\tilde{X}_i$. It may be better that the regulator attempt to maximize expression (4.13) without asking firms to transmit their private information, but rely instead on information that he can obtain easily.

The appendix analyses the structure of pollution taxes which will maximize expression (4.13) when two-way communication between the regulator and the firms is barred (see section 4A.3). The remainder of this section looks at some simple regulatory policies that have often been proposed in the literature.

Conceptually, the simplest by far is a direct generalization of effluent standards which was discussed earlier in this section. For the present example it would mean a pollution quota imposed on each firm. Let $\hat{X}_i$ denote the optimum quota for firm i. Since the regulator is interested in maximizing expression (4.13), $\hat{X}_i$ must be the solution of the equations

$$\partial E[B(X_i, \phi_i)]/\partial X_i = \partial E[D(X, \theta)]/\partial X \qquad \text{for } i = 1, \ldots, N \qquad (4.16)$$

where

$$X = \sum_{i=1}^{N} X_i,$$

or in other words, where expected marginal profit to firm i equals the expected marginal social damage caused by aggregate pollution.

Now, in fact, there is a glaring defect with a scheme of this kind which the reader will have noticed immediately. The point is that social damage, by hypothesis, depends on the aggregate emission of pollution X. Therefore, if effluent standards are to be used, they ought to be imposed on the industry as a whole and not on each firm separately. It seems plausible that in the interest of expected social welfare it would be better if the regulator could devise a scheme in which firms chose their own levels of emission but were subject to the constraint that total emission must equal the optimum quota for the industry as a whole. The point is, of course, that the regulator ought to encourage firms with low clean-up costs to undertake more clean-up that those having high clean-up costs. But by hypothesis the regulator does not know a firm's actual clean-up costs, and so rule (4.16) does not allow for this desirable flexibility.

To make the point more clearly, suppose that the random variables ϕ_i are independent of one another and suppose furthermore that they are identically distributed. It is then clear from rule (4.16) that the optimum quota $\hat{X}_i$ is the same for every firm, say $\hat{X}$. But in fact the true value of ϕ_i typically will vary from firm to firm. It is this lack of flexibility in firm-specific emission standard schemes which has led authors like Dales (1968) to suggest an improvement, namely, a scheme in which the regulator selects the aggregate allowable level of effluent, which is then auctioned off to the firms in the form of licences. In our example, the regulator could sell $N\hat{X}$ licences to the firms.[21] If N is large, the resulting 'market' price for a licence would be akin to a pollution tax. At the resulting market price, say p^*, firms with high actual abatement costs naturally would purchase more licences than those with low actual abatement costs. This shifting of the burden of pollution control across firms is, of course, an improvement on firm-specific pollution quotas. It enables the regulator to retain control over the aggregate level of pollution and at the same time allows for a flexibility that firm-specific quotas do not display.

It will have been noticed that in the scheme just outlined the regulator sets the quota on aggregate pollution in advance. At the instant he selects $N\hat{X}$ he is uncertain about the fee, p^*, which will clear the market for these licences. At the opposite pole is a scheme in which the regulator announces a licence fee or pollution tax (or effluent charge) and firms are allowed to purchase as many licences as they like at the going fee. In this scheme the regulator does not know in advance the eventual level of pollution. The relative merits of these two schemes depend on much the same considerations that were mentioned in the preceding section. If the environmental resource in question displays threshold effects, the scheme in which the regulator auctions away the optimal number of licences is preferable to the imposition of the optimal linear pollution charge.[22]

However, one will recognize that it would be better still to allow more flexibility and have the regulator impose suitably chosen firm-specific pollution-tax schedules of the form $T_i(X_i)$. For the case of a single firm, locating the optimum pollution tax schedule was an easy enough matter and, as noted in the previous section, such a tax schedule in fact sustains the full optimum. Matters are a good deal more complicated here. Suppose that the ϕ_is are independent random variables. The point to note is that the imposition of the optimal tax schedules $T_i(X_i)$ will not sustain the full optimum. The reason is easy to see: by assumption firm i knows only the true value of ϕ_i and not of ϕ_j $(j \neq i)$. Thus faced with a tax schedule $T_i(X_i)$ and a knowledge of ϕ_i, firm i will choose X_i, which will be insensitive to the realized values of ϕ_j $(j \neq i)$. However, we have already noted in this section that the full optimum has firm i polluting at the level $\tilde{X}_i$, which is a function of *all* the ϕ_js. Thus firm-specific pollution tax schedules of the form $T_i(X_i)$ cannot be made to sustain the full optimum but they can be so chosen that in terms of social benefits (equation 4.13) they lead to better results than either the pure licensing scheme or the pure pollution tax scheme. The optimum forms of such tax schedules are studied in the appendix to this chapter.

[21] Or alternatively, each firm could be given $\hat{X}$ transferable licences. In terms of income distribution these two procedures would not be the same though. For a reasoned assessment of the implementability of such a procedure for air-quality control, see Tietenberg (1980).

[22] A further refinement over the aggregate effluent standard scheme which we have just discussed is one in which the regulator auctions away a fixed number of licences and at the same time announces that he will pay a fixed subsidy per licence purchased by firms in excess of their actual emissions. In this scheme the regulator has to compute two parameters in order to maximize expression (4.13), the number of licences to be issued and the rate of subsidy. When optimally chosen, this scheme is superior to the two we have already discussed. For details, see Roberts and Spence (1976) and Kwerel (1977).

Appendix: Imperfect Information and Optimal Pollution Control

4A.1 Introduction

This appendix addresses the problems raised in section 4.7 and analyses the manner in which they may be solved.[23] I take it that there are N polluting firms $(i, j = 1, \ldots, N)$, and that X_i is the level of pollution emitted by firm i. As before, I suppose that the private profit function of firm i is $B(X_i, \phi_i)$, where ϕ_i is a random variable whose realized value is known to i, and where B is strictly concave in X_i for all admissible values of ϕ_i. Social damage due to pollution levels X_i $(i = 1, \ldots, N)$ is given by the function $D(X, \theta)$, where $X = \Sigma_{i=1}^{N} X_i$ and where θ is a random variable. We take it that D is an increasing and strictly convex function of X for every admissible value of θ. Net social welfare is assumed to be given by expression (4.13) which I rewrite as

$$E[\Sigma B(X_i, \phi_i)] - E[D(X, \theta)] \tag{4A.1}$$

Expression (4A.1) is to be maximized. However, the highest attainable level of expression (4A.1) depends on the class of tax schemes that the regulator can choose from. Given that ϕ_i is a variable whose value is known in the first instance only by firm i, there are incentive problems in that firms typically would like to claim that their abatement costs are high at the margin (i.e. that their marginal profitability at high levels of pollution is high). So the question arises whether the regulator can devise tax-subsidy schemes which will neutralize such biases in incentives.

The following section presents tax-subsidy schemes which will enable the full optimum to be attained despite this incentive problem. It requires that the regulator receives messages from firms and then uses them to construct tax schedules that are imposed on firms. Moreover, firms are informed about how their messages will be translated into tax schedules. I shall suppose that firms do not collude. My task will be to show that the regulator can so devise tax schemes that, (a) firms will report their true profit functions (i.e. the true value of ϕ_i), and (b) they will choose the fully optimal pollution levels. In section 4A.3 I suppose that the transmission of messages from firms to the regulator is too costly, so that the regulator is forced to impose tax schedules on firms based only on his knowledge of the probability distribution of the ϕ_is and θ. I show trivially that the full optimum cannot be attained. However,

[23] This appendix is based on Dasgupta, Hammond and Maskin (1980).

I am able to locate the optimum structure of taxes given this communication constraint.

4A.2 Two-way communication

Imagine that the regulator asks the firms to inform him of their profit functions; that is, firm i is asked to report the true value of ϕ_i. The regulator informs the firms that their reports (i.e. the reported values of the N parameters $\phi_1, \ldots, \phi_N$) and their pollution emission levels will be used to compute taxes T_i which will then be imposed on firms. Moreover, firms are told in advance of the manner in which the reported values of ϕ_i and emission levels X_i will be translated into the N taxes, T_i. The idea is to construct tax schedules in such a manner that each firm finds it in its economic interest to report the truth irrespective of what other firms do. That is, the tax schedules are so constructed that truth-telling is a dominant strategy for firms. If this can be achieved, then each firm will tell the truth and in fact the full optimum can be attained. The way to construct such tax schedules is simple enough. The idea is to construct tax schemes in such a way that for every possible set of values of the parameters $\phi_1, \ldots, \phi_N$, the net profit for each firm (net of tax payment) coincides with the social objective function (4.14). We now see how this can be done.

In what follows I suppose firms do not collude. Let $\tilde{X}_i(\phi_1, \ldots, \phi_N)$, $(i = 1, \ldots, N)$, be the full optimum; that is, the solution of equation (4.15). We want to find tax functions $T_i(X_i, \phi_1, \ldots, \phi_N)$ for $i = 1, \ldots, N$ such that if $\hat{\phi}_i$ is firm i's actual parameter value, then for any possible announcement $\phi_i (\neq \hat{\phi}_i)$ it makes to the regulator, and for any pollution level $X_i (\neq \tilde{X}_i)$ it chooses and for any possible announcement ϕ_j that firm j $(j \neq i)$ makes,

$$\begin{aligned} &B[\tilde{X}_i(\phi_1, \ldots, \phi_{i-1}, \hat{\phi}_i, \phi_{i+1}, \ldots, \phi_N), \hat{\phi}_i] \\ &- T_i[\tilde{X}_i(\phi_1, \ldots, \phi_{i-1}, \hat{\phi}_i, \phi_{i+1}, \ldots, \phi_N), \phi_1, \ldots, \phi_{i-1}, \hat{\phi}_i, \phi_{i+1}, \ldots, \phi_N] \\ &\geqslant B(X_i, \hat{\phi}_i) - T_i(X_i, \phi_1, \ldots, \phi_{i-1}, \phi_i, \phi_{i+1}, \ldots, \phi_N) \end{aligned} \tag{4A.2}$$

for $i = 1, \ldots, N$.

If (4A.2) is satisfied, then each firm will announce its true parameter value and also find it most profitable to pollute at the fully optimal level, $\tilde{X}_i$. A set of tax schedules that satisfies (4A.2) is, of course, of the form

$$\begin{aligned} &T_i(X_i, \phi_1, \phi_2, \ldots, \phi_i, \ldots, \phi_N) \\ &= E_\theta \left\{ D\left[\sum_{j \neq i} \tilde{X}_j(\phi_1, \ldots, \phi_N) + X_i, \theta \right] \right\} - \sum_{j \neq i} B[\tilde{X}_j(\phi_1, \ldots, \phi_N), \phi_j] \\ &\pm \text{constant}. \end{aligned} \tag{4A.3}$$

The point then is this. While the regulator does not know the true values of ϕ_i $(i = 1, \ldots, N)$, he can compute the optimal levels of pollution $\bar{X}_i$ for every possible set of values of ϕ_i, by solving equation (4.15). He then asks firms to reveal their ϕ_is and announces that he will impose tax schedules on firms of the form of equation (4A.3). Firms will then be allowed to choose their pollution levels and pay taxes according to equation (4A.3). Since (4A.2) is satisfied for each i if (4A.3) is imposed, each firm will find truth-telling and the optimal level of pollution emission its dominant strategy. It should, however, be noted that the government is unable to balance its budget in this scheme. This is a pervasive problem with incentive schemes of this type (see Green and Laffont, 1977).

4A.3 One-way communication

It may be felt that the foregoing scheme is unduly cumbersome, requiring as it does the transmission of a great deal of information from firms to the regulator (ϕ_i will typically consist of a great many numbers). However, we continue to assume that the regulator can monitor the emission levels costlessly. Much of the literature on environmental control has in fact addressed itself to the problem of designing optimum tax schedules based solely on emission levels. As we recognized in section 4.7, such tax schemes cannot aspire to achieve the full optimum. I now present optimal tax schedules in those circumstances where the regulator does not receive any information from firms about their private abatement costs.

Suppose that the ϕ_is are independent random variables whose probability distributions are public knowledge. The regulator's aim is to maximize (4A.1) by imposing tax schedules of the form $T_i(X_i)$ on firms.

Let $\bar{X}_i(\phi_i)$, where $i = 1, \ldots, N$, be the solution of the problem of maximizing

$$B(X_i, \phi_i) + E\left\{\sum_{j \neq i} B[\bar{X}_j(\phi_j), \phi_j]\right\} - E\left\{D\left[\sum_{j \neq i} \bar{X}_j(\phi_j) + X_i, \theta\right]\right\} \quad (4A.4)$$

A comparison of expressions (4A.1) and (4A.4) immediately makes it clear that $\bar{X}_i(\phi_i)$ is the socially optimal level of pollution for firm i subject to the informational constraint that firm i's private information (i.e., the true value of ϕ_i) remains private. Thus $\bar{X}_i(\phi_i)$, $(i = 1, \ldots, N)$ sustains a second-best social optimum. We must now locate tax functions $T_i(X_i)$ for $i = 1, \ldots, N$, such that for any $X_i \geqslant 0$ such that $X_i \neq \bar{X}_i(\phi_i)$

$$B[\bar{X}_i(\phi_i), \phi_i] - T_i[\bar{X}_i(\phi_i)] \geqslant B(X_i, \phi_i) - T_i(X_i). \quad (4A.5)$$

for all admissible values of ϕ_i. If the regulator imposes tax functions which satisfy expression (4A.5) for all i, then the second-best solution can be derived. It is clear that a set of tax schedules which satisfies (4A.5) is of the form

$$T_i(X_i) = -E\left\{\sum_{j \neq i} B[\bar{X}_j(\phi_j), \phi_j]\right\} + E\left\{D\left[\sum_{j \neq i} \bar{X}_j(\phi_j) + X_i, \theta\right]\right\} \quad (4A.6)$$
$$\pm \text{ constant.}$$

The point to note about equation (4A.6) is that while the regulator does not know the true values of the ϕ_is, he can calculate the functions $\bar{X}_i(\phi_i)$ by differentiating expression (4A.4) with respect to X_i and setting it to zero; that is, by solving the N equations

$$\partial B(X_i, \phi_i)/\partial X_i = E\left\{\partial D\left[\sum_{j \neq i} \bar{X}_j(\phi_j) + X_i, \theta\right] \Big/ \partial X_i\right\}$$

where $\bar{X}_j(\phi_j)$, $j = 1, \ldots, N$, is the solution of the maximizing problem (4A.4). Therefore the regulator can compute the tax functions of equation (4A.6). It is of course, apparent from equation (4A.6) that in general $\mathrm{d}T_i(X_i)/\mathrm{d}X_i$, is not a constant. Nor is it a firm-specific quota.