5

Some Concepts from Welfare Economics

ROBERT DORFMAN

Robert Dorfman is David A. Wells Professor of Political Economy Emeritus at Harvard University.

The essays in this collection are concerned with the economic aspects of current environmental problems. There are other aspects—moral, medical, biological, chemical, and more—but they are treated here only to the extent necessary for understanding the economic aspects. Even with this limitation, the environment remains too large a concept and word to be analyzed in a single volume. We shall therefore confine ourselves to the aspects of the environment that people seem to have in mind when they talk of "the environmental crisis." In that context the subjects of concern are the purity of air and public waters, the plentifulness and vitality of natural landscapes, fauna, and flora, the integrity of certain other natural features such as beaches and, increasingly, widespread conditions, particularly the accumulation of greenhouse gases in the atmosphere, acid precipitation, and deforestation. All of these are comprehended in the *ecosphere*—the living space shared by all living creatures (including man)—and the creatures themselves.

Man has been tampering with the ecosphere for a very long time; one might almost define man as the animal that modifies its environment consciously, not instinctively. The great transformation from primitive hunting and gathering to settled civilization occurred when man learned to convert forests and savannahs into farms and to breed domesticated varieties of plants and animals. This was the most radical change in the environment that mankind has undertaken to this very day. Over the course of time this transformation altered the ecology of entire continents, and there is good reason to believe that ancient civilizations rose and fell as a result of its progress.

Not all of the environmental changes engineered by man have been for the worse, though some writers presume that they have been. Quite apart from the enormous increase in food production, malarial swamps have been drained, deserts have been made habitable, and much more has been achieved. But in the last few generations mankind's propensity to change his environment has itself been transformed, as symbolized by the contrasts between the whaleboat and the radar-equipped factory ship, the waterwheel and the nuclear power plant, or the country road and the interstate highway. The power to use and adapt has become the power to destroy abruptly. Our understanding of the environment has by no means kept pace with our capacity to alter it, and our ability to control our impact has fallen far behind. Therein lies the current threat.

The visible and impending environmental impacts of our newly acquired powers have forced us to recognize that the environment consists of scarce and exhaustible resources. That is where economics enters, for economics is the science of allocating scarce resources among competing ends. What has happened, obviously, is that the economic institutions that sufficed when ecological side effects were mild and gradual have abruptly become inadequate. It is as if we were trying to control automobiles with bridles. The task for economics is to perceive just why the old institutions cannot control the new forces, and to devise new methods of control that can. The papers in this collection are addressed to that task. First we shall examine what they say about the adequacy of the economic system we have inherited.

ENVIRONMENTAL RESOURCES AND PRIVATE PROPERTY

Several of our authors [Coase in Chapter 7, J. H. Dales in Chapter 14] lay great emphasis on the concept of private property and its role in directing the use of resources. Their point is twofold. First, the institution of private property and our other economic institutions have evolved together so that our economic system is well attuned to securing the efficient use of things that are owned. This follows from the familiar argument, going back to Adam Smith's "invisible hand," that competitive markets guide resources into the uses in which they will produce the things that consumers want most. But it should be noticed, first, that this argument applies only to resources that are privately owned and to commodities that consumers buy individually and use as they see fit. Second, if any resource is not privately owned, or if a consumer wants a commodity that he cannot procure and use individually, then the invisible hand doesn't work. On the contrary, ordinary economic institutions do not provide any incentives for furnishing such resources or commodities, or for using them efficiently if they occur naturally. In particular, resources in the environment, which no private person owns, tend to be used heedlessly, with results that are all too obvious.

One explanation, then, of our environmental problems is that many vitally important resources are not owned by anyone and consequently lack the protection and guidance that a private owner normally provides.[1] This state of affairs was tolerable when the unowned resources were plentiful, and careless use had virtually no effect upon them, but now we can no longer be so relaxed about it.

It does not follow that more extensive private ownership is the solution to our environmental problems. There are good reasons that explain why it is either impractical or undesirable to confer private titles to certain resources. Since these are the underlying causes of the misuse of the environment, and any corrective measures must take them into account, it pays to examine them closely.

The resources that make up the environment are unsuitable for private ownership because they lack the "excludability property." That is, it is typically not practical to exclude people from these resources or to prevent people from benefiting from them, either because of physical impossibility, or because controlling access would be inordinately expensive or cumbersome, or because limiting access would be socially unacceptable. This is the case with the atmosphere, and with most public roads, waters, beaches, and so on. It is also the case with flood-protection works, sewage-treatment plants, and many other facilities that improve the condition of the environment. There isn't much point in owning anything from which other people cannot be excluded.

Nonexcludability is not the whole story. Another peculiarity of environmental resources is that there are likely to be enormous economies in the joint consumption or use of the resource as contrasted with individual use. This second characteristic is illustrated by the contrast between housing and streets. It is economical to divide living space into family-size lots and devote each of them to housing one or a few families, but it would be fantastically wasteful to provide each family with its private road to the central business district. Accordingly, most houses are private property, but streets are public property. Most resources in the environment are analogous to streets, not houses, and are therefore used in common by substantial numbers of people. The difficulty that this common use creates is that each user may interfere with the others, reducing the serviceability of the resource to them; moreover, he has little incentive not to do so. Typical examples are road congestion and the use of the atmosphere and waters for discharging waste products. Resources that are most economi-

[1]There is no implication that resources that are owned are automatically used in the public interest. For example, if the owner of a resource is a monopolist, he will not use it in a socially efficient way.

cally used in common frequently are privately owned—ski slopes, for example. But frequently also, it is deemed undesirable to have such resources controlled by private individuals. It is not much help to notice that resources from which people cannot be excluded or which it is very economical for them to share would be better managed if they were ordinary, private economic goods. Their very nature precludes that solution.

Although the characteristics of nonexcludability and of shared, mutually interfering usage are especially prevalent among environmental goods, they do occur elsewhere. For example, fire protection and scientific discoveries have the nonexcludability property. Public and university libraries are used in common, and the users interfere with each other.

In fact, the economic theory of resources that exhibit these characteristics has been worked out in other contexts. We shall outline it briefly here.

This theory rests on two concepts. One is the notion of public goods, also called common-property resources or collective goods. The defining characteristic of a public good is that no member of the community who wants its services can be excluded from them if they are available to any other member. Classic examples are public roads and lighthouses. Environmental examples include the diversity of species, the oceans, and the ozone layer.

The other basic concept of the theory of public goods is externalities. An externality is a direct and unintended side effect of an activity of one individual or firm on the welfares of other individuals or firms. The word "direct" excludes effects that are transmitted by changes in prices. The word "unintended" excludes the effects of both benevolent and aggressive activities. Examples are air pollution generated by power plants or automobiles, water pollution, depletion of fisheries by overfishing, etc. Externalities are not always harmful, though economists are most concerned about the ones that are. The beneficial effects of TV broadcasting are a case in point.

These two concepts work together to produce an important classification of public goods. If the users of a public good neither interfere with each other nor increase the good's usefulness to each other, it is a "pure public good" (or "public bad" if it is harmful, like automobile exhaust). The scenic vistas and protective ozone layer, just mentioned, are pure public goods. If the users of a public good do affect its usefulness to each other, the good is said to be a "congestible public good." Garrett Hardin's commons and the fishery analyzed by Scott Gordon (Chapter 6) are congestible.

The mutual interference of the users of a congestible public good is a special case of the general phenomenon of externalities. An externality occurs when the activities of one person have unintentional effects on the welfares or production functions of other persons.[2]

Furthermore, congestible public goods are of two types. In one type,

[2]Externalities are explained at greater length in Chapter 17 and elsewhere in this collection.

the users are relatively homogeneous and impose reciprocal externalities on each other, like drivers on a crowded street. In the ohter type, there are two or more different kinds of user, and the effects of each type's activities on users of other types are not reciprocated. Thus waste dischargers on a stream impose externalities on swimmers and other users downstream, but are not affected by those users.

This elaborate classification of public goods is needed because the economic analysis and practical management of public goods is different for the different types. We turn first to the simplest case, pure public goods.

Since access to pure public goods cannot be controlled, by definition, the only decision is how much of the goods to provide. If the good occurs entirely naturally, as do oceanic shipping lanes, again there is no decision to be made. But if it has to be created at some trouble and expense, as does a flood-protection dam or a lighthouse, then no single individual is likely to find providing it worthwhile, since he cannot enforce a charge for its services.[3] A social decision therefore has to be made about whether to provide the good and, if so, on what scale.

To be concrete, consider a community consisting of three individuals, A, B, and C. They are all bothered by the smoke from the power plant that supplies their electricity, and the public good in question is a smoke precipitator that might be installed in the power plant's stack. None of them is willing to pay for the precipitator singlehandedly, but each would accept a moderate increase in his electric bill (or taxes) to defray the expense. Whether a precipitator should be installed and, if so, how large a one depend on its cost and on its value to each of the beneficiaries.[4]

The relevant data are shown in Figure 1. The size of the precipitator, measured in terms of the percentage of smoke removed from the discharge plume, is plotted horizontally; the vertical scale is a scale of dollars. The line labeled P is the marginal-cost curve. The smallest conceivable precipitator costs $4 per year to own and operate, and would achieve virtually no reduction in smoke emission. If 20 percent of the smoke is being precipitated, it would cost an additional $7.50 a year to raise the proportion to 21 percent. If 80 percent is precipitated, it would cost an extra $18 a year to raise the proportion to 81 percent. The other points are interpreted similarly. This line is effectually the supply curve for smoke removal. For example, if the consumers were willing to pay $18 for each percentage point of smoke removed, and if the power company could collect smoke-

[3]Exception to this assertion. Wealthy people sometimes provide public goods with no expectation of reimbursement. All charities and nonprofit enterprises depend on such altruism.

[4]Alternatively, this situation could be considered an instance in which four asymmetrical users (the power plant and A, B, and C) share the local airshed. The simpler interpretation, in which the smoke precipitator is a pure public good whose benefits are shared by A, B, and C, brings out the relevant points more clearly.

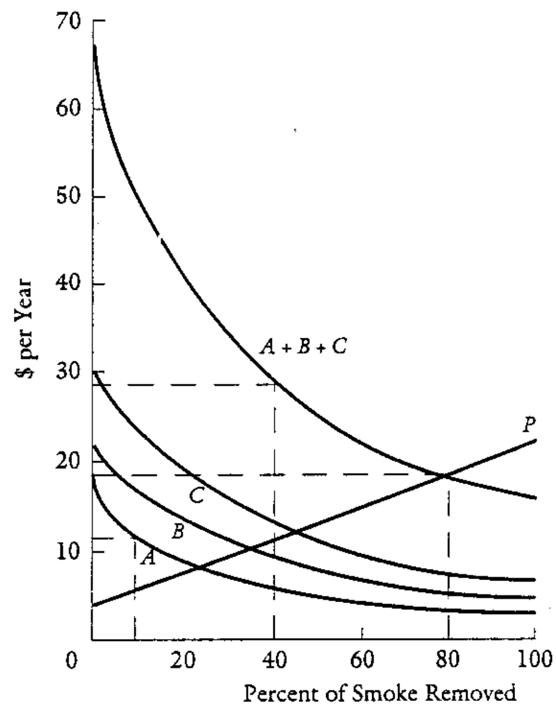

FIGURE 1. Optimal provision of a public good.

removal charges, the power plant would find it worthwhile to remove 80 percent of its smoke from its discharges.

The curves labeled *A, B,* and *C* are willingness-to-pay curves, one for each beneficiary. The curve for *A* shows how much he would be willing to pay for incremental improvements in the atmosphere. If only 10 percent of the smoke were being removed, he would be willing to pay $12 a year to have an additional 1 percent removed, and so on. If *A* were able to buy smoke removal as an individual, this would be his demand curve for it, telling how much he would buy at each stated price. Of course, he cannot do that, since the precipitator is a public good, and its services cannot be supplied to him without providing the same amount to others.

The willingness-to-pay curves for *B* and *C* are interpreted in the same way. To see how the curve labeled $A+B+C$ is constructed, consider removing 40 percent of the smoke. The willingness-to-pay curves show that *A* would be willing to pay $6.00 for this improvement, *B* would be willing to pay $9.50, and *C* $13.00, for a total of $28.50 for the whole community. Since all will benefit from the improvement, its value to the community is $28.50, which is shown as the ordinate of the $A+B+C$ curve at 40 percent. The entire $A+B+C$ curve is constructed similarly as the *vertical* sum of the three willingness-to-pay curves, and shows the aggregate willingness, on the part of all the consumers, to pay for additional reductions in emissions.

In the light of these interpretations, it is clear that the point where P crosses $A+B+C$ has special significance. That point is the level of smoke removal where the cost of an additional percentage is just equal to the amount that all the beneficiaries taken together are willing to pay for it. If less is removed, they would be willing to pay enough to cover the cost of a higher level of removal; if more, they would be forced to pay for a purer atmosphere than they deemed worthwhile. For these data, a precipitator that will remove 80 percent of the discharge is the right size.

That simple example illustrates the essence of the theory of pure public goods: the appropriate level to provide is the one for which the vertical sum of the beneficiaries' willingness-to-pay curves crosses the marginal-cost curve. Of course, there is much more to the theory than this, dealing with such additional matters as how the cost should be distributed among the beneficiaries, but we now have the main idea.[5]

The theory of externalities, including the special kind arising from the shared use of resources, is more involved, and we shall not attempt to summarize it here. The main issues are treated by H. Scott Gordon in Chapter 6 and Ronald Coase in Chapter 7. The basic notion is that since each user imposes a cost (in terms of inconvenience or reduced productivity of the resource) on the others, no use should be allowed unless the utility or benefit of that use is great enough to counterbalance the total cost that it imposes on other users. That simple statement slides over a great many difficulties, among them the problems of measuring both the direct benefits of the use and the costs that it imposes on others. It is the essence of environmental externalities that prices are not charged for access to environmental resources, so that there are no prices to reveal how much users are willing to pay for the privilege, nor are there any prices to indicate how much each user has inconvenienced the others. The selections in Section IV deal with these problems of benefit and cost measurement.

The main point of these concepts and theories for us is that they explain why environmental resources, which partake of these difficulties, tend to be overused, misused, and abused unless special measures are taken. They show that each individual's incentives induce him to use environmental resources more heavily and to contribute less to protecting them than is socially desirable. It should be emphasized that no malign intent or even

[5]For brief and clear discussions of the theory of public-goods provision, see R. A. Musgrave, *The Theory of Public Finance* (New York: McGraw-Hill Book Co., 1959), pp. 74–78, or L. Johansen, *Public Economics* (Chicago: Rand-McNally, 1965), pp. 129–140. The obvious ideal would be to charge each beneficiary in accordance with his willingness to pay, but if this were the policy, the consumers could not be expected to disclose their preferences with sufficient candor. (Contrast this with behavior with respect to ordinary commodities.)

blameworthy negligence is involved. It would be silly to stay home from Coney Island just because our presence would make it even more congested. No individual can make a noticeable reduction in the prevalence of smog by putting a catalytic converter on his car. The only meaningful solutions are collective solutions.

It follows that we have to make collective decisions about the use of the environment, and that one reason for its current state is that thus far we have neglected this obligation. We have to decide both what we want to achieve and what means to use for achieving it. Many of the papers, especially those in Section III, deal with this second problem, but since there are no selections dealing systematically with the first, we shall take it up here.

CRITERIA FOR ECONOMIC PERFORMANCE

So far we have taken it as accepted ground that the environment is not being used properly at present. We are putting too many nitrates and pesticides in the water, too much nitrous oxide and lead in the air, too many chlorofluorocarbons and too much carbon dioxide in the troposphere, and much more. Although it is obvious that what we are doing is wrong, it is by no means obvious what would be right. One thing is certain: we have to go on using the environment, and using it in common. People have to live and congregate somewhere, dispose of waste products, and even use depletable resources. These activities cannot be abolished, though they can be controlled. Controlling them means finding the proper balance between the utility of these activities to the individual and the disutility they impose, via the environment, on others.

It is easier to talk about the proper balance than to define it. Economists have hammered out useful formulations—useful in the sense that they can be applied to specific problems and decisions—only after long and earnest effort. And still, these formulations are not all that might be desired, as we shall see.

The critical difficulty lies in diversity of interest. What is good for Mrs. Goose is not necessarily good for Mr. Gander, and both of them have to be weighed in the proper balance. Reconciling divergent interests is what politics is all about, but economists have something to say, too.

Because of diversity of interest, it is very hard to say which of two policies or situations is better than the other unless, perchance, everyone happens to agree. An appeal to majority rule is a cop-out. It is a way to reach a decision in concrete instances, but it will not serve as a definition of the "right" decision, and no one maintains seriously that it can be relied on to produce the "right" decisions. Our task at the moment is to define, as well as we can, what we mean by a right decision as a foundation for policies for using the environment.

After generations of hard thought, economists have arrived at five criteria for judging policies or decisions. So many criteria are needed because none of them is entirely satisfactory, for one reason or another. Four of the criteria relate to the efficiency of the economic system, the fifth to equity. We shall discuss them all, since all are invoked in the papers that follow.

The four efficiency criteria consist of two pairs. The first, and fundamental, pair relates to the success of the economy in promoting welfare or satisfaction. We shall refer to these as the "utility criteria." The second, and more tractable, pair relates to the success of the economy in producing goods or other physical results. We shall refer to these as the "productivity criteria."

Within each pair there is one fully specific criterion that purports to pick out the best single decision or situation. Unfortunately in both cases the device that singles out the *optimum optimorum* is questionable. So there is a kind of fall-back criterion that identifies the class of decisions or situations in which the best must lie, if there is a best, but does not compare or evaluate the decisions in that class. That is how the four efficiency criteria are related logically. Now we can define them, beginning with the two utility criteria.

The Broad Utility Criterion: Pareto Efficiency

The task of an economy is to produce the combination of goods and services that will promote the welfare of the members of the community as much as possible with the resources and production techniques available. The welfare of the community is some resultant of the welfares of its individual members; these welfares are generally called their utilities.[6] The level of utility of each member of the community is presumed to depend on two things: his own consumption of private goods and services, and the environmental conditions to which he is exposed. We can compare the desirability of two or more different modes of operation of the economy by noting the amount of utility that each affords to each member of the community, and drawing conclusions on the basis of these individual utilities.

A simple diagram is frequently used to illustrate these concepts and make their application more concrete. Suppose that a community consists of only two individuals, Mr. Able and Mr. Baker. The entire output of the

[6]Philosophers sometimes call this assertion "radical individualism." It denies that the community, or the state, or any other group has an objective or welfare distinct from the welfares of its members. Of course, to the extent that members of a group identify with it and derive satisfaction from the attainment of group objectives (e.g., Olympic medals), that attainment is reflected in the welfares of members of the group. Such attainments are a form of public good.

economy is divided between them, and both are affected by whatever public goods or environmental conditions result from the operation of the economy. Then we can depict the welfare results of the economy on a diagram like Figure 2, where Mr. Able's utility is measured along the horizontal axis and Mr. Baker's is measured vertically. Point *A*, for example, represents the results of some mode of operating the economy in which Able's utility level is 100 units and Baker's is 150 units. It may be possible to change the pattern of economic activity so as to increase Able's welfare without diminishing Baker's. For example, Baker may enjoy playing his stereo equipment at a resounding volume that interferes with the yoga meditations to which Able is devoted. Able might be willing to pay Baker as much as $100 to use earphones instead of stereo speakers (i.e., to turn over to Baker $100 worth of the private goods and services to which he would otherwise be entitled), and Baker might be willing to comply for as little as $50. In this case, if Able paid Baker $50, Baker would lose some utility by resorting to earphones but would gain it back by consuming additional private commodities and be just as well off as before, whereas Able's utility would be increased $100 worth by the reduction in the noise level and would be diminished only $50 worth by his reduced consumption of other commodities, so that he would be better off than before. This result is shown in the figure by point *B*. Point *B* is indubitably a better point for the operation of the economy as a whole: one of the members of the community has benefited and none has suffered. Coase analyzes this type of social arrangement at some length in Chapter 7.

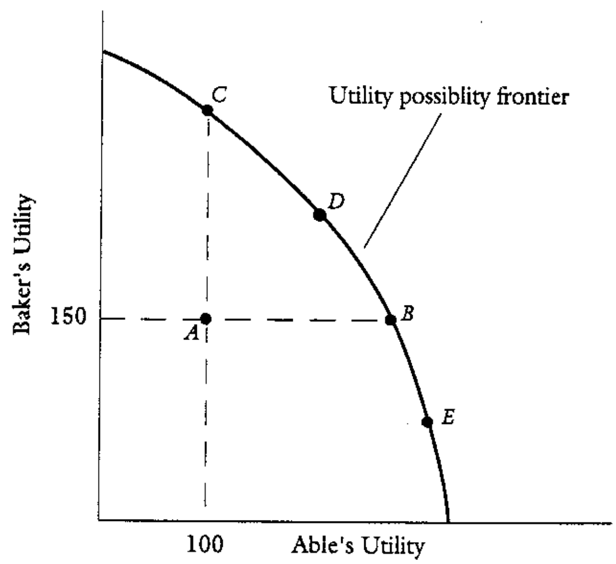

FIGURE 2. The utility-possibility frontier.

Let us suppose that when the economy is at point *B* there are no further possibilities for increasing Able's utility without reducing Baker's. Then the position of the economy corresponding to point *B* is said to be *Pareto optimal* or *Pareto efficient.* As a formal definition, the operation of the economy is Pareto efficient if there is no way to change it that will make some member or members better off by increasing their utilities without reducing the utilities of some other members. The curve on the diagram represents all the Pareto-efficient ways of operating the economy. Each point on it shows the greatest utility that can be provided to Able in conjunction with the given level of utility for Baker, and at the same time the greatest utility than can be afforded Baker in conjunction with the given level of utility for Able. It is called the *utility-possibility frontier.* The utility-possibility frontier can be drawn only for the trivial case of a two-person community, but it can be conceived, and does exist theoretically, for every community.

Now let us go back to point *A,* which lies below the utility-possibility frontier. If the economy is operating there, it is not doing as good a job as possible. For, by assumption, it can be operated so as to attain points such as *B* or *C* or any point on the frontier between them that will make at least one member of the community better off without harming any one. A large part of welfare economics, and of the papers in this collection, is concerned with social arrangements that will enable a real economy to avoid points like *A* that are inefficient from the utility point of view, and enable it to attain points like *D* that are efficient.

One method for resolving the conflict of interest between Able and Baker would be for the community to require Baker to use earphones. This would increase Able's utility even more than if he had to pay for silence, but it would reduce Baker's. The result might be point *E,* which is drawn on the utility-possibility frontier. This assumes that this social arrangement is Pareto efficient, as we defined it. But, to illustrate a terminological pitfall, the *change* from *A* to *E* is not Pareto efficient, because a member of the community is harmed thereby.

The comparison between points such as *C* and *E,* or between social arrangements such as free bargaining between Able and Baker as against government regulation, raises important social issues. As we have drawn them, both social devices lead to Pareto-efficient points, but there is good reason to believe that free bargaining is more likely to produce this result than governmental decrees. We shall consider that issue when we take up various practical methods of control. For the moment we note that the Pareto efficiency criterion does not enable us to discriminate between such points. The next criterion does so.

The Sharp Utility Criterion: Social Welfare

We noted above that the welfare of a community is some resultant of the welfares of its individual members. Suppose we could agree on precisely what the resultant was. Then we should have a *social-welfare function,* that is, a rule by which we could evaluate the welfare of an entire community if we knew the welfares, or utilities, of its individual members. For example, the social-welfare function might state that the welfare of the community is the simple sum of the welfares of its individual members.[7] A slightly more complicated social-welfare function would be the sum of the individual utilities minus one-half the difference between the greatest utility and the lowest one.[8] This social-welfare function would pay attention not only to the aggregate of utility but to the amount of inequality in the society. Any real social-welfare function would undoubtedly be very complicated indeed, and in fact, so far as we know, no satisfactory social-welfare function dependent on individual utilities has ever been constructed. There is good reason to believe that the social-welfare function is a philosopher's stone that does not exist in any real society.

In spite of the difficulties with the concept of a social-welfare function, it is very appealing because it provides the only way to arrive at clearcut, unambiguous social evaluations. Look again at our diagram. Point *D* is clearly socially superior to point *A,* since everyone benefits from a movement from *A* to *D.* But is point *E* socially superior to *A?* It is Pareto efficient, while *A* is not; but, nevertheless, Baker is likely to feel definitely worse off. And Baker is half of this society. Without a social-welfare function it is not even clear that a Pareto-efficient operation of the economy is superior to an inefficient one; and it is even more impossible to compare the merits of a number of points on the utility-possibility frontier. So some people conclude that since social decisions are made in practice, some social-welfare function must be implicit in the social-decision process.

If there were a social-welfare function, we could portray it like an individual consumer's utility map. This is done in Figure 3. Figure 3 shows the same data as Figure 2 plus a number of "social-indifference curves,"

[7]This is an important criterion historically; it is the basis of the utilitarian philosophy of social welfare.

[8]One restriction has to be mentioned. The value of the social-welfare function has to increase whenever the utility of any individual increases, other utilities remaining the same. For example, the following is *not* a social-welfare function: the sum of individual utilities minus twice the difference between the greatest and the lowest. The reason is that this function would actually decrease if the happiest individual became still happier, all others remaining the same. Because of this restriction, it is very dangerous to try to include such considerations as altruism or envy in a social-welfare function. Strange paradoxes can result—for example, a recommended distribution of income in which the most envious people receive the highest incomes and the most generous the lowest.

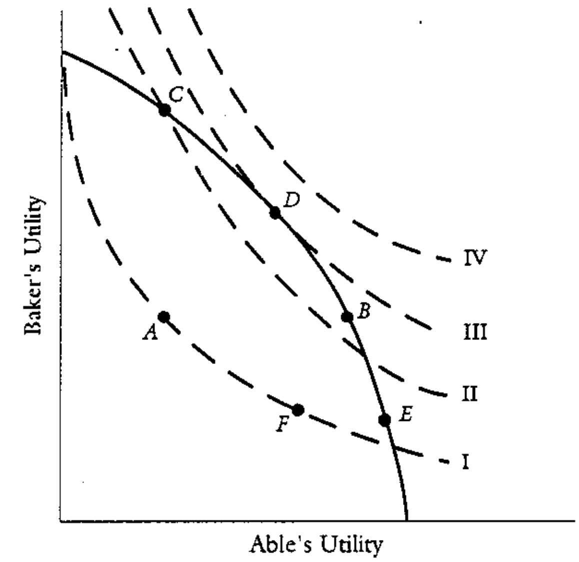

FIGURE 3. The utility-possibility frontier with social-indifference curves.

distinguished by Roman numerals. The value of the social-welfare function is constant along any social-indifference curve and is higher on higher indifference curves. Since points *A* and *F* are both on indifference curve I, they are equally desirable socially according to the social-welfare function depicted. Point *E* is above that social-indifference curve and is superior to either of them. Points *C* and *B* are better still. Point *D* is the most desirable point on the utility-possibility frontier. The points on social-indifference curve IV are even more desirable, but they are unattainable. From this diagram one can conclude that according to this social-welfare function it would be better to let Able and Baker bargain than to forbid Baker to use his loudspeakers, but that it would be better still to adopt the social arrangements that lead to point *D,* whatever they are. As the example illustrates, if there is a social-welfare function, the social desirability of any two configurations of individual utilities, Pareto efficient or not, can be compared.

But we have already cast doubt on the existence of such functions.[9] Indeed, there is even good reason to doubt that individual utilities can be defined or measured, either in principle or in practice. So criteria are needed for assessing economic performance without appeal to these dubi-

[9]For a more profound and decisive critique, see Kenneth J. Arrow, *Social Choice and Individual Values,* second edition (New York: John Wiley & Sons, 1963).

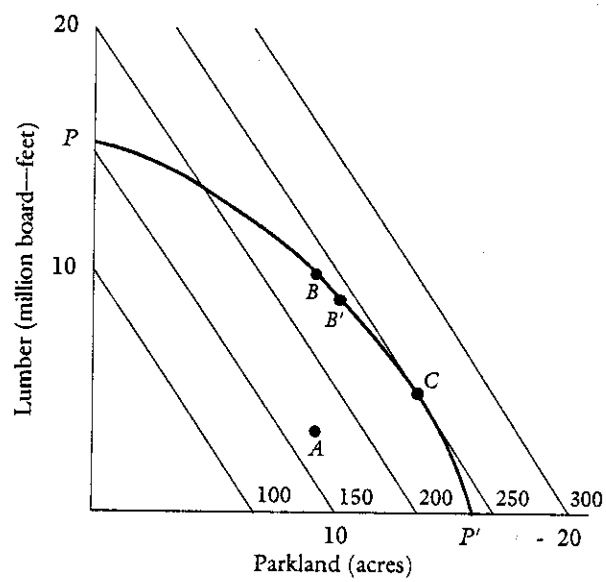

FIGURE 4. The production-possibility frontier.

ous concepts.[10] The two productive-efficiency criteria are intended to fill this need.

The Broad Productivity Criterion

The broad productivity criterion is a more hardheaded analog of Pareto efficiency. Instead of looking for that possible will-o'-the-wisp, utility, it concentrates on measurable outputs of goods and services and, possibly, environmental quality. Otherwise it is similar. Specifically, an economy is said to be productively efficient if it is producing as much of every good and service (rather than utility) as is technically possible, given the outputs of all other goods and services and the amounts of resources used.

As an example, consider an economy that produces only two commodities (there may be any number of consumers). Suppose, to be concrete, that the two commodities are lumber and parkland. The productive possibilities for the community are shown in Figure 4, with acres of parkland plotted horizontally and output of lumber vertically. The key feature of the diagram is the arc *PP′*, called the *production-possibility frontier.* It

[10]The Pareto-efficiency criterion does not depend on having a social-welfare function or on being able to measure individual utilities. All it requires is that we be able to tell for each individual which situation in any pair he prefers. But even this is excessively demanding for most practical purposes.

shows the greatest amount of lumber that can be produced in conjunction with a specified acreage of parkland, and the greatest amount of parkland consistent with a specified output of lumber.[11]

Any combination of lumber and parkland that is on or below the production-possibility frontier is technically feasible (for example, the combinations shown by points *A*, *B*, and *C*), but combinations that lie above the frontier are physically impossible. A productively efficient economy will produce an output right on the frontier (*B* or *C*, not *A*). A badly managed economy could find itself at *A*, perhaps by turning its most productive stands of timber into a park. If those stands were turned over to lumbering and an equal, less heavily wooded area turned into a park, more lumber could be produced without diminishing the amount of parkland, as at efficient point *B*. An important criterion of economic efficiency, therefore, is whether the rules and procedures of an economy tend to lead it to produce at a point like *A* or at one like *B*.

The relationship between productive efficiency (i.e., producing on the production-possibility frontier) and Pareto efficiency is clear, but not simple. On the simple assumption that everyone would like both more lumber and more parkland, an economy cannot be operating Pareto efficiently unless it is productively efficient. For if it were at a point like *A*, it would be possible to change its operations so as to produce more lumber without sacrificing any parkland. There would then be more lumber, and more utility, to distribute; and some people could be made better off without harming anyone. On the other hand, it is quite possible for an economy to be productively efficient without being Pareto efficient. This would happen if the economy were producing everything efficiently but were producing the wrong combination of things, in the sense that the outputs of some commodities could be reduced and other outputs increased in such a way that the new bundle of commodities would satisfy everybody better than the old. In terms of the graphs, an economy could produce at point *B* on the production-possibility frontier in Figure 4 and yield the utilities of a point inside the utility-possibility frontier, such as point *A* of Figure 3. To illustrate, notice that, as the figure is drawn, if the economy were producing at *B*, it could have about 1.5 acres more parkland by giving up 1.5 million board-feet of lumber (thus moving to point *B′*).[12] If there were a thousand families in the community and each would be glad to reduce its consumption of lumber by 1,000 board-feet or more in order to expand the parks by an acre, then it would not be Pareto efficient for the economy to produce at point *B*; everyone's welfare could be increased by moving to *B′*.

In short, productive efficiency is a necessary condition for Pareto effi-

[11] The diagonal lines will be explained below.

[12] The amount of lumber that has to be given up to obtain one more acre of parkland is known as the *marginal rate of transformation* of lumber for parkland.

ciency, but it is not sufficient. Pareto efficiency is the more fundamental and the more demanding achievement. On the other hand, it is much easier to judge the productive efficiency of any policy or economy than to assess its Pareto efficiency.

The criterion that an efficient economy produces an output that lies on its production-possibility frontier enables us to reject an output like point *A* but does not enable us to choose between outputs like *B* and *C*. The final efficiency criterion is addressed to this question.

The Sharp Productivity Criterion: GDP

The relative desirability of any two privately owned commodities is reflected in their prices. On familiar grounds, an additional dollar's worth of anything a consumer buys gives him as much satisfaction as an additional dollar's worth of anything else. Therefore, if an economy revises its operations so as to provide consumers with commodities that they value more highly, in dollars-and-cents terms, it is providing them with more utility (is moving closer to the utility-possibility frontier) even though it must reduce the quantities of some commodities to provide more of others. This remark discloses the social significance of prices as guides to economic activity, as well as suggesting how to use prices to choose among the points on the production-possibility frontier. It suggests that the best point on the frontier is the one at which the value of the goods and services produced is as great as possible, using market prices for the private goods and services. This gives rise to the criterion that the market value of goods and services produced should be as great as possible.

The market value of final goods and services produced in a country is, by definition, the gross domestic product, GDP for short. The criterion just deduced can therefore be expressed by saying that the output of marketable goods and services is optimal when the GDP is as great as possible, using the market prices corresponding to that output to evaluate the various goods and services. But that criterion is inadequate for guiding environmental decisions, or decisions about public goods in general, since it ignores the values of unmarketed public goods and services. For example, in Figure 4 the GDP account would include the value of the lumber produced but would not include the value derived from using the parkland, which is a public good for which no price is charged.

In order to bring environmental values, and the values of other public goods and services, into the picture, a more inclusive concept is needed, which we shall call augmented GDP. Augmented GDP, ADP for short, is defined as the ordinary GDP plus the value of public goods and services of all sorts, including, of course, those provided by the environment. The papers in Section IV explain at some length the issues involved in estimating the values of public goods and services, and the reasons why it is useful to have both the GDP and the AGDP concepts. For the rest of this article, when we refer to gross domestic product we shall always mean the augmented version of the concept unless we state the contrary.

This concept enables us to take due account of the values of environmental and other public goods, merely by choosing the point on the production-possibility frontier at which the AGDP is as great as possible. The use of the AGDP criterion is illustrated in Figure 4. If lumber is worth \$10 per million board-feet and a year's use of parkland is worth \$15 per acre, then

$$\text{AGDP} = \$10 \times \text{millions of board-feet} + \$15 \times \text{acres of park.}$$

The slanting lines in the figure are lines of constant AGDP at the prices assumed. The output at point *C* is worth \$250, and this is greater than the value of output at any other point on or under the production-possibility frontier. According to the AGDP criterion, point *C* should be chosen.

When we invoke the AGDP criterion we not only extend the concept of GDP to include the value of public goods, we also extend the reasoning that justifies its use. The ordinary GDP criterion, remember, was justified by noting that consumers could be better pleased if the output of the economy could be changed so as to replace some commodities for which they would be willing to pay just \$1 by others for which they would be willing to pay more than \$1. But consumers do not have to pay for the use of public goods or the environment, so the AGDP criterion cannot be justified by appeal to consumers' market choices. In fact, however, when the government provides public goods it is acting as an agent for the consumers by buying something for them collectively that they cannot purchase individually. The reasoning, then, is that consumers will be better pleased by the provision of some public goods if (and only if) they value those public goods more highly than the marketable goods that have to be relinquished in order to obtain them. The best mix of public and market goods is then the one with the highest possible AGDP, where private goods are valued at market prices and public goods at what the public would be willing to pay for them if it had to.

No point below the production-possibility frontier can satisfy the AGDP criterion, since if the economy were at such a point AGDP could be increased by increasing the output of some commodity without cutting back any other output. Thus an output point that satisfies the AGDP criterion automatically satisfies the broad productivity criterion. It also satisfies the broad utility criterion—i.e., is Pareto efficient, because it leaves no room for improving anyone's consumption bundle without detracting from someone else's.

The AGDP criterion is the most practicable of the four, in fact the only one that can be applied with much assurance. We have followed a long road to arrive at it, and it is time to reconsider what we have learned. The ideal would be the sharp utility criterion, which maximizes the society's social-welfare function. But a society may not have a social-welfare function; in any event we do not know it, so that is impracticable. Next best is the broad utility criterion, but that one is indecisive among many options and, besides, depends upon unascertainable data for the most part. Next we considered the broad productivity criterion, the demand that the bundle of outputs produced by the economy should lie on the production-possibility frontier. This is a modest requirement, and one that can often be implemented, but it is insufficient as a principle of social choice because it does not instruct us which point on the production-possibility frontier is to be preferred. Finally, we arrived at the AGDP criterion, which does select a decision that satisfies both the broad utility criterion and the broad productivity criterion. These, indeed, are its main justifications.

For these reasons, the AGDP criterion is the main reliance of practical economic analysis in the environmental field. Perfection cannot be claimed for it. In particular, it pays no attention to the equity or inequity of the distribution of income or of environmental benefits. It is therefore quite possible for a society to prefer a decision that is inconsistent with this criterion to one that is consistent and, by the same token, for a proposal that will increase AGDP to be socially undesirable.