CHAPTER 34

EXTERNALITIES

We say that an economic situation involves a **consumption externality** if one consumer cares directly about another agent's production or consumption. For example, I have definite preferences about my neighbor playing loud music at 3 in the morning, or the person next to me in a restaurant smoking a cheap cigar, or the amount of pollution produced by local auto- ... consumption externalities. On

sibilities of one firm are influenced by the choices of another firm or consumer. A classic example is that of an apple orchard located next to a beekeeper, where there are mutual positive production externalities—each firm's production positively affects the production possibilities of the other firm. Similarly, a fishery cares about the amount of pollutants dumped into its fishing area, since this will negatively influence its catch.

The crucial feature of externalities is that there are goods people care about that are not sold on markets. There is no market for loud music at 3 in the morning, or drifting smoke from cheap cigars, or a neighbor who keeps a beautiful flower garden. It is this lack of markets for externalities that causes problems.

Up until now we have implicitly assumed that each agent could make consumption or production decisions without worrying about what other agents were doing. All interactions between consumers and producers took place via the market, so that all the economic agents needed to know were the market prices and their own consumption or production possibilities. In this chapter we will relax this assumption and examine the economic consequences of externalities.

In earlier chapters we saw that the market mechanism was capable of achieving Pareto efficient allocations when externalities were *not* present. If externalities are present, the market will not necessarily result in a Pareto efficient provision of resources. However, there are other social institutions such as the legal system, or government intervention, that can "mimic" the market mechanism to some degree and thereby achieve Pareto efficiency. In this chapter we'll see how these institutions work.

34.1 Smokers and Nonsmokers

It is convenient to start with an example to illustrate some of the main considerations. We'll imagine two roommates, A and B, who have preferences over "money" and "smoke." We suppose that both consumers like money, but that A likes to smoke and B likes clean air.

We can depict the consumption possibilities for the two consumers in an Edgeworth box. The length of the horizontal axis will represent the total amount of money the two agents have, and the height of the vertical axis will represent the total amount of smoke that can be generated. The preferences of agent A are increasing in both money and smoke, while agent B's preferences are increasing in money and clean air—the absence of smoke. We'll measure smoke on a scale from 0 to 1, where 0 is no smoke at all, and 1 is the proverbial smoke-filled room.

This setup gives us a diagram like that depicted in Figure 34.1. Note that the picture looks very much like the standard Edgeworth box, but the interpretation is quite different. The amount of smoke is a good for A and a bad for B, so that B is moved to a more preferred position as A consumes less smoke. Be sure to note the difference in the way things are measured on the horizontal and vertical axes. We measure A's money horizontally from the lower left-hand corner of the box, and B's money horizontally from the upper right-hand corner. But the total amount of smoke is measured vertically from the lower left-hand corner. The difference occurs because money can be divided between the two consumers, so there will always be two amounts of money to measure, but there is only one amount of smoke that they must both consume.

In the ordinary Edgeworth box diagram B is made better off when A reduces his consumption of good 2—but that is because B then gets to consume more of good 2. In the Edgeworth box in Figure 34.1 B is also better off when A reduces his consumption of good 2 (smoke), but for a very different reason. In this example, B is better off when A reduces his consumption of smoke since both agents must consume the same amount of smoke and smoke is a bad for agent B.

We've now illustrated the consumption possibilities of the two roommates and their preferences. What about their endowments? Let's assume that they both have the same amount of money, say $100 apiece, so that their endowments will lie somewhere on the vertical line in Figure 34.1. In order to determine exactly where on this line the endowments lie, we must determine the initial "endowment" of smoke/clean air.

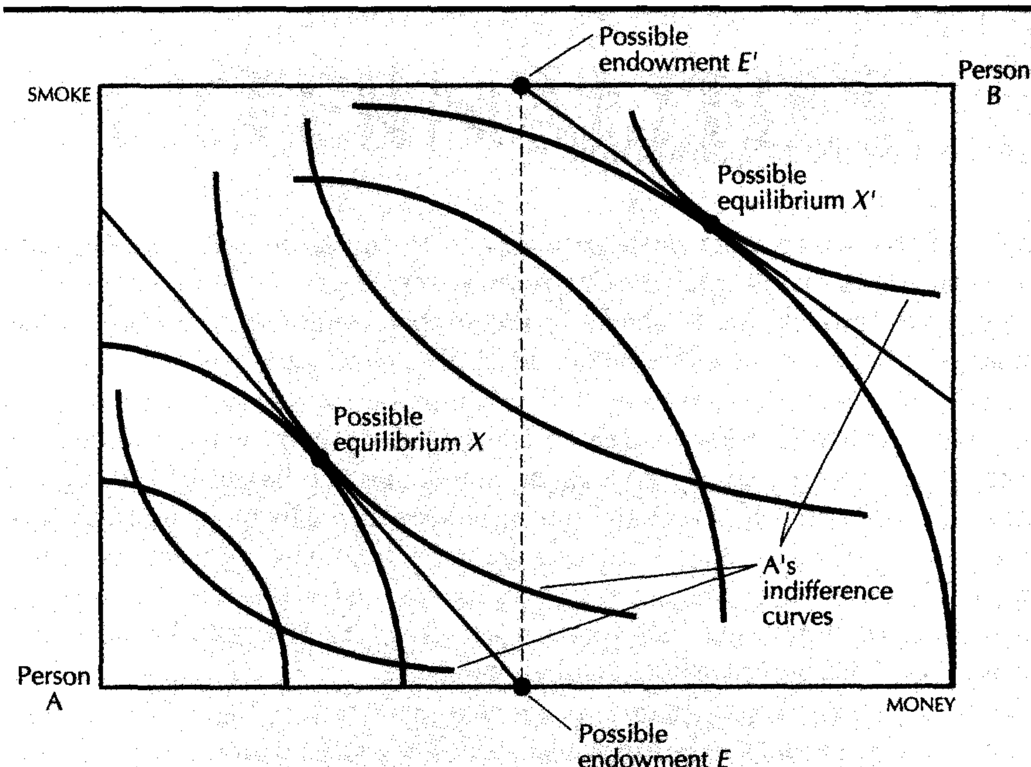

Preferences for money and smoke. Smoke is a good for person A but a bad for person B. Which equilibrium we end up at depends on which endowment we start at.

The answer to this question depends on the legal rights of smokers and nonsmokers. It may be that A has a right to smoke as much as he wants, and B just has to put up with it. Or, it could be that B has a right to clean air. Or the legal right to smoke and clean air could be somewhere between these two extremes.

The initial endowment of smoke depends on the legal system. This is not so different from the initial endowment of ordinary sorts of goods. To say that A has an initial endowment of \$100 means that A can decide to consume the \$100 himself, or he can give it away or trade it to any other individual. There is a legal definition of property involved in saying that a person "owns" or "has a right to" \$100. Similarly if a person has a property right to clean air, it means that he can consume clean air if he wants to, or he can give it away or sell that right to someone else. In this way, having a property right to clean air is no different from having a property right to \$100.

Let's start by considering a legal situation where person B has a legal right to clean air. Then the initial endowment in Figure 34.1 is labeled E; it is where A has $(100, 0)$ and B has $(100, 0)$. This means that both A and B have \$100, and that the initial endowment—what there would be in the absence of trade—is clean air.

Just as before, in the case with no externalities, there is no reason why the initial endowment is Pareto efficient. One of the aspects of having a property right to clean air is having the right to trade some of it away for other desirable goods—in this case, for money. It can easily happen that B would prefer to trade some of his right to clean air for some more money. The point labeled X in Figure 34.1 is an example of such a case.

As before, a Pareto efficient allocation is one where neither consumer can be made better off without the other being made worse off. Such an allocation will be characterized by the usual tangency condition that the marginal rates of substitution between smoke and money should be the same between the two agents, as illustrated in Figure 34.1. It is easy to imagine A and B trading to such a Pareto efficient point. In effect, B has the right to clean air, but he can allow himself to be "bribed" to consume some of A's smoke.

Of course, other assignments of property rights are possible. We could imagine a legal system where A had a right to smoke as much as he wanted, and B would have to bribe A to reduce his consumption of smoke. This would correspond to the endowment labeled E' in Figure 34.1. Just as before, this would typically not be Pareto efficient, so we could imagine the agents trading to a mutually preferred point such as the one labeled X'.

Both X and X' are Pareto efficient allocations; they just come from different initial endowments. Certainly the smoker, A, is better off at X' than at X, and the nonsmoker, B, is better off at X than at X'. The two points have different distributional consequences, but on grounds of efficiency they are equally satisfactory.

In fact, there is no reason to limit ourselves to just these two efficient points. As usual there will be a whole contract curve of Pareto efficient allocations of smoke and money. If agents are free to trade both of these goods, we know that they will end up somewhere on this contract curve. The exact position will depend on their property rights involving smoke and money and on the precise mechanism that they use to trade.

One mechanism that they could use to trade is the price mechanism. Just as before we could imagine an auctioneer calling out prices and asking how much each agent would be willing to buy at those prices. If the initial endowment point gave A the property rights to smoke, he could consider selling some of his smoking rights to B in exchange for B's money. Similarly, if the property rights for clean air were given to B, he could sell some of his clean air to A.

When the auctioneer manages to find a set of prices where supply equals demand everything is fine: we have a nice Pareto efficient outcome. If there is a market for smoke, a competitive equilibrium will be Pareto efficient. Furthermore, the competitive prices will measure the marginal rate of substitution between the two goods, just as in the standard case.

This is just like the usual Edgeworth box analysis, but described in a slightly different framework. As long as we have well-defined property rights in the good involving the externality—no matter who holds the property rights—the agents can trade from their initial endowment to a Pareto efficient allocation. If we want to set up a market in the externality to encourage trade, that will work as well.

The only problem arises if the property rights are *not* well defined. If A believes that he has the right to smoke and B believes that he has the right to clean air, we have difficulties. *The practical problems with externalities generally arise because of poorly defined property rights.*

My neighbor may believe that he has the right to play his trumpet at 3 in the morning, and I may believe that I have the right to silence. A firm may believe that it has the right to dump pollutants into the atmosphere that I breathe, while I may believe that it doesn't. Cases where property rights are poorly defined can lead to an inefficient production of externalities—which means that there would be ways to make both parties involved better off by changing the production of externalities. If property rights are well defined, and mechanisms are in place to allow for negotiation between people, then people can trade their rights to produce externalities in the same way that they trade rights to produce and consume ordinary goods.

34.2 Quasilinear Preferences and the Coase Theorem

We argued above that as long as property rights were well defined, trade between agents would result in an efficient allocation of the externality. In general, the amount of the externality that will be generated in the efficient solution will depend on the assignment of property rights. In the case of the two roommates, the amount of smoke generated will depend on whether the smoker has the property rights or the nonsmoker has them.

But there is a special case where the outcome of the externality is independent of the assignment of property rights. If the agents' preferences are **quasilinear**, then every efficient solution must have the same amount of the externality.

This case is illustrated in Figure 34.2 for the Edgeworth box case of the smoker versus the nonsmoker. Since the indifference curves are all horizontal translates of each other, the locus of mutual tangencies—the set of Pareto efficient allocations—will be a horizontal line. This means that the amount of smoke is the same in every Pareto efficient allocation; only the dollar amounts held by the agents differ across the efficient allocations.

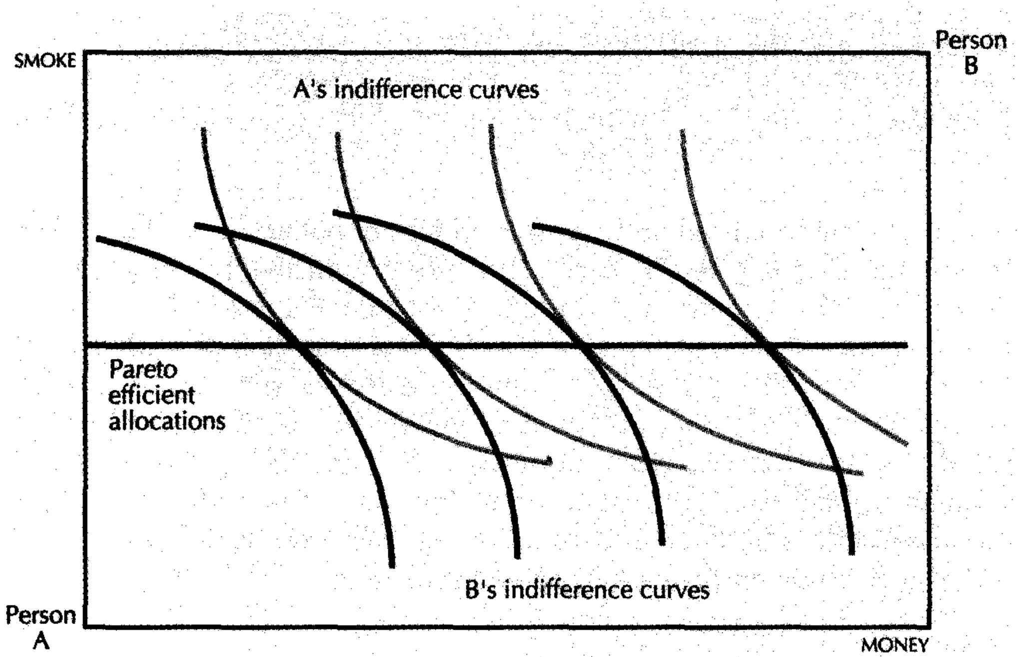

Quasilinear preferences and the Coase theorem. If each consumer's preferences are quasilinear, so that they are all horizontal translates of each other, the set of Pareto efficient allocations will be a horizontal line. Thus there will be a unique amount of the externality, in this case smoke, at each Pareto efficient allocation.

The result that under certain conditions the efficient amount of the good involved in the externality is independent of the distribution of property rights is sometimes known as the **Coase Theorem**. However, it should be emphasized just how special these conditions are. The quasilinear preference assumption implies that the demands for the good causing the externality doesn't depend on the distribution of income. Therefore a reallocation of endowments doesn't affect the efficient amount of the externalities. This is sometimes expressed by saying that the Coase theorem is valid if there are no "income effects."[1]

In this case, the Pareto efficient allocations will involve a unique amount of the externality being generated. The different Pareto efficient allocations will involve different amounts of money being held by the consumers; but the amount of the externality—the amount of smoke—will be independent of the distribution of wealth.

34.3 Production Externalities

Let us now consider a situation involving production externalities. Firm S produces some amount of steel, s, and also produces a certain amount of pollution, x, which it dumps into a river. Firm F, a fishery, is located downstream and is adversely affected by S's pollution.

Suppose that firm S's cost function is given by $c_s(s,x)$, where s is the amount of steel produced and x is the amount of pollution produced. Firm F's cost function is given by $c_f(f,x)$, where f indicates the production of fish and x is the amount of pollution. Note that F's costs of producing a given amount of fish depend on the amount of pollution produced by the steel firm. We will suppose that pollution increases the cost of providing fish $\Delta c_f/\Delta x > 0$, and that pollution *decreases* the cost of steel production, $\Delta c_s/\Delta x \leq 0$. This last assumption says that increasing the amount of pollution will decrease the cost of producing steel—that reducing pollution will increase the cost of steel production, at least over some range.

The steel firm's profit-maximization problem is

$$\max_{s,x}\ p_s s - c_s(s,x)$$

and the fishery's profit-maximization problem is

$$\max_{f}\ p_f f - c_f(f,x).$$

Note that the steel mill gets to choose the amount of pollution that it generates, but the fishery must take the level of pollution as outside of its control.

[1] Ronald Coase is an emeritus professor at the University of Chicago Law School. His famous paper, "The Problem of Social Costs," *The Journal of Law & Economics*, 3 (October 1960), has been given a variety of interpretations. Some authors suggest that Coase only asserted that costless bargaining over externalities achieves a Pareto efficient outcome, not that the outcome will be independent of the assignment of property rights. Coase received the 1991 Nobel Prize in Economics for this work.

The conditions characterizing profit maximization will be

$$p_s = \frac{\Delta c_s(s^*, x^*)}{\Delta s}$$
$$0 = \frac{\Delta c_s(s^*, x^*)}{\Delta x}$$

for the steel firm and

$$p_f = \frac{\Delta c_f(f^*, x^*)}{\Delta f}$$

for the fishery. These conditions say that at the profit-maximizing point, the price of each good—steel and pollution—should equal its marginal cost. In the case of the steel firm, one of its products is pollution, which, by assumption, has a zero price. So the condition determining the profit-maximizing supply of pollution says to produce pollution until the cost of an extra unit is zero.

It is not hard to see the externality here: the fishery cares about the production of pollution but has no control over it. The steel firm looks only at the cost of producing steel when it makes its profit-maximizing calculation; it doesn't consider the cost it imposes on the fishery. The increase in the cost of fishing associated with an increase in pollution is part of the **social cost** of steel production, and it is being ignored by the steel firm. In general, we expect that the steel firm will produce too much pollution from a social point of view since it ignores the impact of that pollution on the fishery.

What does a Pareto efficient production plan for steel and fish look like? There is an easy way to see what it should be. Suppose that the fishery and the steel firm merged and formed one firm that produced both fish and steel (and possibly pollution). Then there is no externality! For a production externality only arises when one firm's actions affect another firm's production possibilities. If there is only one firm, then it will take the interactions between its different "divisions" into account when it chooses the profit-maximizing production plan. We say that the externality has been **internalized** by this reassignment of property rights. Before the merger, each firm had the right to produce whatever amount of steel or fish or pollution that it wanted, regardless of what the other firm did. After the merger, the combined firm has the right to control the production of both the steel mill and the fishery.

The merged firm's profit-maximization problem is

$$\max_{s,f,x} \; p_s s + p_f f - c_s(s,x) - c_f(f,x),$$

which yields optimality conditions of

$$p_s = \frac{\Delta c_s(\hat{s}, \hat{x})}{\Delta s}$$

$$p_f = \frac{\Delta c_f(\hat{f}, \hat{x})}{\Delta f}$$

$$0 = \frac{\Delta c_s(\hat{s}, \hat{x})}{\Delta x} + \frac{\Delta c_f(\hat{f}, \hat{x})}{\Delta x}.$$

The crucial term is the last one. This shows that the merged firm will take into account the effect of pollution on the marginal costs of both the steel firm and the fishery. When the steel division decides how much pollution to produce, it considers the effect of this action on the profits of the fish division; that is, it takes the social cost of its production plan into account.

What does this imply about the amount of pollution produced? When the steel firm acted independently, the amount of pollution was determined by the condition

$$\frac{\Delta c_s(s^*, x^*)}{\Delta x} = 0. \tag{34.1}$$

That is, the steel mill produced pollution until the marginal cost was zero:

$$MC_S(s^*, x^*) = 0.$$

In the merged firm, the amount of pollution is determined by the condition

$$\frac{\Delta c_s(\hat{s}, \hat{x})}{\Delta x} + \frac{\Delta c_f(\hat{f}, \hat{x})}{\Delta x} = 0. \tag{34.2}$$

That is, the merged firm produces pollution until the *sum* of the marginal cost to the steel mill and the marginal cost to the fishery is zero. This condition can also be written as

$$-\frac{\Delta c_s(\hat{s}, \hat{x})}{\Delta x} = \frac{\Delta c_f(\hat{f}, \hat{x})}{\Delta x} > 0 \tag{34.3}$$

or

$$-MC_S(\hat{s}, \hat{x}) = MC_F(\hat{f}, \hat{x}).$$

In this latter expression $MC_F(\hat{f}, \hat{x})$ is positive, since more pollution increases the cost of producing a given amount of fish. Hence the merged firm will want to produce where $-MC_S(\hat{s}, \hat{x})$ is positive; that is, it will want to produce *less* pollution than the independent steel firm. When the true social cost of the externality involved in the steel production is taken into account, the optimal production of pollution will be reduced.

When the steel firm considers minimizing its **private costs** of producing steel, it produces where the marginal cost of extra pollution equals zero;

but the Pareto efficient level of pollution requires minimizing the **social costs** of the pollution. At the Pareto efficient level of pollution, the *sum* of the two firm's marginal costs of pollution must be equal to zero.

This argument is illustrated in Figure 34.3. In this diagram $-MC_S$ measures the marginal cost to the steel firm from producing more pollution. The curve labeled MC_F measures the marginal cost to the fishery of more pollution. The profit-maximizing steel firm produces pollution up to the point where its marginal cost from generating more pollution equals zero.

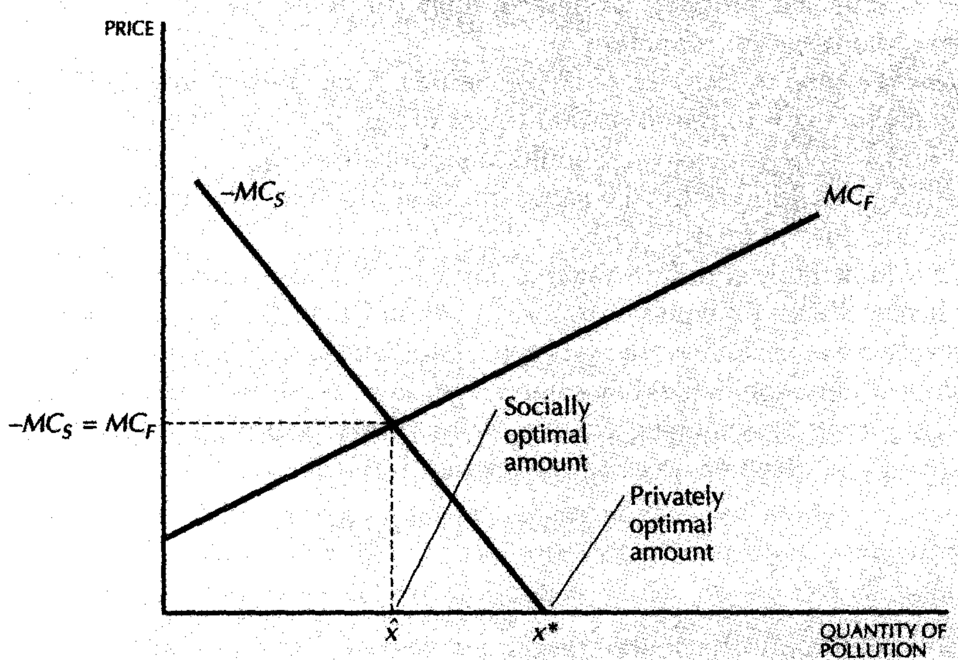

Social cost and private cost. The steel firm produces pollution up to the point where the marginal cost of extra pollution equals zero. But the Pareto efficient production of pollution is at the point where price equals marginal social cost, which includes the cost of pollution borne by the fishery.

But at the Pareto efficient level of pollution, the steel firm pollutes up to the point where the effect of a marginal increase in pollution is equal to the marginal social cost, which counts the impact of pollution on the costs of both firms. At the efficient level of pollution production, the amount that the steel firm is willing to pay for an extra unit of pollution should equal the social costs generated by that extra pollution—which include the costs it imposes on the fishery.

This is perfectly consistent with the efficiency arguments given in earlier chapters. There we assumed that there were no externalities, so that private costs and social costs coincided. In this case the free market will determine a Pareto efficient amount of output of each good. But if the private costs and the social costs diverge, the market alone may not be sufficient to achieve Pareto efficiency.

EXAMPLE: Pollution Vouchers

Everyone wants a clean environment ... as long as someone else pays for it. Even if we reach a consensus on how much we should reduce pollution, there is still the problem of determining the most cost-effective way to achieve the targeted reduction.

Take the case of nitrogen oxide emissions. One emitter may find it relatively inexpensive to reduce its emissions of this pollutant, whereas another may find it very expensive. Should they both be required to reduce their emission of pollutants by the same physical amount, by the same proportional amount, or by some other rule?

Let's look at a simple economic model. Suppose that there are only two firms. Firm 1's emission quota is x_1 and firm 2's is x_2. The cost of achieving an emission quota x_1 is $c_1(x_1)$ and similarly for firm 2. The total amount of emission is fixed at some target level X. If we want to minimize the total costs of achieving the emissions target, subject to the aggregate constraint, we need to solve the following problem:

$$\min_{x_1, x_2} c_1(x_1) + c_2(x_2)$$
$$\text{such that } x_1 + x_2 = X.$$

A by now standard economic argument shows that the marginal cost of emission control must be equalized across the firms. If one firm had a higher marginal cost of emission control than the other, then we could lower total costs by reducing its quota and increasing the quota of the other firm.

How can we achieve this outcome? If the government regulators had information on the cost of emissions for all firms, they could calculate the appropriate pattern of production and impose it on all the relevant parties. But the cost of gathering all this information, and keeping it up-to-date, is staggering. It is much easier to characterize the optimal solution than to actually implement it!

Many economists have argued that the best way to implement the efficient solution to the emission control problem is to use a market. It appears that such a market based emissions control system will soon be put into effect in Southern California. Here is how the California plan works.[2]

[2] See Richard Stevenson, "Trying a Market Approach to Smog," *New York Times*, March 25, 1992, C1.

Each of the 2700 largest polluters in Southern California is assigned a quota for their emissions of nitrogen oxide. This quota is initially set to be 8 percent less than their previous year's emission. If the firm exactly meets its emissions quota it faces no fines or penalties. However, if it reduces its emissions by *more* than its emissions quota, it can sell the extra "right to emit" on the open market.

Suppose that a firm's quota is 95 tons of nitrogen oxide emissions per year. If it manages to produce only 90 tons in a given year, then it can sell the right to emit 5 tons of nitrogen oxide to some other firm. Each firm can compare the market price of an emission credit to the cost of reducing its emissions and decide whether it was more cost-effective to reduce emissions further or purchase emission credits from other firms.

Firms that find it easy to reduce emissions will sell credits to firms that find it costly to reduce emissions. In equilibrium, the market price of the right to emit one ton of pollution should just equal the marginal cost of reducing emissions by one ton. But this is exactly the condition characterizing the optimal pattern of emissions! The market for emission permits produces the efficient pattern of emissions automatically.

34.4 Interpretation of the Conditions

There are several useful interpretations of the conditions for Pareto efficiency derived above. Each of these interpretations suggests a scheme to correct the efficiency loss created by the production externality.

The first interpretation is that the steel firm faces the wrong price for pollution. As far as the steel firm is concerned, its production of pollution costs it nothing. But that neglects the costs that the pollution imposes on the fishery. According to this view, the situation can be rectified by making sure that the polluter faces the correct social cost of its actions.

One way to do this is to place a tax on the pollution generated by the steel firm. Suppose that we put a tax of t dollars per unit of pollution generated by the steel firm. Then the profit-maximization problem of the steel firm becomes

$$\max_{s,x} \; p_s s - c_s(s,x) - tx.$$

The profit-maximization conditions for this problem will be

$$p_s - \frac{\Delta c_s(s,x)}{\Delta s} = 0$$
$$-\frac{\Delta c_s(s,x)}{\Delta x} - t = 0.$$

Comparing these conditions to equation (34.3), we see that setting

$$t = \frac{\Delta c_f(\hat{f},\hat{x})}{\Delta x}$$

will make these conditions the same as the conditions characterizing the Pareto efficient level of pollution.

This kind of a tax is known as a **Pigouvian tax.**[3] The problem with Pigouvian taxes is that we need to know the optimal level of pollution in order to impose the tax. But if we knew the optimal level of pollution we could just tell the steel firm to produce exactly that much and not have to mess with this taxation scheme at all.

Another interpretation of the problem is that there is a missing market—the market for the pollutant. The externality problem arises because the polluter faces a zero price for an output good that it produces, even though people would be willing to pay money to have that output level reduced. From a social point of view, the output of pollution should have a *negative* price.

We could imagine a world where the fishery had the right to clean water, but could sell the right to allow pollution. Let q be the price per unit of pollution, and let x be the amount of pollution that the steel mill produces. Then the steel mill's profit-maximization problem is

$$\max_{s,x} \; p_s s - qx - c_s(s,x),$$

and the fishery's profit-maximization problem is

$$\max_{f,x} \; p_f f + qx - c_f(f,x).$$

The term qx enters with a negative sign in the profit expression for the steel firm since it represents a cost—the steel firm must buy the right to generate x units of pollution. But it enters with a positive sign in the expression for the profits of the fishery, since the fishery gets revenue from selling this right.

The profit-maximization conditions are

$$p_s = \frac{\Delta c_s(s,x)}{\Delta s} \tag{34.4}$$

$$q = -\frac{\Delta c_s(s,x)}{\Delta x} \tag{34.5}$$

$$p_f = \frac{\Delta c_f(f,x)}{\Delta f} \tag{34.6}$$

$$q = \frac{\Delta c_f(f,x)}{\Delta x}. \tag{34.7}$$

[3] Arthur Pigou (1877–1959), an economist at Cambridge University, suggested such taxes in his influential book *The Economics of Welfare.*

Thus each firm is facing the social marginal cost of each of its actions when it chooses how much pollution to buy or sell. If the price of pollution is adjusted until the demand for pollution equals the supply of pollution, we will have an efficient equilibrium, just as with any other good.

Note that at the optimal solution, equations (34.5) and (34.7) imply that

$$-\frac{\Delta c_s(s,x)}{\Delta x} = \frac{\Delta c_f(f,x)}{\Delta x}.$$

This says that the marginal cost to the steel firm of reducing pollution should equal the marginal benefit to the fishery of that pollution reduction. If this condition were not satisfied, we couldn't have the optimal level of pollution. This is, of course, the same condition we encountered in equation (34.3).

In analyzing this problem we have stated that the fishery had a right to clean water and that the steel mill had to purchase the right to pollute. But we could have assigned the property rights in the opposite way: the steel mill could have the right to pollute and the fishery would have to pay to induce the steel mill to pollute less. Just as in the case of the smoker and nonsmoker, this would also give an efficient outcome. In fact, it would give precisely the *same* outcome, since exactly the same equations would have to be satisfied.

To see this, we now suppose that the steel mill has the right to pollute up to some amount $\overline{x}$, say, but the fishery is willing to pay it to reduce its pollution. The profit-maximization problem for the steel mill is then

$$\max_{s,x}\ p_s s + q(\overline{x} - x) - c_s(s,x).$$

Now the steel mill has two sources of income: it can sell steel, and it can sell pollution relief. The price equals marginal cost conditions become

$$p_s - \frac{\Delta c_s(s,x)}{\Delta s} = 0 \qquad (34.8)$$

$$-q - \frac{\Delta c_s(s,x)}{\Delta x} = 0. \qquad (34.9)$$

The fishery's maximization problem is now

$$\max_{f,x}\ p_f f - q(\overline{x} - x) - c_f(f,x),$$

which has optimality conditions

$$p_f - \frac{\Delta c_f(f,x)}{\Delta f} = 0 \qquad (34.10)$$

$$q - \frac{\Delta c_f(f,x)}{\Delta x} = 0. \qquad (34.11)$$

Now observe: the four equations (34.8)–(34.11) are precisely the same as the four equations (34.4)–(34.7). In the case of production externalities, the optimal pattern of production is independent of the assignment of property rights. Of course, the distribution of profits will generally depend on the assignment of property rights. Even though the social outcome will be independent of the distribution of property rights, the owners of the firms in question may have strong views about what is an appropriate distribution.

34.5 Market Signals

Finally we turn to the third interpretation of externalities, which in some respects is the most profound. In the case of the steel mill and the fishery there is no problem if both firms merge—so why don't they merge? In fact, when you think about it, there is a definite incentive for the two firms to merge: if the actions of one affect the other, then they can make higher profits together by coordinating their behavior than by each going alone. *The objective of profit maximization itself should encourage the internalization of production externalities.*

Said another way: if the joint profits of the firms with coordination exceed the sum of the profits without coordination, then the current owners could each be bought out for an amount equal to the present value of the stream of profits for their firm, the two firms could be coordinated, and the buyer could retain the excess profits. The new buyer could be either of the old firms, or anybody else for that matter.

The market itself provides a signal to internalize production externalities, which is one reason this kind of production externality is rarely observed. Most firms have *already* internalized the externalities between units that affect each other's production. The case of the apple orchard and the beekeeper mentioned earlier is a case in point. Here there *would* be an externality if the two firms ignored their interaction ... but why would they be so foolish as to do so? It is more likely that one or both of the firms would realize that more profits could be made by coordinating their activities, either by mutual agreement or by the sale of one of the firms to the other. Indeed, it is very common for apple orchards to keep honey bees for the purpose of fertilizing the trees. That particular externality is easily internalized.

EXAMPLE: Bees and Almonds

Many varieties of fruit and nut trees need bees to pollinate their blossoms, thereby allowing the trees to produce crops.

According to the Carl Hayden Bee Research Center in Tucson, Arizona, honeybees pollinate about one-third of the human diet and more than 50 different agricultural crops valued at more than \$20 billion a year in the United States.[4]

Some owners of orchards keep their own bees; some rely on their neighbors' bees or wild bees. However, as the theory of externalities suggests, the most natural solution to the problem of inadequate bee supply is a market for bee services.

Consider, for example, the California almond market. There are 530,000 acres of almond trees in California, and every year, more than 1 million honeybee hives are needed to pollinate the trees. But California only has 440,000 resident bee hives. There aren't enough California bees to pollinate all those almond trees!

The solution is to import bees from other nearby states. There is, in fact, a ready market for such services, with beekeepers bringing hives from North Dakota, Washington, and Colorado to supplement the native California bees. The almond growers pay well for these services: in 2004, bee pollination services sold for \$54 per hive.

34.6 The Tragedy of the Commons

We have argued above that if property rights are well defined, there will be no problem with production externalities. But if property rights are not well defined, the outcome of the economic interactions will undoubtedly involve inefficiencies.

In this section we will examine a particularly well-known inefficiency called "the tragedy of the commons."[5] We will pose this problem in the original context of a common grazing land, although there are many other possible illustrations.

Consider an agricultural village in which the villagers graze their cows on a common field. We want to compare two allocation mechanisms: the first is the private ownership solution where someone owns the field and decides how many cows should graze there; the second is the solution where the field is owned in common by the villagers and access to it is free and unrestricted.

Suppose that it costs a dollars to buy a cow. How much milk the cow produces will depend on how many other cows are grazed on the common land. We'll let $f(c)$ be the value of the milk produced if there are c cows grazed on the common. Thus the value of the milk per cow is just the average product, $f(c)/c$.

[4] Anna Oberthur, "Almond Growers Face Need for Bees," Associated Press, February 29, 2004.

[5] See G. Hardin, "The Tragedy of the Commons," *Science*, 1968, 1243–47.

How many cows would be grazed on the common if we wanted to maximize the total wealth of the village? In order to maximize the total amount of wealth, we set up the following problem:

$$\max_{c} \; f(c) - ac.$$

It should be clear by now that the maximal production will occur when the marginal product of a cow equals its cost, a:

$$MP(c^*) = a.$$

If the marginal product of a cow were greater than a, it would pay to put another cow on the commons; and if it were less than a, it would pay to take one off.

If the common grazing ground were owned by someone who could restrict access to it, this is indeed the solution that would result. For in this case, the owner of the grazing grounds would purchase just the right amount of cows to maximize his profits.

Now what would happen if the individual villagers decided whether or not to use the common field? Each villager has a choice of grazing a cow or not grazing one, and it will be profitable to graze a cow as long as the output generated by the cow is greater than the cost of a cow. Suppose that there are c cows currently being grazed, so that the current output per cow is $f(c)/c$. When a villager contemplates adding a cow, the total output will be $f(c+1)$, and the total number of cows will be $c+1$. Thus the revenue that the cow generates for the villager will be $f(c+1)/(c+1)$. He must compare this revenue to the cost of the cow, a. If $f(c+1)/(c+1) > a$, it is profitable to add the cow since the value of the output exceeds the cost. Hence the villagers will choose to graze cows until the average product of a cow is driven to a. It follows that the total number of cows grazed will be $\hat{c}$, where

$$\frac{f(\hat{c})}{\hat{c}} = a.$$

Another way to derive this result is to appeal to free entry. If it is profitable to graze a cow on the common field, villagers will purchase cows. They will stop adding cows to the common only when the profits have been driven to zero, that is, when

$$f(\hat{c}) - a\hat{c} = 0,$$

which is just a rearrangement of the condition in the last paragraph.

When an individual decides whether or not to purchase a cow, he looks at the extra value he will get $f(c)/c$ and compares this to the cost of the cow, a. This is fine for him, but what has been left out of this calculation is the fact that his extra cow will reduce the output of milk from all the *other* cows. Since he is ignoring this **social cost** of his purchase, too many cows will be grazed on the common ground. (We assume that each individual has a number of cows that is negligible relative to the total number grazed on the common.)

This argument is illustrated in Figure 34.4. Here we have depicted a falling average product curve, since it is reasonable to suppose that the output per cow declines as more and more cows are grazed on the common land.

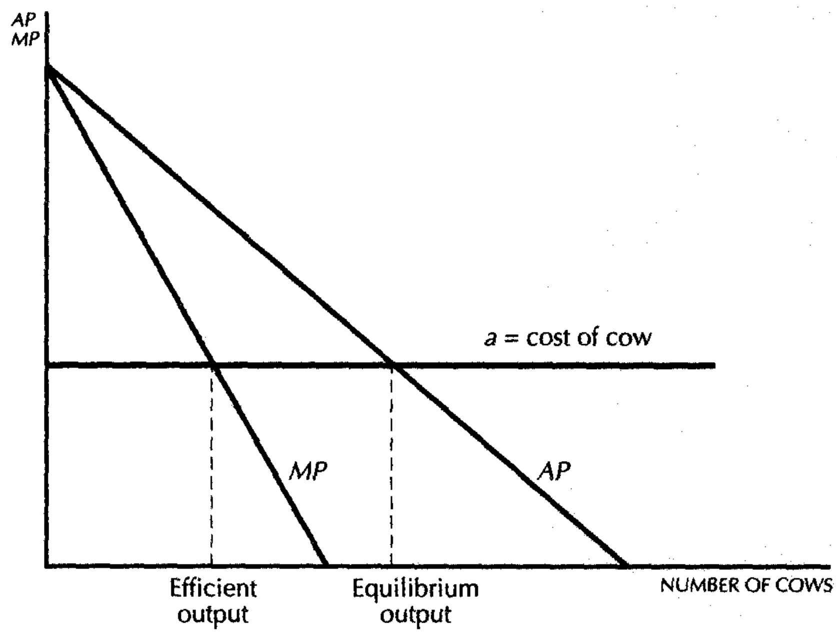

The tragedy of the commons. If the grazing area is privately owned, the number of cows will be chosen so that the marginal product of a cow equals its cost. But if grazing area is common property, cows will be grazed until the profits are driven to zero; thus the area will be overgrazed.

Since the average product is falling, it must be that the marginal product curve always lies below the average product curve. Thus the number of cows where the marginal product equals a must be less than where the average product equals a. The field will be overgrazed in the absence of a mechanism to restrict use.

Private property provides such a mechanism. Indeed, we have seen that if everything that people care about is owned by someone who can control its use and, in particular, can exclude others from overusing it, then there are by definition no externalities. The market solution leads to a Pareto efficient outcome. Inefficiencies can only result from situations where there is no way to exclude others from using something, a topic that we will investigate in the next chapter.

Of course, private property is not the only social institution that can encourage efficient use of resources. For example, rules could be formulated about how many cows can be grazed on the village common. If there is a legal system to enforce those rules, this may be a cost-effective solution to providing an efficient use of the common resource. However, in situations where the law is ambiguous or nonexistent, the tragedy of the commons can easily arise. Overfishing in international waters and the extermination of several species of animals due to overhunting are sobering examples of this phenomenon.

EXAMPLE: Overfishing

According to a report in the *New York Times*, "...overfishing has decimated the stocks of cod, haddock and flounder that have sustained New Englanders for centuries."[6] According to one expert, fishermen in New England are taking 50 to 70 percent of the available stock, over twice the sustainable amount.

This overfishing is a prime example of the problem of the commons: each fisherman has a negligible impact on the total stock of fish, but the accumulated efforts of thousands of fishermen results in serious depletion. The New England Fisheries Management Council is attempting to alleviate the problem by banning new entry to the industry, requiring fishermen to limit their days at sea, and increasing the mesh size of their nets.

It appears that the supplies of fish could be restocked in as little as 5 years if conservation measures were undertaken. The present value of profits to the industry as a whole would be larger with regulation to prevent overfishing. However, such measures would almost certainly imply a substantial reduction in the number of fishing boats in the industry, which is highly unpopular with the small fishermen, who would likely be forced to leave the industry.

EXAMPLE: New England Lobsters

Some fishing industries have already applied stringent rules to avoid overfishing. For example, lobster fishermen work under carefully designed rules to ensure that they do not fish themselves out of a livelihood. For example, they are required to toss back any female lobster bearing eggs, any lobster

[6] "Plenty of Fish in the Sea? Not Anymore," *New York Times*, March 25, 1992, A15.

shorter than a minimum size, and any lobsters bigger than a maximum size.

The "eggers" give birth to more lobsters and the small "tiddlers" grow up to mate. But why throw back the big lobsters? According to marine biologists, large lobsters produce more offspring and larger offspring. If fishermen always took the largest lobsters, the remaining small lobsters would pass their genes onto their progeny, resulting in smaller and smaller lobsters in each generation.

With lobsters there is good news and bad news. First the good news. The 2003 Maine lobster harvest was 5.4 million pounds, more than 2.5 times the 1945–85 average. This suggests that the careful husbandry practiced by the industry has yielded a significant growth in the lobster population.

However, it appears that conservation isn't the only factor. There have also been considerable changes in the population of other species of marine life off the Maine coast, such as sea urchins, and some observers believe that these changes are the primary driver of change in the lobster population.[7]

This leads to the bad news. Further south, in Massachusetts and New York, the lobster catch has fallen dramatically. No one is quite sure why one region is doing so well and the other so poorly. Ironically, Maine may be doing well due to increased harvesting of finned fish and of sea urchins, both of which eat young lobsters. Massachusetts' problems may be due to specific factors, such as a large oil spill and a disfiguring shell disease. Another culprit is warming water: Narragansett Bay temperatures have risen almost two degrees Celsius in the last 20 years.

Ecologies can be very complex and can change rapidly. The efforts to avoid overfishing are to be applauded, but they are only part of the story.

34.7 Automobile Pollution

As suggested above, pollution is a prime example of an economic externality. The activity of one consumer operating an automobile will typically lower the quality of the air that other consumers breathe. It seems unlikely that an unregulated free market would generate the optimal amount of pollution; more likely, if the consumer bears no cost in generating pollution, too much pollution would be produced.

One approach to controlling the amount of automobile pollution is to require that automobiles meet certain standards in the amount of pollution that they generate. This has been the basic thrust of U.S. antipollution policy since the Clean Air Act of 1963. That act, or, more properly, the subsequent amendments, set automobile emission standards for the manufacturers of vehicles in the United States.

7 See *The Economist*, "Claws!" August 19, 2004, and Cornelia Dean, "Lobster Boom and Bust," *New York Times*, August 9, 2004.

Lawrence White has examined the benefits and costs of this program; most of the following discussion is drawn from this work.[8]

White estimates that the cost of emission control equipment is about \$600 per car, the extra maintenance costs are about \$180 per car, and the costs of the reduced gasoline mileage and the necessity for unleaded gasoline come to about \$670 per car. Thus the total cost per car of the emission control standards is about \$1450 over the lifetime of the car. (All figures are in 1981 dollars.)

He argues that there are several problems with the current approach to the regulation of automobile emissions. First, it requires that all automobiles meet the same standards. (California is the only state with different standards for emission control.) This means that *everyone* who buys a car must pay an extra \$1450 whether they live in a high pollution area or not. A 1974 National Academy of Sciences study concluded that 63 percent of all U.S. cars did not require the stringent standards now in effect. According to White, "almost two-thirds of car buyers are spending ... substantial sums for unnecessary systems."

Secondly, most of the responsibility for meeting the standards falls on the manufacturer, and little falls on the user. Owners of cars have little incentive to keep their pollution control equipment in working order unless they live in a state with required inspections.

More significantly, motorists have no incentive to economize on their driving. In cities such as Los Angeles, where pollution is a significant hazard, it makes good economic sense to encourage people to drive less. Under the current system, people who drive 2000 miles a year in North Dakota pay exactly the same amount for pollution control as people who drive 50,000 miles a year in Los Angeles.

An alternative solution to pollution would be *effluent fees.* As described by White, effluent fees would require an annual inspection of all vehicles along with an odometer reading and tests that would estimate the likely emissions of the vehicle during the past year. Different communities could then levy fees based on the estimated amount of pollution that had actually been generated by the operation of the vehicle. This method would ensure that people would face the true cost of generating pollution and would encourage them to choose to generate the socially optimal amount of pollution.

Such a system of effluent fees would encourage the vehicle owners themselves to find low-cost ways of reducing their emissions—investing in pollution control equipment, changing their driving habits, and changing the kinds of vehicles that they operate. A system of effluent fees could impose even higher standards than are now in effect in communities where pollution is a serious problem. Any desired level of pollution control can

[8] See Lawrence White, *The Regulation of Air Pollutant Emissions from Motor Vehicles* (Washington, D.C.: American Enterprise Institute for Public Policy Research, 1982).

be achieved by appropriate effluent fees ... and it can be achieved at a substantially lower cost than the current system of mandated standards.

Of course, there is no reason why there might not also be some federally mandated standards for the two-thirds of the vehicles that are operated in localities where pollution is not a serious problem. If it is cheaper to impose standards than to require inspections, then by all means that should be the proper choice. The appropriate method of pollution control for automobiles should depend on a rational analysis of benefits and costs—as should all social policies of this nature.

Summary

1. The First Theorem of Welfare Economics shows that a free, competitive market will provide an efficient outcome in the absence of externalities.

2. However, if externalities are present, the outcome of a competitive market is unlikely to be Pareto efficient.

3. However, in this case, the state can sometimes "mimic" the role of the market by using prices to provide correct signals about the social cost of individual actions.

4. More importantly, the legal system can ensure that property rights are well defined, so that efficiency-enhancing trades can be made.

5. If preferences are quasilinear, the efficient amount of a consumption externality will be independent of the assignment of property rights.

6. Cures for production externalities include the use of Pigouvian taxes, setting up a market for the externality, simply allowing firms to merge, or transferring property rights in other ways.

7. The tragedy of the commons refers to the tendency for common property to be overused. This is a particularly prevalent form of externality.

REVIEW QUESTIONS

1. True or false? An explicit delineation of property rights usually eliminates the problem of externalities.

2. True or false? The distributional consequences of the delineation of property rights are eliminated when preferences are quasilinear.

3. List some other examples of positive and negative consumption and production externalities.

4. Suppose that the government wants to control the use of the commons, what methods exist for achieving the efficient level of use?