environment. No nation can be excluded from the benefits of the emission reduction efforts of another nation. Biodiversity preservation is another example of a public good. No person can be excluded from the public benefits created from a stable ecosystem created by preserving species. For wildlife viewing, unless the species exist on private reserves, it is difficult to exclude potential beneficiaries from enjoying the prospect of seeing them. If people derive value from the existence of a species, exclusion is impossible. This holds for many environmental resources—if air quality in a city improves, nobody who lives in or visits the city can be excluded from the benefits. Other environmental resources are excludable in consumption. For example, market prices control recreational access to downhill skiing at a resort.

Non-rivalry depends on the characteristics of the good. Non-rivalry means that the benefits gained from the resource are independent of the number of other people who wish to use it. Better air quality, for instance, increases one person's well-being, no matter how many other citizens there are. If an endangered species is protected, the number of people affected does not reduce the benefits to each person. Some environmental resources do not possess this property. If designating a wilderness area as a recreational resource protects it from future development and yet also increases its profile, then the presence of more people implies more congestion and fewer benefits per trip for many visitors. Many people visiting Yellowstone National Park, for example, feel that congestion is a 'bad': too many people, too little solitude.

What is the relevance of these properties for the socially optimal provision of public goods and market failure? Let's revisit the Centennial open space issue as a public goods problem. Open spaces are a public good if defined as an aesthetic available to all people in the same amount—my visual consumption of open space does not reduce your access, nor reduce the amount of open space. The potential problem with using the market to provide this public good voluntarily is *free-riding*—since he or she cannot be excluded from the same amount of the good, each person has an incentive to let someone else provide the public good. Everyone has an incentive to free-ride off the efforts of others.

Free-riding can lead to the well-known case of the *prisoner's dilemma*, also called the *social trap* or the *tragedy of the commons* (for more details, see Box 2.3 and Chapter 7). The dilemma exists when people find that individual incentives lead them to the worst outcome possible—for themselves and society. The idea is compelling: a person looks out for his or her own self-interest and knows that other people are doing the same. Given the incentives, he or she cannot avoid the worst outcome, because the private incentives always push the person towards the non-cooperative outcome. In public good terms, this means that a person knows he or she should not free-ride, because it is better to be a good citizen. Yet the incentives still exist to free-ride. If he does not free-ride but everyone else does, he ends up in the worst position.

In reality, people do contribute to the provision of some public good voluntarily. Whether people contribute to the public good is not the question. The question is whether they voluntarily contribute 'enough'—do people pay enough to create the social optimal level of the public good? If so, a decentralized market approach has succeeded in allocating a public good. If not, by definition the market fails. The market fails because people have undersupplied the social optimal level of the public good, or oversupplied a public 'bad' such as pollution. Free-riding can be partial or complete; the market might fail by a little or a lot. But

BOX 2.3 The Tragedy of the Commons

Biologist Garrett Hardin coined the phrase 'the tragedy of the commons' in the journal *Science* in 1968. The word 'tragedy' meant a 'remorseless working of things'. Hardin was concerned with unchecked global population growth, which he saw as unsustainable for common property. His example of a pasture open to all graziers to let their cattle roam over leads to overgrazing. Each grazier gets the full benefit of adding one more animal to the common (the profit per beast). But as cattle numbers rise, overgrazing occurs, which then causes everyone to lose. Each person rushes headlong down a path that leads to social ruin, destroying the idea of Adam Smith's beneficial invisible hand. This tragedy gets serious, Hardin argues, as population rises, since rising population equates to more people wanting to graze their animals.

Hardin's analysis has been praised, criticized, and extended. We now know that overgrazing of common areas is not new—it happened in ancient Mayan civilization and ancient Egypt. We also know that the price of the good produced on the commons matters more to the rate of exploitation than Hardin thought. We also know that many commons throughout the world are regulated by rules set by those who have access to the area.

Yet Hardin's message is powerful. We can find many examples of open-access resources damaged by overuse. An example is a study of common grazing areas in the Northern Isles off the Scottish coast. On Shetland, common grazings seem to suffer higher levels of environmental damage than comparable privately owned areas.

technically, the market still fails. Market failure occurs with voluntarily public good provision when people contribute any amount less than their true benefits for the good.

Consider the incentives to free-ride on open space provision. Suppose that Ole and the citizens of Centennial are considering pooling their resources and using the market to buy out Riley, to stop him from building his new mansion on the hogback. Since no one can be excluded from the potential benefits of open space, the marginal benefit to society from one more unit of visibility equals the summed marginal benefits from each person who gains enjoyment from the view. Marginal social benefits reflect how society values each extra unit of the public good. The 'optimal level' of the public good is the level that maximizes the net benefits to society. The optimal level is determined by comparing the marginal social benefits to the marginal social cost to provide one more unit of the good. The public good should be provided up to the point at which the marginal social benefits equal the marginal social costs. This optimal level is called the *Lindahl equilibrium*, named after the Swedish economist who first identified it.

If the market fails to generate the optimal level of the public good, how can society attain the optimal level? One possibility is to build a club: everyone who benefits from the good gets together to provide the good, and the benefits of this provision are then restricted to those people belonging to the club. This seems attractive since the beneficiaries of the good pay for it, and government is not forced to intervene. Organizations such as the Royal Society for the Protection of Birds (RSPB) in the United Kingdom and the Sierra Club in the USA, are partial examples of this. The RSPB buys and safeguards nature reserves using funds generated by club members.

But while club members get reduced fees for entry to sites, non-members also benefit from environmental preservation. The non-excludability of benefits means that fewer nature reserves would be provided than optimal if their provision was left to the RSPB.

What is more, club financing is likely to be less than optimal, in that people do not have to join and pay, even though they know they will benefit anyway. The strategic incentive to free-ride still exists.

Common property. Common property, or the *commons*, can be managed in many ways depending on how property rights are defined and enforced. Open-access commons exist when people cannot be excluded from accessing a resource, such that one person's use rivals another's use. If a person's use reduces the total available to all, everyone has an incentive to capture the benefits before someone else gets them. This open-access free-for-all leads to inefficient use of the resource. The moratorium on fishing for cod and witch flounder off the Grand Banks in the North Atlantic is a prime example. By inefficiency, we mean that the fishers harvest to the point at which the marginal costs exceed the marginal revenue (i.e. the market price) of harvesting. Overuse implies that the market price has failed to signal the true scarcity of the asset. Since a fish caught by one fisher is one fewer fish for all others to catch, fishers have no private incentive to account for the scarcity value of the resource. They expend too much effort and end up overharvesting the fish stock relative to what is economically efficient and potentially biologically inefficient. Again, we return to this problem in Chapter 7.

What happens in open-access commons? Each fisher has an incentive to catch as many fish as possible before someone else catches the same fish. He has no incentive to value the scarcity of the fish, because if he does not catch them, someone else will. His decision to let the fish be is not respected by others, because they have as much right to the fish as he does. He starts expending effort and in doing so his effort is such that his marginal costs end up exceeding his marginal revenue. This violates the standard efficiency condition that says that net benefits are maximized when marginal revenue equals marginal costs.

Why? Efficiency is breached because each fisher imposes an external cost on all the other fishers. This external cost arises because each time a fish is caught, it makes it more costly for everyone else to catch the next one. This is like a variant of the infamous 90–10 rule of diminishing returns to effort—you get 90 per cent of the fish with 10 per cent of the effort, but catching the remaining 10 per cent takes 90 per cent of the effort because they are harder to find.

When do the fishers stop their effort in open access? A fisher stops when his marginal costs equal the average revenue, or when his economic profits are zero. In this case, even though net profits to fishers are zero, the net social value is negative, because the scarcity value of the fish has been ignored in the free-for-all. In this case, each fisher does not earn the economic rent that would reflect the scarcity value of the fish to society.

Open access is the classic case of a 'tragedy of the commons'. Since everyone has access, all have the rights to the resource and scarcity value is ignored. But the reality is that most commons have a property-right scheme, either formal or informal, that works to allocate resources in a more efficient manner. There are numerous documented examples of self-governing commons in which people work as a collective unit and respect the scarcity value of the resource. These groups succeed when they establish common-property rights that include sharing rules, exclusion principles, and enforcement and punishment schemes (see Ostrom's classic 1990 book of examples of alternative management schemes for common property).

Market failure need not always occur with the commons. People have defined—and continue to define—rules to capture the scarcity value of a resource shared by many people. Be aware that market failure is associated with the commons in situations in which people use the term 'the commons' or the 'global commons' when they mean a global or regional public good that is both non-rival and non-excludable. Pure public goods are much more of a challenge to handle because of the need for a larger coalition of affected parties. We expand on this in Chapter 9, on the economics of climate change.

Hidden information. Market failure occurs when people cannot observe either the actions or types of other people. *Moral hazard* confounds market operations because one person cannot observe the hidden actions of others. *Adverse selection* frustrates markets because a person cannot observe the hidden type of a person (e.g. whether a farmer has high or low opportunity costs for increasing biodiversity on his or her land) or the hidden quality of some good or service (e.g. whether or not a loaf of bread is organic). Both types of hidden information slow the creation of markets that could be used to allocate resources to more efficient use, such as the reduction of environmental risk.

Moral hazard implies that a regulator cannot perfectly monitor pollution abatement, and a firm could shirk on pollution control. The firm has an incentive to shirk if it bears all the control costs in return for a fraction of the benefits. The result is again too much pollution relative to the efficient level. Also, moral hazard can lead to an inefficient 'pooling' of environmental risk, a topic that we take up in Chapter 5.

Environmental risks are part of life, and it would be better to find markets to allow those who are less willing to bear risk to sell the risk to those who are willing to buy it. But moral hazard reduces the ability to reallocate risk among different economic agents. When the private market cannot monitor actions, an insurer withdraws from the pollution liability market because insurance affects the individual's incentives to take precautions.

Given that accidental spills or storage of pollution can create potential financial liabilities (e.g. clean-up costs, medical expenses), a firm would like to pay to pass these risks on to a less risk-averse agent such as an insurer. But since there is a trade-off between risk-bearing and incentives, the market for pollution liability insurance is incomplete, as insurers reduce the information rents of the better-informed individual. The market produces an inefficient allocation of risk.

Adverse selection affects the environment too. One key recommendation of environmental policy is sustainable production of products. These products should be produced using methods that protect the environment. The challenge facing these sustainable products is that these products are more expensive and demand a higher market price. But if buyers cannot be guaranteed that they are getting their 'sustainability' bang for their buck, they will not buy them. Reporting that a product is sustainable is not enough for many consumers—they need a warranty or label.

If a consumer cannot distinguish the sustainable product from a similar product produced using standard practices, he has no incentive to pay the price premium. Why should he or she pay more than average when he or she cannot distinguish high-quality from low-quality products? The problem is that if the buyers disappear, the sellers have to withdraw from the market. The sellers with an above-average quality product who cannot get a price premium have no incentive to stay in the market, because they could do better investing elsewhere with greater returns. And if enough buyers who are willing to pay a high price disappear, the sellers disappear too, and the market collapses. Preventing this collapse requires a voluntary approach to certification (e.g. the AB scheme in France, or the Soil Association scheme in the UK) or a government-sponsored certification scheme.

2.3 Markets for the Environment

Market failure can create environmental problems. But even if markets fail, we can still use the ideas behind markets to address the problems that might exist. We do not need to turn to command-and-control government intervention or collaborative stakeholder processes. Rather, people can rewrite the rules to create new markets to address the failings of existing markets. They can use market-based policy as a substitute for technocratic or stakeholder processes, which have their own successes and failures (see Box 2.4).

Think about other kinds of real issues—such as financial assets, for instance—in which we are much less willing to delegate decision-making authority to the government or stakeholders. That people have been creating and using markets to manage many assets for the past three centuries signals their power. For example, the relative stakes per percentage risk are larger in financial assets than in risks from environmental dilemmas. But we do not ask the

BOX 2.4 Government Intervention Failure: The Common Agricultural Policy

In this chapter, we are discussing how market failure can lead to undesirable consequences for the environment. Governments have a case for some level of intervention to do something about it. Sometimes, however, government intervention can make things worse for the environment. A good recent historical example of this is the Common Agricultural Policy of the European Union, known as the CAP.

Through a complex system of import levies, export subsidies, and intervention buying, the CAP offered farmers throughout the EU higher prices for their outputs than the free market would generate. The official reasons given were to support farm incomes, stabilize crop prices, and increase self-sufficiency rates for major food products. Casting aside the dubious economic arguments that could be made in support of these goals, let us focus instead on what the implications were for the environment. Farmers responded to higher prices by increasing output, both by cultivating more land and increasing the intensity of production on all farmland. Coupled with technological progress, this policy altered the countryside, especially in a country such as the United Kingdom, where farmland accounts for about 80 per cent of the total land area. The most important impacts on farming historically are as follows:

- The removal of hedgerows, the ploughing of field margins, and the loss of farm woodlands
- Replacement of permanent pasture with temporary pastures and arable
- Land drainage
- Significant increases in the use of pesticides, herbicides, and fertilizers

The effects of these changes in farm practices on wildlife and landscape were extensive.

Figures from the Nature Conservancy Council published in 1983 showed, for example:

- A 50–60 per cent loss in lowland heaths since 1949
- A 95 per cent loss in herb-rich flower meadows
- An 80 per cent loss in chalk grasslands of high ecological value

The CAP has since been reformed, with the removal of many aspects of direct commodity price support and the introduction of a range of agri-environmental schemes that offer farmers payments to undertake actions likely to produce environmental benefits, such as reducing grazing pressures or cutting fertilizer or pesticide application rates. For an analysis of the effects of agricultural land use on one indicator of biodiversity—birds—see Dallimer et al. (2009).

government to dictate the price of stocks and bonds. We ask the government to help establish, monitor, and enforce the trading rules of the market, but not the market price itself.

We now consider three ways to create new markets to address market failure associated with the environment. A market has a supply side and a demand side, which together produce quantity exchange at a market price that reflects the value of the asset. Working from this basic construction, we consider three options. First, we can assign property rights for environmental assets and let people negotiate over the price and quantity of the good. Second, we can work through regulators to set a market price per unit of the environmental asset and let people decide how much of the asset they want to buy. Third, we can use regulators to set the quantity of the asset that can be bought and let the people decide what price they are willing to pay for the fixed quantity. Let us consider each in turn.

Assign property rights and bargain over price and quantity. In 1960, economist Ronald Coase argued that we can create new markets for non-market goods such as the environment as long as we are willing to remove institutional constraints to assigning well-defined property rights. Coase noted that two parties have an incentive to negotiate an economically efficient and mutually advantageous solution to a dispute provided that one party is given unilateral property rights to the asset in question. The key point for efficiency is that property rights are assigned, not to whom they are assigned (that is an equity issue). The outcome is the same—an efficient allocation of resources. This is the Coase theorem as traditionally defined. The theorem only holds if transaction costs are low and legal entitlements can be exchanged and enforced. The transaction costs are the price paid to organize the economic activity, including information, negotiation, writing and enforcing contracts, specifying property rights, and changing institutional designs.

Consider our open space example. Riley and Ole disagree about the amount of open space in Centennial valley. Riley owns the right to build his house on the hogback and Ole owns his own land below. All the citizens own the right to the view of the Centennial valley and ridge around the Snowy Range Mountains. But again, the rub is that if everyone owns the right to the view, no one owns the right.

One solution to the conflict is for the government to intervene and restrict Riley's development, or to tell Ole to live with it. This command-and-control approach allows a third party to select the winner and loser(s). In contrast, the citizens or government could work with the community to create a dispute resolution process in which everyone could sit down at the bargaining table, and try to come to some agreement. This would require the stakeholders to agree to some solution. It would promote collaboration, and by avoiding polarization it would produce creative solutions with political momentum. This could support local leadership and collaborative efforts to help Riley and the others enhance the environment and achieve economic productivity—or it could collapse into a bitter confrontation, with no resolution.

Alternatively, we could promote a Coasean solution that would create a market for open space by assigning property rights to either Riley or Ole. First, suppose that a third party assigns the property rights to Ole. Ole now has the right to keep the open spaces as he sees fit. This means that he could choose to keep the open spaces to his liking, and have the legal power to ask Riley to build a house that would not disturb the open spaces. This would make Ole the supplier of development, and he would have to decide how much open space to surrender to development under a specific price schedule. Ole's supply schedule of development space reflects his marginal costs from development, which increases as development increases: the more development, the greater loss of well-being as measured by more marginal costs.

Riley would now have to come to Ole to buy the use of open space for development. Riley would have to decide how much open space he would demand for a specific price. His demand schedule for development space would reflect his marginal benefits from development, which are assumed to be decreasing in development—more development implies lower marginal benefits. The market for development space would clear at the market price at which all the space demanded would equal all supplied. Both Ole and Riley would benefit from the trade: Ole would earn benefits pay through receiving a price for open space that exceeded his opportunity cost; Riley would benefit from paying for space that is less than his benefits for developing. The market leads to an efficient outcome—the marginal benefits from development equal the marginal costs.

Now suppose instead that the third party assigned the property rights to Riley. The roles are reversed—Riley has the right to use up open space as he sees fit; Ole has the right to try to buy open space. The power of the Coase theorem is that the outcome would be the same as before—same market price, same market quantity, same efficiency result. The economic efficiency is the same. Note, however, that the distribution of wealth differs. Now Riley receives the payment from Ole. Whether this is preferable to Ole receiving the wealth is a question of ethics.

What happens is that we redefine Ole and Riley's schedules to reflect the new property-right structure. Ole's supply curve for development now becomes his demand curve for open space. Ole now demands open space—different levels for specific prices. Riley's demand curve for development now becomes his supply curve for open space. He chooses to sell different amounts of open space at specific prices.

The market for open space would again clear at the identical market price at which all the space demanded would equal all supplied. Ole and Riley would capture the same benefit from the trade: Ole would earn benefits by paying a price for open space that was less than his opportunity cost of development; Riley benefits from selling open space at a price that exceeds his costs for forgoing developing. Again the market is efficient—the marginal benefits from open-space preservation equal the marginal costs. The smaller the number of people involved in the dispute, the more likely the Coase theorem is to work. More people increase the transaction costs necessary to come to an agreement, and make the market less efficient. It is also hard to identify responsible parties in some cases, and people have more incentive to free-ride in larger groups. In the case of larger numbers, we can still use the market. Now, we can set either the price or the quantity traded, which will allow the market to work.

Set the price of social damage—Pigovian taxes or green taxes. For over a century, economists have promoted the idea that we can adjust market prices to fix environmental dilemmas. We create new economic incentives in the market by altering the relative price of pollution or an otherwise unpriced environmental asset. The economist Alfred Pigou suggested in the early twentieth century that an effective solution to pollution problems is to add a tax on to the market price. This tax, now called a *Pigovian tax* or a *green tax*, would be set to equal the external cost, or marginal damage, suffered by those affected by the pollution.

In principle, society can alter a person's choices by imposing a green tax. The person continues to produce and pay the tax as long as the marginal benefits he or she receives from the output exceed the tax. Once the green tax exceeds the marginal benefits, he or she cuts off production. Ideally, if the green tax is set such that it reflects the equilibrium level of marginal damage, the person voluntarily selects the level of output that is the social optimum.

For Riley and Ole, a regulator could set the green tax at the equilibrium level of Ole's marginal damage suffered from the loss of open space. Now Riley would develop the space up to the point at which the tax exceeded his marginal benefits, and then he would stop. This would result in the same level of development as the Coasean solution.

Economists have debated the idea of whether a *double dividend* exists with green taxes. A double dividend exists: (i) when a green tax reduces the amount of pollution emitted; and (ii) when the revenues raised by the taxes are used to offset other distortionary taxes, such as income or capital gains taxes. A distortionary tax is a charge on activities that society wants to promote rather than discourage, such as working and investing. If we can use green taxes to reduce pollution and then use the extra revenue to reduce income taxes, society has two hits with one shot.

Unfortunately, if you compare it to the optimal benchmark, the double dividend story has a twist. We know that environmental protection benefits society. But this protection also raises prices and the overall cost of living, which cuts into a person's wages. Environmental protection is like a second labour tax added on to the already distortionary income tax. Now, our worker has even more reason to decrease his labour supply, which is an additional cost to society—less labour, less wealth created. The open question remains whether the environmental benefits of the green tax offset the losses from the decreased labour supply. The double dividend moral is this: correcting one mistake can make a second even worse.

In addition, green taxes have historically been set to raise small amounts of revenue, not induce big changes in behaviour. Taxes have been set too low to induce people to increase significantly either pollution abatement or environmental protection (Hanley et al., 2007). This reality is in part affected by the lack of information required to successfully implement an incentive to approach some social goal. Setting an efficient green tax requires information on the marginal costs and benefits schedules. We also need information on the environmental fate and transport systems and the monetary value of risks to life and limb. The information required to find this marginal damage function is not free, as had been presumed in the original green tax models.

An important extension to the idea of an optimal green tax was suggested by Baumol and Oates four decades ago. They showed that if society wants to achieve a given target level of emissions reduction, a tax on emissions could be the lowest-cost way of achieving this reduction, compared with a command-and-control approach such as design standards (the government tells firms how to control pollution) or a performance standard (the government tells them how much pollution they can emit). To understand the Baumol and Oates approach, we introduce the idea of a *marginal abatement cost* curve, or *MAC*.

Abatement costs are the costs of reducing emissions. Polluters can reduce emissions by several alternative means:

- installing 'end-of-pipe' treatment plants;
- changing their production processes—for example, by using cleaner inputs or recycling waste; and
- reducing output.

We assume that polluters know the range of options, and that they always choose the lowest-cost means. This may vary as emissions are reduced. A useful way of thinking about how abatement costs vary with the level of emission reduction is shown in Figure 2.2(a). This shows a *MAC* for a firm, the Jones Company, for reducing emissions by installing end-of-pipe treatment. The graph is read right to left, since this shows falling emissions. As can be seen, the *MAC* rises as the firm progressively cuts back on its emissions.

What does this mean?—That as emissions fall, the additional cost of reducing emissions increases still further: in other words, that the marginal abatement costs are rising. This rising *MAC* curve is an almost-universal empirical finding. Marginal abatement costs 'take off' at a 75 per cent cut, since end-of-pipe technology cannot cut emissions by more than this. The area under the *MAC* curve at any point gives the total abatement cost (e.g. area 'a' shows the total abatement cost of going from 100 per cent emissions to 75 per cent emissions). Figure 2.2(b) shows the *MAC* curve for the firm defined across all emission reduction options; this is now flatter past the 75 per cent reduction level, as the firm can choose to use other methods, such as changing its inputs. The curve takes off at extreme levels of emission reduction (95%), as emission reductions become costly at high levels of pollutant removal.

Another important empirical finding in the literature is that *MAC* curves can vary across firms for the same pollutant. For example, for Biological Oxygen demand (BOD) discharges into the Forth Estuary in Scotland, the actual abatement costs varied by as much as 2,500 times per kilogram of BOD removed. This may be due to firms that emit the same pollutant having different production processes and plant; having different levels of managerial skills; or being located in different areas (e.g. facing different transportation costs for bringing in cleaner inputs).

In Figure 2.2(a), we show the *MAC* curves for Jones PLC and another firm (Bloggs) that both emit BOD into a river. Jones has higher abatement costs than Bloggs, since it operates a different production process. Assume for simplicity that, in the absence of spending any money on pollution control, both firms discharge the same level of emissions, which we show as e^0 in the figure, equal to 10,000 tonnes/BOD/week each. The total unregulated discharge is 20,000 tonnes/week.[1] Figure 2.2(b) shows the aggregate *MAC* curve, which is just the horizontal summation of the *MAC* curves of Jones and Bloggs, referred to as MAC_1. Suppose that a control authority, the Environmental Protection Agency (EPA), wants to get an overall reduction of 10,000 tonnes/week. How could it use economic incentives to achieve this cut, and what might the outcomes be?

Suppose that the EPA sets a tax of t on every unit of emissions from Jones. This means that if Jones emits level of emissions e', it pays $(e' \cdot t)$ in taxes. What should the managers of Jones do? Imagine that they are emitting at e^0. The best they can do is to reduce emissions to e^t, since above e^t the marginal benefits of cutting emissions are greater than the marginal costs ($t > MAC$); whereas below e^t the marginal benefits of increasing emissions (savings in abatement costs, measured by *MAC*) are greater than the marginal costs (increased tax payments of t per unit). Setting emissions equal to e^t is the firm's best response: it implies an equilibrium of

$$t = MAC.$$

[1] In practice, it is unlikely that firms would engage in *zero levels of* pollution control in the absence of regulation, since some control might result from changes in operations motivated by cost saving, such as the recycling of waste streams.

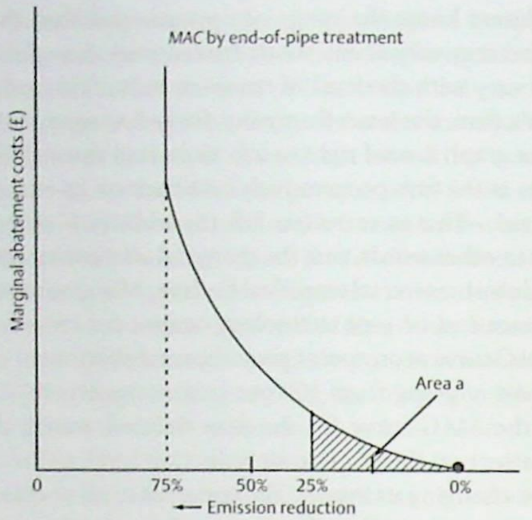

Figure 2.2a A *MAC* for the Jones Company, for reducing emissions by installing end-of-pipe treatment.

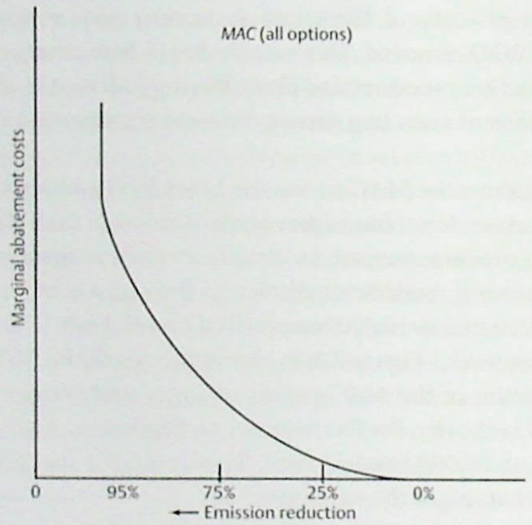

Figure 2.2b The *MAC* curve for the Jones Company defined across all emission reduction options.

At e', the firm pays tax revenues equal to area 'b' and, relative to the pre-tax position, has increased abatement spending by area 'a'. How does the control authority know what tax rate to set? Suppose that it knows the marginal abatement cost schedule MAC_1 from Figure 2.3. The target level of aggregate emissions is shown at E^*, equal to 10,000 tonnes of BOD. At this emission level, the aggregate *MAC* schedule has a value of *t*. This is the correct level at which to set the tax to achieve the target reduction.

But what if the EPA does not have this information? Information, or the lack of the needed information, is a likely scenario in reality. The EPA has to guess the tax rate and observe firms' reactions. If it sets the tax rate too high, the firms cut emissions by too much and the target is overshot;[2] if the tax rate is set too low, the opposite occurs; emission

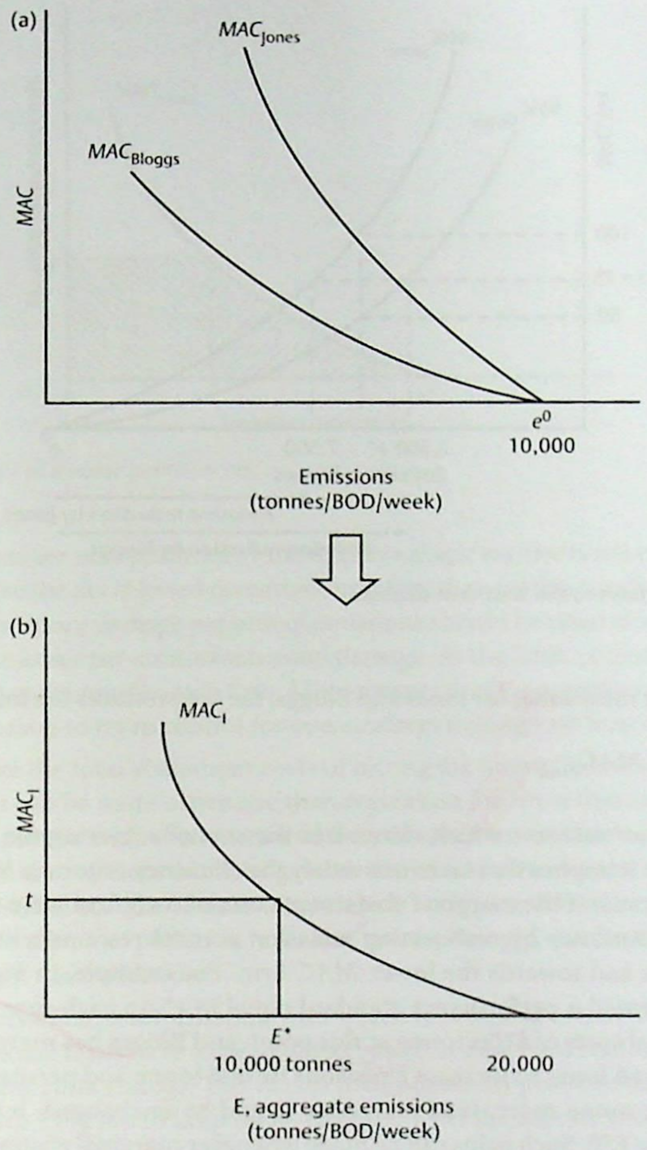

Figure 2.3 From firm's emissions to aggregate emissions.

reductions are undersupplied. The EPA iterates on to the correct tax rate. Firms dislike this approach, however, since the future tax rate is uncertain, which makes planning for investment more difficult.

The main attraction of taxes over regulation can now be explained. Under a tax, and as already shown, each firm's best response is to adjust its emissions so that we get

$$t = MAC$$

[2] While it might seem odd to talk about achieving too much pollution control, remember that each extra reduction in pollution that we aim for imposes a cost on society.