Market Failure

15

There's a joke that goes something like this: how many economists does it take to screw in a light bulb? Answer: none, the invisible hand will take care of it.

Well, not exactly. It is true that economists tend to view the world through the prism of markets. The majority of them probably think the Market Welfare Model, the claim that market equilibrium is also the best outcome for society, is *mostly* correct, and they tend to have an above-average inclination to let markets run without interference. Nevertheless, economists do not go so far as to think themselves out of a job, for if markets could always be relied on to run smoothly on autopilot there would be no need for the services of economists.

It would be more accurate to say that most economists give markets the benefit of the doubt, but they are always scanning the horizon for exceptions, for situations in which markets either don't work properly or otherwise fail to meet the needs of society. (Note: this discussion applies only to microeconomics, the subject of this book. In the realm of macroeconomics most economist recognize the need for some sort of permanent economic management.) The core of this enterprise is the theory of **market failure**, a systematic analysis of why markets sometimes come up short and how they can be fixed.

In this chapter we will examine two of the main forms of market failure, public goods and externalities. Pay attention: nearly all microeconomic policy draws on one or the other of these concepts.

15.1 Public Goods

The theory of **public goods** is a bit of a composite, since it began imprecisely and only gained careful definition after many years of use. When the ideas got sorted out it became clear that there were two *different* criteria, and a good could meet one standard and not the other.

The first criterion is usually called **nonexclusion**; it describes goods for which it is impractical to deny use or access to those who do not pay for them. To understand the importance of this point, consider a typical item you might buy at a store, such

P. Dorman, *Microeconomics*, Springer Texts in Business and Economics,
DOI 10.1007/978-3-642-37434-0_15, © Springer-Verlag Berlin Heidelberg 2014

as a loaf of bread. The bread has a price; it is expected that you will stop at the checkout counter before you leave and pay this money to a store employee. If you don't you will be committing an offense, shoplifting, for which you could be arrested. The store stocks bread because it anticipates that those who want it enough will pay for it. Suppose, however, that for some reason it is not possible to prevent the shoplifting of bread or even discourage it: customers can take the bread home with them without giving it a thought and never pay a cent. Stores will discontinue carrying bread altogether, because they would suffer a loss with each loaf they "sell". Yet people may still desire bread and might even be willing to pay for it, if it were required of them.

The example is absurd, of course, because stores *can* make people pay for bread and similar items. (During natural disasters, riots and other disruptions, when stores have been abandoned by their employees, this ability to require payment vanishes, and the result is often looting.) But there are other goods that are very difficult to charge for. A classic example is national defense, the "good" that consists of defending a country against external attack. Could this be sold the way bread is? You could try: imagine a defense "company" which sells certificates entitling the bearer to freedom from attack by a foreign power. What would you do if someone chooses not to pay? Could you organize a defense of the entire country except this one non-buyer? ("We will defend our country, except that you can bomb this one house over here.") It's just not practical, and this is why it is difficult to organize national defense along the lines of a consumer market.

The same logic applies to clean air. Most countries now have agencies whose job is to regulate air pollution; could this be "sold" to customers on a cash basis? If someone decides not to pay their clean air bill could you cut them off—give everyone else clean air but make their air dirty? It turns out that there are a number of important goods in our society that resist the market approach of denying access to those who don't pay.

There are two subtleties to be aware of, however. First, goods are sometimes given away as if the nonexclusion property were applicable, when in fact it is simply a choice of the provider. A band might offer a free concert to its fans, for instance, but this doesn't mean that their music is nonexcludable. They can make the decision to do this, but it would also be practical to sell tickets and exclude those who don't buy them; bands do this all the time. So: just because users are not excluded doesn't mean that the good they are using is nonexcludable.

Second, there is a spectrum of nonexcludability in the real world. Take the case of police services. You might think this should be nonexcludable: can the police deny protection to individuals who haven't purchased a protection "ticket"? Well, no and yes. If an officer sees a crime being committed on the street, there will usually not be an opportunity to see who has paid what; protection must be automatic. But the payment of police for particular protection services (including from the police) is a common form of corruption, and, in addition, it is entirely possible for individuals or businesses to purchase "extra" protection from private police forces. (There are now more privately employed police in the United States

than police working for all levels of government combined.) So police services are in a grey zone, partially falling under the nonexcludability criterion, partially not.

The second criterion is **nonrivalry**, but a more precise description would be that the good or service can be provided at (or near) zero marginal cost. Again, consider national defense. If the population of a country expands by one (a birth or the arrival of an immigrant), does defending the country become more expensive? Presumably the answer is no, so there is no marginal cost to defending the additional resident. The same goes for clean air: more breathers does not entail more cost. (More potential polluters does, but they are not *users* of the clean air regulations.) An early example in the economics literature was lighthouses: it doesn't cost more to maintain them if additional boats are guided by them through shallow or rocky passages.

As with nonexclusion, nonrivalry presents a large grey zone, since many goods have a low, but not zero, marginal cost of provision. A national park which is mainly used by hikers has some additional cost per user, because more hikers require more labor to keep the trails and campsites from being degraded. Nevertheless, the largest cost (by far) is the provision of the land itself, and this does not vary with use. For instance, a scenic mountain which has been preserved in its original state (more or less) for the satisfaction of hikers could otherwise have been developed into a ski resort. This might constitute the **opportunity cost** of using the resource for a park, and this cost is unrelated to the number of hikers who take advantage of the park. But perhaps the land has little use other than hiking (low opportunity cost) and maintenance is the main expense; in that case, on balance, the nonrivalry property may not apply. There is a lot of room for careful analysis and judgment.

These two properties, nonexclusion and nonrivalry, together define what it means for a good to be public in the way this term is used by economic theorists. If both properties hold it is said that we have a **pure public good**. But this is quite different from the way the word "public" is used in everyday contexts. Normally, "public" refers to the public sector; that is, government. A public agency is a branch of the government, unlike, say, the "private sector", which includes businesses but excludes the government. We are not going to rewrite the dictionaries, but it is important to clear up this source of confusion. Note well: *public goods in the economic sense are not necessarily provided by the public sector, and goods provided by the government are not necessarily public goods as economics defines them*. It may be helpful to reread this sentence and give it some extra thought; perhaps no misunderstanding in introductory economics is as widespread as the confusion between these two uses of "public".

Examples are not difficult to come by. Radio broadcasts (but not radios) are pure public goods. Nonexcludability applies, because it would be difficult to prevent someone with a radio from picking up a signal because they haven't paid for it. Nonrivalry applies because it is no more expensive to provide a broadcast to more people (within a given geographic area) than fewer. Yet most radio stations are private, for-profit enterprises. The trick is that they sell their audiences to advertisers, so that, from a business standpoint, the advertisers and not the listeners

are the true users. When you look at it that way both criteria no longer apply: advertisements can be taken off the air if they are not paid for, and it is costly to provide airtime for each additional ad. Does this mean that the broadcasts were not public goods in the first place? To listeners they still are. Moreover, some radio stations are listener-sponsored: rather than selling ads, they plead with their listeners to contribute money during periodic pledge drives. And in most cases enough listeners send in money even though their listening privileges would not be revoked if they didn't.

Meanwhile, postal services are provided by public agencies in most countries. (In the US this is the job of the US Postal Service.) According to everyday usage, this would make their services "public". Yet neither criterion of economic publicness applies: the post office certainly has the ability to deny services to those who don't purchase stamps, and the delivery of each additional letter or package is costly. So this public outfit is *not* in the business of providing a public good.

So, if the theory of public goods does not predict which goods will be provided by the public sector, what's the point? The answer is that each criterion is associated with a particular sort of market failure. Either the failure is fixed, or there is a cost in terms of economic efficiency.

First consider nonexclusion. The problem that arises is that some people, maybe most, will use the good without paying; economists call this the **free rider problem**. That's fine for the happy freeloaders, but how will the good in question be financed? It costs money to provide, and if not enough can be raised by selling access, there may be no provision at all. If the clean air agency tries to fund itself by selling people the right to breathe clean air, each person may think, "Why should I pay when I can get the same air by paying nothing?" And if enough people think this way (which they will if they have the self-interested values economists ascribe to them), the agency will run short of money and have to shut down (or regulate less effectively). Free ridership can also lead to overuse of services that have significant marginal costs. As we will see in a later chapter, this has been the story of the world's fisheries. For each fisher, the population of fish has been a free good; no one has made them pay the true cost of depleting it. The result has been catastrophic overuse, to the point where it has become necessary to shut down entire marine regions so that the fish can repopulate.

The problem with nonrivalry, on the other hand, is the possibility that there will be *too little* use of the good. If the marginal cost is zero, then any price that may be charged will discourage some users whose personal benefit does not justify paying for it. Yet, as long as they derive any benefit, it is economically inefficient to exclude them. This point is illustrated in Fig. 15.1, where marginal cost is zero for all users and the marginal benefit is given by the demand curve.

The economically efficient level of supply is Q*, since up to this point the marginal benefit exceeds the marginal cost for all consumers. Of course, that means that there is no revenue to defray any of the *fixed* costs the supplier incurs. So suppose a price P_1 is charged to cover these costs. Now only Q_1 units are sold to consumers. Yet consider a consumer who falls between these two points, to the right of Q_1 but the left of Q*. This person would receive positive benefit, as

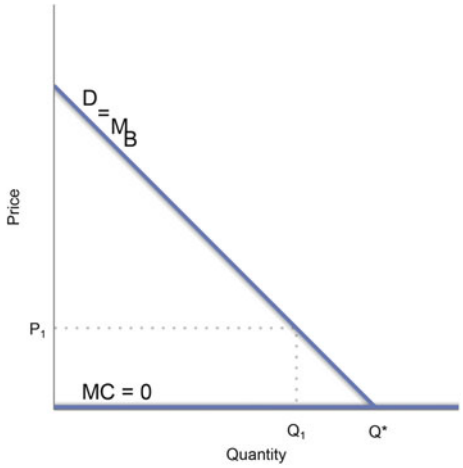

Fig. 15.1 Marginal cost and benefit for a perfectly nonrival good. When the MC = 0, the principle of MB = MC means that the optimal quantity is Q*. By charging price P_1, however, the seller excludes individuals (Q^*–Q_1) from the market

indicated by the demand curve, and she would not impose any additional cost. It would therefore be economically rational to enable her to acquire the good (MB > MC), but she has been shut out by the price. So Fig. 15.1 portrays a dilemma: either no money is raised to cover fixed costs, or a price is set which dissuades consumers who would be more than willing to pay their marginal cost (0). That's the problem.

The dilemma of nonrivalry is at the center of one of today's most hotly-debated economic questions, what to do about the explosion of digital reproduction over the internet. With each revolution in technology the relationship between the fixed and marginal costs of providing products like recorded music has shifted. Thirty years ago the fixed cost of paying musicians and studio engineers were matched by the expense of producing and packaging bulky LP's. Then came CDs, which were cheaper to produce and distribute. Now, with digital compression, inexpensive data storage and distribution through the internet, there is nearly no marginal cost of making an additional reproduction of music, movies or other information products. The situation closely resembles Fig. 15.1; users would like to download for free, bringing the quantity of files close to Q*. Companies holding copyrights in music, film and literature want to charge for each access, imposing a price like P_1. This is a true dilemma: if the downloaders get their way, there is a risk that artists, for example, will not be paid for their work. If the copyright holders win, many people who would receive a personal benefit from downloading these files, and who would not add costs to any member of society, would be priced out.

This just scratches the surface of the issue; there are many more complications, so I will leave it to you to think about possible solutions. In general, however, economists tend to favor one particular approach to public goods: tax the community of potential users and provide the good free of charge (at marginal cost). This is what we do for national defense; we levy a tax on the entire country and ask for no other form of payment. (Some—perhaps a lot of—military spending is not for defense per se; if it serves special interests, in principle it may be possible to ask those who benefit to pay for it. But waging war for the private benefit of a few also

raises serious ethical issues!) This addresses both problems we described earlier. It overcomes the free-rider problem by requiring everyone to pay. It overcomes the nonrivalry problem by not charging for use.

This sounds like just the ticket, but it is not so easy in practice. First, there is the question of just how much of the public good to provide. With goods sold in the market the question is answered by supply and demand, but who decides how much is enough of any public good? Second, there are usually qualitative aspects of these goods that have to be determined. There are many ways to provide national defense; a clean air agency has to prioritize some health risks over others. Once again, there is no market guidance available, and choices have to be made. Third, it is rarely the case that public goods benefit everyone equally; there are usually some who have a stronger interest than others. For instance, while all of us presumably value clean air, people with asthma or other respiratory disorders have an even greater interest. Should those who benefit more pay more? How would you know who they are and how much greater their benefit is? Finally, all of the preceding questions take for granted the desire of the government to do the right thing, but as we saw in Chap. 9, it is not so simple. Governments may lack the capacity to do the job properly, or they may be captured by special interests with their own agenda. To remove public goods from the market and place them in the hands of government is not to solve the problem, but to replace one set of problems with another—which *might* be the right thing to do anyway. (Economists have given a lot of attention to the questions raised in this paragraph, and you will discover some of their answers if you take a course in Public Finance or Public Policy Analysis.)

15.2 Externalities

A market is essentially just the sum of lots of bilateral (two-party) transactions. Individual buyers and sellers find agreement and exchange money and goods. Each participant presumably acts in what he or she (or it, such as a company) believes is his or her best interest. All actions are voluntary in the narrow sense that if a transaction does not offer at least as much as the status quo it will not be accepted. So, putting it all together, it is plausible that market exchanges should never make anyone worse off than they were before, and that if every agreement that is potentially beneficial to two parties is agreed to, markets should work optimally to promote economic well-being.

This is the intuition behind the Invisible Hand hypothesis. Early economists like Adam Smith thought it was so self-evidently true that there needed to be little analysis of or argument for it. Nevertheless, it ignores a crucial possibility: what if an agreement between two people, say A and B, has significant effects on a third party, C, who is not part of the deal? Now the Invisible Hand falters. If these third-party, or **spillover**, effects are undesirable, then transactions could be made that lower the well-being of the community. And if the spillover effects are positive, then perhaps too few such transactions will be made, since A and B are not taking into account the benefits received by C.

In a nutshell, this is the insight that underlies the theory of **externalities**. Of course, there is much more to say than this, and we will discuss some of the wrinkles in the paragraphs to come.

For many years the topic of externalities was confusing even for high-level economists. After all, *every* market transaction affects third parties in some way. For instance, if you fill up your car with gas, the effects go beyond you and the gas station (or oil company). By adding to demand, you help increase the price (or slow down its decline), and this effects not only everyone else who might buy or sell gas, but also those in related markets, like autos, air travel, etc. In a well-developed market economy, just about everything is connected to everything else, so how can there not be third-party effects? Does this mean that all market transactions entail externalities? But the whole point to markets is that prices *should* reflect the decisions of buyers and sellers; that's what makes them tick. So what is the difference between an "externality" that reflects the proper operation of markets and one that undermines them?

All confusion was swept aside when Ronald Coase (who we met in the chapter on businesses) published a remarkable article in 1961, "The Problem of Social Cost". In it he not only defined an externality in a way that was both extremely simple and perfectly precise, but also placed it in a rich context of law, social institutions and individual interests. First the definition: an externality is an effect, positive or negative, of a missing market on those who would otherwise have been parties to it.

Suppose, for example, that you and I are neighbors. You have grass in front of your house, and I have a sheep. You might come to me and ask, "Would you be willing to have your sheep graze on my grass once a week? It gets too long, and I don't want to have to cut it myself. I'll pay you if you want." Then we could discuss how much you should pay, and the result would be that you get the benefit of having your grass cut, and I get some money to help maintain the sheep. This is a "normal" market transaction: there is a buyer, a seller, a service, a cost, a quantity, a time frame, a price. I might even take the prospect of a deal like this (or many of them, if I have many neighbors) into consideration if I am thinking about getting a sheep in the first place. The more you (and others) benefit and are willing to pay, the more it makes financial sense to acquire a sheep.

What if, however, you and I never made an agreement, but the sheep likes to wander around (or jump its fence) and graze on your grass anyway. You still receive the same benefit, but now you are not paying for it. This benefit is now an externality. There is no market (buying, selling) between us, but a benefit has been delivered nevertheless. It is a potential economic issue because I may be less likely to buy the sheep initially if I don't anticipate being paid for the good it does for others. In Market Welfare Model terms, there is a difference between the marginal benefits generated by the sheep and the marginal benefits *I* receive from it, so my calculation of personal cost vs benefit does not reflect the wider neighborhood interest.

Of course, having a sheep graze randomly wherever it wants can be a problem as well. Perhaps you have planted vegetables by your house, and now the sheep is

nibbling on them too. If I insist on letting the sheep roam freely (or not fixing my fence), you could come to me and say, "You have no right to ruin my vegetable garden with your sheep. If you want to continue this, I think you need to compensate me." And we could discuss just how much compensation you would need. Such a transaction would convert your personal cost, vegetables, into my cost, money. I would have to take it into consideration in all sheep-related decisions: whether to buy a sheep at all, whether to fix the fence, etc. The marginal costs and benefits of each decision would each play their proper role.

The other possibility, however, is that my sheep ruins your vegetables, but you are not in a position to make me pay for this infringement. Perhaps we live in a society whose laws do not give you any right to take action against me and my sheep. In that case there would be a missing market, and you suffer a negative externality. Then, if I am self-interested (as always, the default assumption in economics), your cost no longer enters my calculations. I will buy and raise sheep without any consideration of the impact it has on you. I will be too willing to purchase sheep and not willing enough to confine them.

In both instances, the critical feature that separates a "normal" market outcome from an externality is the presence of a market at all. If someone gains a benefit or bears a cost, and if there is no market to purchase the first or sell the second, then we have an externality—and therefore a market failure. Figure 15.2 on the next page depicts these two cases, where the cost of sheep is the cost the owner pays to acquire them, while the benefit is the sum of all the benefits me and my neighbors get from the sheep *minus* any costs they impose on us. For simplicity it is assumed that the marginal cost of buying sheep is constant; each animal costs the same.

In the first diagram we assume there are **positive externalities**. This means that for any given sheep, the marginal (additional) benefit to society (MSB) will be greater than the marginal private benefit (MPB). The vertical distance between these two curves, for instance the height of the triangle at Q_1, reflects this difference. (As drawn the size of the externality does not change from one sheep to the next—the curves are parallel—but this does not have to be the case.) The sheep owner, being self-interested (assumption), considers only his or her private benefits. This leads to Q_1 sheep being purchased: every sheep to the left of this provides a private benefit greater than the cost; every sheep to the right a cost greater than the private benefit. Nevertheless, if one were to add in the external benefits and use the MSB curve, it is clear that the desirable number of sheep is larger, at Q_2. Each individual sheep between Q_1 and Q_2 generates a marginal social benefit above the cost; adding these surpluses of benefit over cost all together gives us the area of the shaded triangle, the net benefits society would have had at Q_2 and foregoes at Q_1.

The situation is slightly different in the diagram on the bottom. Now the externality is assumed to be **negative**, so the marginal social benefit curve is below the private benefit curve. Failure to take externalities into account (because they don't have to be paid for) leads the sheep owner to buy Q_2 sheep when the social optimum would have been Q_1. Now the shaded triangle represents the net costs to society of having too many sheep.

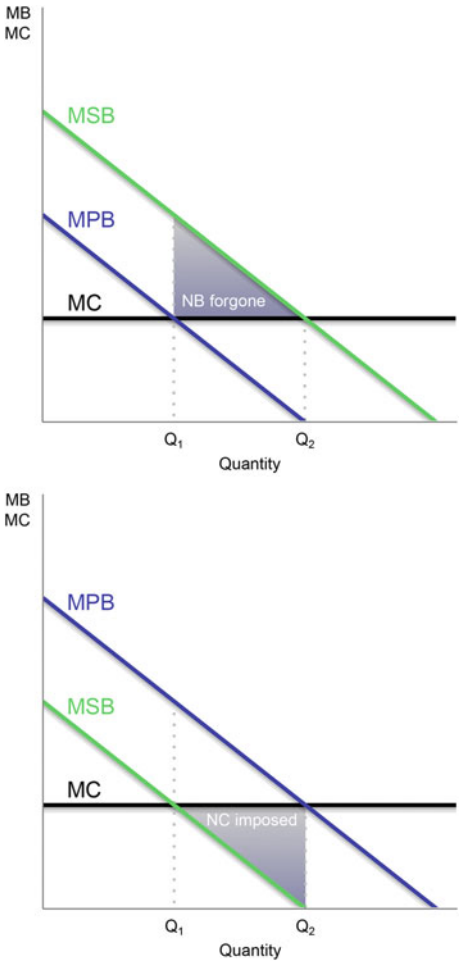

Fig. 15.2 Private vs social benefits of owning sheep. (**a**) Due to a positive externality, the marginal social benefit (*MSB*) of sheep ownership is greater than the marginal private benefit (*MPB*). Self-interest leads to Q_1 sheep being acquired, rather than the social optimum, Q_2. The shaded triangle shows the amount of net social benefits (total benefit minus total cost) forgone at Q_1. (**b**) Due to a negative externality, the marginal social benefit (*MSB*) of sheep ownership is less than the marginal private benefit (*MPB*). Self-interest leads to Q_2 sheep being acquired, rather than the social optimum, Q_1. The shaded triangle shows the amount of net social cost (total cost minus total benefit) being imposed at Q_2

To summarize the argument so far, the logic of a market economy is that, if decisions about what and how to produce and consume are made by individuals acting separately in the market based on their personal benefits and costs, then there needs to be a market for every one of these impacts. If there are beneficial effects that no one has to pay for (due to a missing market), there will be an undersupply of them. If there are harmful effects that no one has to be paid for accepting, there will be an oversupply. It's a problem to have "holes" in the market system.

To illustrate the logic of externalities it was useful to have a purely hypothetical example that comes to us without all the usual real-world complications, but to see the theory in action, let's switch back to reality.

The most important negative externality in the real world is pollution. Many production processes, and many consumption processes as well, generate harmful byproducts that foul the environment and threaten our health. This represents a

market failure insofar as these harmful effects do not have to be paid for by those who produce them. That is, a factory has to pay workers for giving up other uses for their time and accepting the authority of the employer. It has to pay the local electricity company for using electricity. It has to pay a bank for any loans it may have taken out—the use of the bank's money. But, in general, it does not have to pay for the use of the air, water and other natural resources it may damage through pollution. When it issues its quarterly report it has to calculate its profits, and these include the revenues it gets from sales and the costs it has to pay to operate. Costs *not* paid for, like pollution externalities, are not part of this calculation. If the business is guided by the profit motive it will choose production methods that minimize the costs that have to be paid and shift more of the impact to those that don't. If a factory has to pay more for materials and labor to cut down on pollution, in the absence of any force outside the market, it will pollute more and keep a lid on its monetary costs.

Looking at the situation from the perspective of an economist, we can recognize quantity effects and price effects. The quantity effect is that there is too much pollution. If firms had to pay the true cost of pollution to society there would be a lot less. (If firms could employ labor for free they would use too many workers as well.) The price effect is that the cost of the goods we buy often does not reflect the true marginal cost to society—the resources that are actually used in production. Burning petroleum products produces air pollution and contributes to global warming, but these costs are not included in the price. This means that gas oline, diesel and similar fuels are priced below the true cost to society of burning them. Insofar as the cost of fuel enters into the pricing of many other goods, the entire system of prices is divorced from the costs and benefits they are supposed to reflect. Thus, key externalities ripple outward through the economy, causing distortions in the pattern of what we produce and consume.

Box 15.1: A Humourous Look at Pollution Externalities

A Cuban animation released in the 1970s begins with a man working in a factory. He has a bad cough, and it just gets worse. Finally he puts his tools down, clocks out, and leaves. You watch him walking down the street, which is enveloped in clouds of soot and smoke coming from his factory. He coughs all the way to a drug store, where he buys some cough medicine. Glug, glug: he drinks it down, and his throat is soothed. Now he can go back to work. You see him walk past the factory again, belching its pollution. Then the camera pulls back, and you see what the factory produces: cough medicine!

But not all externalities are negative. Education produces enormous positive externalities, benefits to society for which students, whose choices make this possible, are not paid. Very loosely, every time a student increases her education level there are private benefits and external ones. Private benefits include better job prospects, the satisfaction from learning more about the world, and greater self-knowledge.

External benefits include the improved functioning of democracy (which thrives in an educated citizenry) and additions to the fund of knowledge available to everyone, such as in the sciences and arts. But the decision to acquire more education is made at the individual level on the basis of private costs and benefits—the financial cost to the student, the opportunity costs, the effort required to succeed in class, the personal gains expected to result. Since no one pays the students for the services they also provide to society this doesn't enter the calculation. Presumably there are some students who, on the basis of their own costs and benefits, would choose to pass up school, but who, if one considers the social as well as the personal aspects, *should* go to school. The result would be an under-educated society. (Note: this analysis applies even if each student has perfect foresight about the future benefits of education, which is hardly the case.)

15.3 Remedies for Externalities

Since externalities are extremely common (some would say ubiquitous), and since many of them have important consequences, economists have given lots of attention to the search for solutions. Very generally, we can speak of four types of responses:

1. Moral suasion. In many cases individuals can be urged to consider the impacts their actions have on others; in other words, they can be asked to be less self-interested. Since such appeals normally coincide with social norms in most societies, they can be effective. For instance, studies have shown that one of the most powerful anti-pollution policies instituted by the US government was the Environmental Protection Adminstration's Toxic Release Inventory. This program named names: it listed all the major manufacturers and provided a record of how much (and what kind of) toxic chemicals they emitted into the environment, even if perfectly legal. This exposed the companies to public pressure, and the evidence is that they took extra steps to lower their profile as polluters. To at least some extent this represented a retreat from raw self-interest, although it is possible that a bad environmental reputation might also hurt sales, employee recruitment and other bottom-line factors. Much of the corporate social responsibility movement, which we surveyed in the chapter on firms, is predicated on the belief that moral suasion can work.

2. Direct regulation. Often we turn to the government to issue directives telling those who create externalities how to modify their behavior. There are laws limiting the amount and type of pollution firms can emit into the air, water and soil. Neighborhood associations lobby for regulations that discourage some land uses and encourage others, since development can have either positive or negative external effects on those living next door. There are other regulations on the content of radio and TV broadcasting, based on the belief that externalities arise here too. Many economists are skeptical of the value of these regulations, believing that they tend to be cumbersome and inefficient, and that they create opportunities for special interests to gain unfair advantages. Changing a few words in a regulation can

provide a bonanza for some producers and an insuperable handicap for others. Economists who feel this way tend to support market approaches (see below), but the general public tends to be more inclined toward regulation because of the ethical message it sends: "it is right to do it this way, so we will make you do it".

3. Taxes and subsidies. These are also sometimes called "Pigovian" taxes and subsidies after their renowned advocate, A. C. Pigou, a professor of economics at Cambridge University (England) during the first decades of the twentieth century. To see the logic, look again at Fig. 15.2a, b. In the first case, for instance, too little is produced or purchased because of positive externalities. The marginal private benefit curve functions as a demand curve, representing the benefit perceived by the decision-maker(s). Pigou's solution is to *subsidize* such a good by an amount equal to the vertical distance between the two curves. By doing this, the decision-maker would now face the direct gains from sheep-buying (or whatever), as well as the financial benefits of the subsidy. The combination of the two would add up to the MSB curve, which would now be the new demand curve. Thus Q_2 rather than Q_1 would be the amount produced.

Alternatively, look at the bottom figure. Here the externality is detrimental, and the demand curve, the private benefits on which the decision-maker bases his or her decision, is above the true marginal social benefit curve. In this case the idea is to *tax* the good by this same vertical difference, pulling the demand curve down to the level of MSB. Facing such a tax, the decision-maker would now select the appropriate quantity Q_1 rather than the excessive quantity Q_2.

In either case, the job of the public authority is to estimate the size of the externality and impose a tax or subsidy equal to it. There is no need to precisely specify what decisions ought to be made: if the prices once again reflect the true costs and benefits of each action, it can safely be left to each decision-maker to choose the option that seems best. Indeed, under Pigou's approach, there is no need for the government to know what the best corrective actions are; they need to know only the size of the externalities. Individuals looking out for their own interests may well come up with innovations—new ways of producing or consuming—that the government could not have predicted. On the other hand, switching from direct regulation to taxes and subsidies mutes the ethical message the public may wish to send, and it also requires that a monitoring apparatus be put into place to accurately determine who should be taxed or subsidized and how much.

A hybrid of direct regulation and the Pigovian approach that has proved popular in pollution control is **cap-and-trade**: regulators establish a maximum allowable quantity of pollution and allow those who pollute less than their share to sell their surplus to those pollute more. Like taxes and subsidies, these markets in pollution permits (the right to emit a certain quantity of pollution) create incentives for companies to find less-polluting techniques and products (so they can sell more permits or buy less of them). One disadvantage is that, whereas a traditional regulation will typically result in less pollution than allowed, because some companies will "underpollute", cap-and-trade virtually guarantees that the full allotment will be utilized, since for each "underpolluter" there will now be a corresponding "overpolluter". Of course, regulators can anticipate this and set the

allowable pollution level lower under cap-and-trade. One practical question with potentially large financial implications is how the permits will be allocated initially—whether they will be handed out for free or sold in an auction. On this choice may ride billions of dollars in such economically crucial pollution markets as sulfur and carbon dioxide.

4. Assigning property rights. Now we move into Coase's own territory, as set out in his 1961 article. Let's go back to the negative externality version of the sheep story, the one with the sheep gnawing on your lettuce, since it is similar to one told by Coase. We can imagine that there is a law in the community we live in that says each person has the right to grow vegetables without interference from their neighbors or their neighbors' sheep. The law is enforced with great severity, and I would never think of violating it. Of course, if I could get you to *agree* to have my sheep eat your vegetables, I wouldn't be a criminal, and I could save myself the expense of building higher, stronger fences. So I might enter into a negotiation with you: how much, I ask, would you be willing to accept as compensation for the damage caused by my sheep? You might propose a figure, I would counteroffer, and perhaps we could arrive at an agreement. Whether an agreement is possible depends above all on whether the damage caused by the sheep (the least you will accept) is less than the cost of enhancing the fence (the most I am willing to pay). If this is the case there is a potential zone of agreement, and we are in something like the bargaining world explored in the previous chapter. If not, all deals are off—and they *should* be off, since the cost of restraining the sheep is less than the damage they do if they run free.

As Coase pointed out, the above story does *not* describe an externality, because the requisite market, rather than going missing, is there is broad daylight. You are indeed in a position of selling, and I in a position of buying, the right to impose harm. The market exists for two reasons. First, property rights have been unambiguously assigned to you. You "own" your land in the sense that you have the right to deny me and my sheep the use of it. This gives you something to sell and me something to buy. Second, it is not difficult for the two of us to enter into a negotiation over the price of this transaction. In Coase's terminology, which we encountered in an earlier chapter, there are low **transaction costs** to this process. We are able to communicate easily enough, and it would not be too difficult to draft and enforce a potential agreement.

But Coase noticed something else. Suppose the legal context is different; now there is a presumption that sheep have the right to roam freely, despite their bad habits. The shoe would then be on the other foot: it would be up to you to ask me to confine my sheep, and to make it worth my while you would propose a payment. Would I agree? It turns out that the decisive criterion is the same, whether the cost of building up my fence is less than the damage done by the sheep. If so, we can bargain; if not, bargaining will get us nowhere. Note again: *the potential for an agreement depends on the cost of the negative impact and the cost of preventing it; it is the same irrespective of whether rights are assigned to the one who suffers the impact (and who must be paid to accept it) or the one who imposes it (and must be paid to prevent it).* This is a striking insight, one that is obvious when you think

about it, but which might easily be overlooked. (In fact, Coase went one step further and tried to show that the actual agreement—the price paid, the amount of damage agreed to—would be the same under either system of rights. This turns out to be false for reasons we will return to shortly.) In addition, of course, an agreement in which you pay me to confine my sheep also requires the same two premises as before, that my right to let my sheep run free is unambiguous (and can therefore be bought or sold), and that there are few transaction costs to the process of bargaining.

So Coase recognizes that any clear assignment of property rights can be the basis for a bargaining process—that is, a market—that eliminates the problem of externalities. But Coase is a realist. He knows perfectly well that externalities occur quite commonly, which means that the sort of bargaining mechanism we have been talking about often fails to occur in the real world. Why is this? Sometimes it is due to a poor definition of property rights. We may not know who actually has the right to permit or restrict, and this makes it difficult to set up markets for buying and selling these permits or restrictions. But the main culprit tends to be transaction costs.

We set up an apocryphal situation with sheep and vegetables, but most real-world spillover problems are more convoluted. Take the case of air pollution. Here it is often difficult to identify both the perpetrators and the victims. Many pollute the air simultaneously, each in different ways. Many breathe the air, and each has somewhat different interests and preferences. Assembling both groups into single entities that can bargain as a collective is extremely difficult and expensive. Moreover, as we saw above clean air is a public good, so a bargaining process would have to overcome the free rider problem as well as the zero marginal cost problem. The fact that we seldom see such negotiations taking place is a sign that the costs of surmounting all these difficulties is simply too great.

But once we see the issue through the lens provided by Coase, it becomes possible to envision another approach to remedying externalities: with a big assist from government we might indeed set up a bargaining arrangement that, by constituting a market, eliminates the externality. Japan, for instance, has used this approach in its policies to limit the pollution of its coastal waters. Coasts are important to Japan: as an island nation it has lots of coastline, and it depends on fish and seaweed from these regions for a significant part of its diet. At the same time, a large percentage of the population lives in coastal cities, and the combination of industrial, residential and agricultural runoff has caused severe pollution episodes in the past. As one part of its policy system, the Japanese government assisted the formation of fishing cooperatives, with one coop representing each of the major estuaries. These coops were assigned rights to water quality within their jurisdiction; so anyone who want to dump waste into a river or stream (which will eventually flow to an estuary) has to pay the cooperative for permission. It is not a perfect system, but it helped reverse a serious threat to the country's natural resources.

Coase originally presented his approach as an alternative to the system of Pigovian taxes and subsidies we considered above, but most economists today see them as two sides of the same question. Suppose we find a town situated on a river

that also serves as a waste receptacle for a paper mill. It is more profitable for the mill to emit a higher level of pollution, but residents of the town prefer a lower level. If the town passes a law requiring the mill to pay a pollution tax, what it has done is to effectively claim property rights to the water quality of the river. It is as if the people of this town said, "This is our river, and we are going to charge you to use it." It is not difficult to imagine a period of de facto bargaining, which could take the form of the mill threatening to shut down unless the tax is reduced, and some compromise tax finally being agreed to.

The interpretation of a Pigovian subsidy is exactly the reverse. If the town passes a law offering payments to the mill in return for its adoption of a less-polluting production process, they are in effect saying, "We recognize that you have the right to continue polluting if you wish, but we want to pay you so that you will pollute less anyway." Again, this could be the opening salvo in a bargaining process, paving the way for ultimate agreement.

The point is that there is a correspondence between the type of financial mechanism used to reduce an externality problem and the implicit property rights on which they are based. If the public pays the polluter, the polluter has the implicit right to pollute; if the polluter is compelled to pay the public, it is the public that is acting on its ownership rights. The right policy depends to some extent on what you think ought to be the assignment of rights between the polluters and those who have to cope with the pollution.

One final note: would the agreement between the town and the mill turn out to be the same in either case, as Coase originally thought? Almost certainly not: (1) The assignment of rights affects the disagreement position of the two parties. If the mill has the right to pollute, the default position (if an agreement is not reached) is that the pollution continues. This is better for the mill and worse for the town. If the town has the right to be free of pollution, the default is worse for the mill and better for the town. As we saw in the previous chapter, such a large change in the disagreement position would almost certainly lead to a change in the final bargaining outcome. (2) There are dynamic effects to the assignment of rights. If the town has the right to prohibit pollution and imposes a tax, it will be less profitable to produce paper there. In the long run there will be less investment in paper-making than would otherwise occur. If the mill has the right to pollute and becomes the recipient of subsidies, this will attract new investment to the region—why not open a new mill and apply for these subsidies too? So in the long run there will be more mills and therefore more pollution at any given level of subsidy. (3) There is plenty of evidence that people tend to demand more to give up something they have than they will offer to get something they don't. That is, they are likely to demand more in taxes than they would offer in subsidies. The reasons for this are still under active debate, but the pattern itself is not doubted.

Bottom line: policies to limit detrimental externalities through financial incentives are also expressions of property rights. The choice of whose right should prevail has large economic consequences. These are summarized in Table 15.1.

Table 15.1 The assignment of property rights and economic outcomes

Property rights	Assigned to polluters	Assigned to those affected by pollution
Financial instrument	Subsidies to polluters	Taxes on polluters
Size of financial incentive	Generally smaller	Generally larger
Resulting level of pollution	Generally higher	Generally lower

15.4 New Types of Externalities

Actually, the externalities are old, but the thinking about them is new! Much recent research has gone into **network externalities** and **positional externalities**. These concepts offer new ways to think about longstanding social issues.

Network externalities refer to the spillover effects of individual consumption decisions when the value of a good or service depends on how many users it has. A familiar example is the telephone. If you were the only person in the whole world who had one, it would have little value except perhaps as an art object. It takes two to have a conversation, and the more people who have their own phones, the more benefit this device will give you. To take the opposite extreme, if you were the only person who did *not* have a phone, you would probably feel you were missing something important. In societies where telephone access is nearly universal, all sorts of activities require phone contact.

The same goes for computer software. The value of a program increases as more people use it. With the vast majority of computer users adopting Microsoft word-processing and other programs, for example, those who prefer to use competing programs have to worry about the compatibility of their files: will they be able to exchange files with friends and coworkers? Also, as the number of Microsoft users increases, so does the availability of other programs that add to or take advantage of the features in Microsoft products. (There is a bigger market, which attracts more businesses.) The same issue has appeared in cellphones and tablets with the availability of apps.

It is reasonable to call this effect an externality, because it results from the lack of a market in spillovers. My choice of software (or in the case of the telephone, hardware) changes the benefit you get from your choices, but you cannot negotiate with me to offer incentives so I will make the choice you prefer. True, the spillover for any individual choice is small, but over millions of users the cumulative effect can be enormous. Obviously, the transaction costs required to set up spillover markets on the scale necessary would be beyond calculation—the whole idea seems bizarre.

And why would anyone care about such externalities? One reason is that network effects can lead to the adoption of inefficient technologies. Once a product with network externalities is in wide use, it will be more advantageous for new consumers to adopt it, even though another product might be even better if it were adopted as widely. Without a market in externalities, there is no way to coordinate a mass migration from the lesser to the better product. The product that arrives first on the market and has an opportunity to build up its network can lock out a later, better alternative.

Positional externalities arise in situations where what matters is not how much you have or how well you do, but how much or how well in comparison to others. Suppose, for example, that the best students in a particular country all want to be admitted to Elite University, which is known for its unsurpassed resources, the brilliance of its teachers, and the top jobs awarded to its graduates. Suppose also that admission to Elite depends almost entirely on a student's performance in a standardized test. Since there are a limited number of openings at Elite, what counts is not the score a student gets, but how well that score ranks in comparison to all the other scores. Getting 90 % of the answers right is no consolation if so many students do better than this that all the available slots at Elite will be taken by them.

Aware of the situation, you might hire a professional coach to help improve your score. If you are the only student who does this, your chances for admission will go up (if the coaching works). But others may have the same idea. If everyone hires a coach, and if everyone's score goes up by 5 %, this makes everyone look smarter, but the same people end up getting into Elite as before. The problem is that, since it is one's *position* in the rank-order of test-takers, and not the score itself, that matters for admission, each person's improvement comes at the expense of everyone else. If all improve, these effects cancel out, and no one has gained. The effects are externalities of particular sort—positional externalities.

This also explains the inefficiency, and in most cases the impossibility, of replacing the administrative organization of firms with lots of individual contracts between workers and owners. Each worker would try to bargain for what he or she wants: better working conditions, lighter or more interesting work assignments, a higher status within the enterprise. Many of these interests, however, come at the expense of other workers. If there is a limited number of good job assignments available, for instance, my getting one comes at the expense of you or someone else. These positional externalities would overwhelm any attempt at organizing production via one-on-one negotiations. Instead, firms have administrative structures that attempt, for better or worse, to take all these interconnections into consideration and allocate tasks and other job-related matters in a systematic manner.

When positional externalities are allowed to proliferate without restraint, the result is often an *arms race*. All participants invest more resources in getting ahead, but the outcome is only that the total amount of investment has gone up. What is rational from each individual's perspective is socially wasteful, much as universal defection (which it resembles) was seen to be an irrational outcome in the Prisoner's Dilemma.

15.5 Taking Stock

From the beginning, this book has argued that it is a misunderstanding to see economics only as a hymn of praise to the wonders of free markets. There is a whiff of that in some economic writings, since economists often feel that the general public fails to see the positive aspects of a well-functioning market system. Yet one could also say that economics is centrally focused on the *failure* of real-world markets, the defects that prevent them from achieving the nirvana promised

by the Market Welfare Model. This provides the justification for most of the policies that economists analyze and debate.

But just as a scrupulous economist should not accept market outcomes without carefully inspecting the process for public goods, externalities and other distortions, neither should she throw up her hands if she finds that market failure has occurred. Economic theory has classified and dissected these failures, and it has much to say about how, and even whether, to remedy them.

Defining a problem as precisely as possible is the biggest step toward identifying a solution. The theories outlined in this chapter provide a starting point for much of the work applied economists do in the policy arena. They estimate the value of public goods that fail to be produced due to free-rider problems and the difference between private and social benefits or costs in the presence of externalities. *Economics does its best work when markets need a helping hand.* The deep understanding of market failures offered by economics would not have been possible without a vision of ideal markets to compare them to. Logically, this chapter should be seen as a culmination, not a revision, of Chap. 6 on the Market Welfare Model.

At the same time, however, the problems that now shadow the Market Welfare Model should also be seen as casting doubt on the market failure framework. As we have learned in earlier chapters, the traditional theories of utility and rational choice which underpin the Market Welfare Model have been buffeted by new findings in the field of economic psychology. Individuals do not always make the choices that maximize their well-being, and happiness, as measured in surveys and physiological responses, does not correspond very well to utility as this has been defined by economists. This means that we do not necessarily need a theory of market failure to identify shortcomings in the way markets work: misguided decision-makers or simply the biases of market incentives that may impinge on happiness could lead to the same conclusion. Even more disturbing, it is entirely possible that correcting a market failure could lead to even *worse* outcomes as measured by happiness or capabilities. For instance, a negative externality in production like pollution could lead to less food crops being harvested, But what if our food consumption choices do not improve our health or happiness? Counteracting the externality, and thereby increasing the supply and reducing the price of "junk food" ingredients, for instance, might actually make us worse off.

Taken on their own terms, the criticisms emanating from economic psychology and happiness studies are deeply at odds with the perspective of both the Market Welfare Model *and* the market failure theory sketched in this chapter. Some economists would argue that utility theory should be rejected as faulty and anachronistic. The majority, however, continue to use utility theory and base much of their policy advice on market failure reasoning, even though they are aware that the basis for this approach has been called into question. They are guided by the faith that traditional economic ideas remain approximately right, in spite of their demonstrated shortcomings. At the same time it would not be an exaggeration to say that the current situation remains fluid, and that economics is in the midst of some sort of transition on these issues. Where this transition will take us, and how much weight will be given to notions like market failure in the years to come, it is too soon to tell.

The Main Points

1. Public goods are those which have one of two characteristics: either there is a zero or near-zero marginal cost of provision (nonrivalry) or it is impractical to prevent access to them if no payment is made (nonexclusion). Some public goods are provided by the public sector, but not all. Many goods provided by the public sector are *not* public goods in the economic sense.
2. The problem with nonexclusion is that there is an incentive for users to not pay; this is referred to as the free-rider problem. The problem with nonrivalry, on the other hand, is the possibility that there will be too little use of the good. If the marginal cost is zero, any price that may be charged will discourage some users whose personal benefit does not justify paying for it. Yet, as long as they derive any benefit, it is economically inefficient to exclude them.
3. Externalities arise as a result of missing markets. If an activity produces beneficial goods or services that users need not pay for (due to a missing market), there is a positive externality and a tendency for underprovision of that benefit. If an activity produces harmful outcomes that suppliers do not need to pay for (due to a missing market), there is a negative externality and a tendency for overprovision of that harm.
4. There are several potential remedies for externalities. One is moral suasion—persuading individuals to produce more goods with positive externalities or fewer with negative. Another is direct regulation by the government, such as laws prohibiting or limiting certain forms of pollution. A third is the use of monetary incentives—taxes and subsidies—to induce more provision of goods with positive externalities and less of goods with negative externalities. Finally, it may be possible to create the missing market by establishing new property rights or otherwise encouraging bargaining between creators and recipients of externalities.
5. The so-called Coase Theorem (which Ronald Coase himself did not propose) states that the assignment of property rights, whether rights belong to the creator or the recipient of an externality, does not alter the amount the externality that will be arrived at through bargaining. This hypothesis is false, however, due to the impact of the assignment of rights on bargaining power, its effect on entry or exit of those who create the goods being bargained over, and the psychological tendency referred to as status quo bias.
6. New types of externalities are attracting the interest of economists. One is network externalities, the effect that one person's choice of a good has on others whose benefit from that good depends on how many people use it. Another is positional externalities, where individuals invest in goods that increase their rank or place in a queue at the expense of others in the same ranking or queuing order.

► Terms to Define

Cap-and-trade
Externalities
Free-rider problem

Market failure
Negative externalities
Network externalities
Nonexclusion
Nonrivalry
Pigovian taxes and subsidies
Positional externalities
Positive externalities
Public goods
Pure public good
Social vs private costs (or benefits)
Spillovers
Zero marginal cost problem

Questions to Consider

1. Older economics textbooks sometimes mentioned streets and roads as examples of nonexcludable goods: while a few toll roads might charge for access, most didn't, and it did not seem remotely practical to set up toll booths at every intersection. This meant that road-building and maintenance would have to be financed out of taxes rather than user payments. Now, however, the technology exists to charge drivers for every stretch of road they drive on, even tailored to the time of day or season of the year. Each license plate can have a transmitter that sends signals to receptors placed by the side of the road, so that the number of times the vehicle passed (and when it passed) can be saved and stored. Periodically the driver can be charged for the roads he or she used. What do you think of this system? Would you be in favor of dropping tax support for roads and replacing it with user fees along these lines? Should the government subsidize research in new technologies that have the potential to keep track of each person's use of other goods, like parks, pedestrian walkways, etc.?
2. As we saw in this chapter, the potential problem of zero marginal cost has been solved by broadcasting companies that sell advertising based on the number of viewers or listeners. This same approach now shows up on the internet, where free content, like on-line newspapers or search results, is financed by ads. Is this a satisfactory solution to the nonrivalry characteristic of digital information? What are the advantages and disadvantages of relying on ad revenue? Can you think of any alternatives?
3. Airplanes produce a lot of noise when they take off and land, and this is a problem for people who live near airports. Is it an externality? Explain, using the definition provided in this chapter. Should there be financial incentives to reduce this noise? If so, should they be based on an assignment of "quiet rights" to nearby residents or "noise rights" to airports and airplanes? What difference would the choice of rights make for the amount of the incentive, the level of noise, the location of airports and the location of residential communities? If you don't favor taxes or subsidies, what do you propose—some other policy to control noise or no policy at all?

4. It could be argued that one of the most important examples of a network externality is the adoption of a language. Explain why the number of people who learn and speak various languages is influenced by network externalities. Is this a problem or a solution to a problem or both? Should something be done for languages that might decline as global communications become more integrated?
5. Some consumption items are regarded as "status goods": their main value is to confer social status on those who own or use them. Can you think of any examples that fit this description? If so, does the problem of positional externalities apply? Is there an "arms race"? Should some corrective action be taken?