CHAPTER NINETEEN

Externalities and Public Goods

In Chapter 13 we looked briefly at a few problems that may interfere with the allocational efficiency of perfectly competitive markets. Here we will examine two of those problems—externalities and public goods—in more detail. This examination has two purposes. First, we wish to show clearly why the existence of externalities and public goods may distort the allocation of resources. In so doing it will be possible to illustrate some additional features of the type of information provided by competitive prices and some of the circumstances that may diminish the usefulness of that information. Our second reason for looking more closely at externalities and public goods is to suggest ways in which the allocational problems they pose might be mitigated. We will see that, at least in some cases, the efficiency of competitive market outcomes may be more robust than might have been anticipated.

DEFINING EXTERNALITIES

Externalities occur because economic actors have effects on third parties that are not reflected in market transactions. Chemical makers spewing toxic fumes on their neighbors, jet planes waking up people, and motorists littering the highway are, from an economic point of view, all engaging in the same sort of activity: they are having a direct effect on the well-being of others that is outside market channels. Such activities might be contrasted to the direct effects of markets. When I choose to purchase a loaf of bread, for example, I (perhaps imperceptibly) increase the price of bread generally, and that may affect the well-being of other bread buyers. But such effects, because they are reflected in market prices, are not externalities and do not affect the market's ability to allocate resources efficiently.[1] Rather, the increase in the price of bread that results from my increased purchase is an accurate reflection of societal preferences, and the price increase helps ensure that the right mix of products is produced. That is not the case for toxic chemical discharges, jet noise, or litter. In these cases, market prices (of chemicals, air travel, or disposable containers) may not accurately reflect actual social costs because they may take no account of the damage being done to third parties. Information being conveyed by market prices is fundamentally inaccurate, leading to a misallocation of resources.

As a summary, then, we have developed the following definition.

[1]Sometimes effects of one economic agent on another that take place through the market system are termed *pecuniary* externalities to differentiate such effects from the *technological* externalities we are discussing. Here the use of the term *externalities* will refer only to the latter type, because these are the only type with consequences for the efficiency of resource allocation by competitive markets.

DEFINITION

Externality. An *externality* occurs whenever the activities of one economic actor affect the activities of another in ways that are not reflected in market transactions.

Before analyzing in detail why failing to take externalities into account can lead to a misallocation of resources, we will examine a few examples that should clarify the nature of the problem.

Interfirm externalities

To illustrate the externality issue in its simplest form, we consider two firms: one producing good x and the other producing good y. The production of good x is said to have an external effect on the production of y if the output of y depends not only on the inputs chosen by the y-entrepreneur but also on the level at which the production of x is carried on. Notationally, the production function for good y can be written as

$$y = f(k, l; x), \tag{19.1}$$

where x appears to the right of the semicolon to show that it is an effect on production over which the y-entrepreneur has no control.[2] As an example, suppose the two firms are located on a river, with firm y being downstream from x. Suppose firm x pollutes the river in its productive process. Then the output of firm y may depend not only on the level of inputs it uses itself but also on the amount of pollutants flowing past its factory. The level of pollutants, in turn, is determined by the output of firm x. In the production function shown by Equation 19.1, the output of firm x would have a negative marginal physical productivity $\partial y/\partial x < 0$. Increases in x output would cause less y to be produced. In the next section we return to analyze this case more fully, since it is representative of most simple types of externalities.

Beneficial externalities

The relationship between two firms may be beneficial. Most examples of such positive externalities are rather bucolic in nature. Perhaps the most famous, proposed by J. Meade, involves two firms, one producing honey (raising bees) and the other producing apples.[3] Because the bees feed on apple blossoms, an increase in apple production will improve productivity in the honey industry. The beneficial effects of having well-fed bees are a positive externality to the beekeeper. In the notation of Equation 19.1, $\partial y/\partial x$ would now be positive. In the usual perfectly competitive case, the productive activities of one firm have no direct effect on those of other firms: $\partial y/\partial x = 0$.

Externalities in utility

Externalities also can occur if the activities of an economic actor directly affect an individual's utility. Most common examples of environmental externalities are of this type. From an economic perspective it makes little difference whether such effects are created by firms (in the form, say, of toxic chemicals or jet noise) or by other individuals (litter or, perhaps, the noise from a loud radio). In all such cases the amount of such activities would enter directly into the individual's utility function in much the same way as firm x's output entered into firm y's production function in Equation 19.1. As in the case of firms, such externalities may sometimes be beneficial (you may actually like the song being played on your neighbor's radio). So, again, a situation of zero externalities can be

[2] We will find it necessary to redefine the assumption of "no control" considerably as the analysis of this chapter proceeds.

[3] J. Meade, "External Economies and Diseconomies in a Competitive Situation," *Economic Journal* 62 (March 1952): 54–67.

regarded as the middle ground in which other agents' activities have no direct effect on individuals' utilities.

One special type of utility externality that is relevant to the analysis of social choices arises when one individual's utility depends directly on the utility of someone else. If, for example, Smith cares about Jones's welfare, then we could write his or her utility function (U_S) as

$$\text{utility} = U_S(x_1, \ldots, x_n;\ U_J), \tag{19.2}$$

where $x_1, \ldots, x_n$ are the goods that Smith consumes and U_J is Jones's utility. If Smith is altruistic and wants Jones to be well off (as might happen if Jones were a close relative), $\partial U_S/\partial U_J$ would be positive. If, on the other hand, Smith were envious of Jones, then it might be the case that $\partial U_S/\partial U_J$ would be negative; that is, improvements in Jones's utility make Smith worse off. The middle ground between altruism and envy would occur if Smith were indifferent to Jones's welfare ($\partial U_S/\partial U_J = 0$), and that is what we have usually assumed throughout this book (for a brief discussion, see the Extensions to Chapter 3).

Public goods externalities

Goods that are "public" or "collective" in nature will be the focus of our analysis in the second half of this chapter. The defining characteristic of these goods is nonexclusion; that is, once the goods are produced (either by the government or by some private entity), they provide benefits to an entire group—perhaps to everyone. It is technically impossible to restrict these benefits to the specific group of individuals who pay for them, so the benefits are available to all. As we mentioned in Chapter 13, national defense provides the traditional example. Once a defense system is established, all individuals in society are protected by it whether they wish to be or not and whether they pay for it or not. Choosing the right level of output for such a good can be a tricky process, because market signals will be inaccurate.

EXTERNALITIES AND ALLOCATIVE INEFFICIENCY

Externalities lead to inefficient allocations of resources because market prices do not accurately reflect the additional costs imposed on or benefits provided to third parties. To illustrate these inefficiencies requires a general equilibrium model, because inefficient allocations in one market throw into doubt the efficiency of market-determined outcomes everywhere. Here we choose a very simple and, in some ways, rather odd general equilibrium model that allows us to make these points in a compact way. Specifically, we assume there is only one person in our simple economy and that his or her utility depends on the quantities of x and y consumed. Consumption levels of these two goods are denoted by x_c and y_c, so

$$\text{utility} = U(x_c, y_c). \tag{19.3}$$

This person has initial stocks of x and y (denoted by x^* and y^*) and can either consume these directly or use them as intermediary goods in production. To simplify matters, we assume that good x is produced using only good y, according to the production function

$$x_o = f(y_i), \tag{19.4}$$

where subscript o refers to outputs and i to inputs. To illustrate externalities, we assume that the output of good y depends not only on how much x is used as an input in the production process but also on the x production level itself. Hence this would model a situation, say, where y is downriver from firm x and must cope with the pollution created by production of x output. The production function for y is given by

$$y_o = g(x_i, x_o), \tag{19.5}$$

where $g_1 > 0$ (more x input produces more y output), but $g_2 < 0$ (additional x output reduces y output because of the externality involved).

The quantities of each good in this economy are constrained by the initial stocks available and by the additional production that takes place:

$$x_c + x_i = x_o + x^*, \tag{19.6}$$

$$y_c + y_i = y_o + y^*. \tag{19.7}$$

Finding the efficient allocation

The economic problem for this society, then, is to maximize utility subject to the four constraints represented by Equations 19.4–19.7. To solve this problem we must introduce four Lagrange multipliers. The Lagrangian expression for this maximization problem is

$$\begin{aligned}\mathscr{L} = {} & U(x_c, y_c) + \lambda_1[f(y_i) - x_o] + \lambda_2[g(x_i, x_o) - y_o] \\ & + \lambda_3(x_c + x_i - x_o - x^*) + \lambda_4(y_c + y_i - y_o - y^*),\end{aligned} \tag{19.8}$$

and the six first-order conditions for a maximum are

$$\begin{aligned}
\partial\mathscr{L}/\partial x_c &= U_1 + \lambda_3 = 0, && [\text{i}]\\
\partial\mathscr{L}/\partial y_c &= U_2 + \lambda_4 = 0, && [\text{ii}]\\
\partial\mathscr{L}/\partial x_i &= \lambda_2 g_1 + \lambda_3 = 0, && [\text{iii}]\\
\partial\mathscr{L}/\partial y_i &= \lambda_1 f_y + \lambda_4 = 0, && [\text{iv}]\\
\partial\mathscr{L}/\partial x_o &= -\lambda_1 + \lambda_2 g_2 - \lambda_3 = 0, && [\text{v}]\\
\partial\mathscr{L}/\partial y_o &= -\lambda_2 - \lambda_4 = 0. && [\text{vi}]
\end{aligned} \tag{19.9}$$

Eliminating the λs from these equations is a straightforward process. Taking the ratio of Equations i and ii yields the familiar result

$$MRS = \frac{U_1}{U_2} = \frac{\lambda_3}{\lambda_4}. \tag{19.10}$$

But Equations iii and vi also imply

$$MRS = \frac{\lambda_3}{\lambda_4} = \frac{\lambda_2 g_1}{\lambda_2} = g_1. \tag{19.11}$$

Hence optimality in y production requires that the individual's *MRS* in consumption equal the marginal productivity of x in the production *of* y. This conclusion repeats the result from Chapter 13, where we showed that efficient output choice requires that dy/dx in consumption be equal to dy/dx in production.

To achieve efficiency in x production, we must also consider the externality that this production poses to y. Combining Equations iv–vi gives

$$\begin{aligned}MRS = \frac{\lambda_3}{\lambda_4} &= \frac{-\lambda_1 + \lambda_2 g_2}{\lambda_4} = \frac{-\lambda_1}{\lambda_4} + \frac{\lambda_2 g_2}{\lambda_4} \\ &= \frac{1}{f_y} - g_2.\end{aligned} \tag{19.12}$$

Intuitively, this equation requires that the individual's MRS must also equal dy/dx obtained through x production. The first term in the expression, $1/f_y$, represents the reciprocal of the marginal productivity of y in x production—this is the first component of dy/dx as it relates to x production. The second term, g_2, represents the negative impact that added x production has on y output—this is the second component of dy/dx as it relates to x production. This final term occurs because of the need to consider the externality from x production. If g_2 were zero, then Equations 19.11 and 19.12 would represent essentially the same condition for efficient production, which would apply to both x and y. With the externality, however, determining an efficient level of x production is more complex.

Inefficiency of the competitive allocation

Reliance on competitive pricing in this simple model will result in an inefficient allocation of resources. With equilibrium prices p_x and p_y, a utility-maximizing individual would opt for

$$MRS = p_x/p_y \tag{19.13}$$

and the profit-maximizing producer of good y would choose x input according to

$$p_x = p_y g_1. \tag{19.14}$$

Hence the efficiency condition (Equation 19.11) would be satisfied. But the producer of good x would choose y input so that

$$p_y = p_x f_y \quad \text{or} \quad \frac{p_x}{p_y} = \frac{1}{f_y}. \tag{19.15}$$

That is, the producer of x would disregard the externality that its production poses for y and so the other efficiency condition (Equation 19.12) would not be met. This failure results in an overproduction of x relative to the efficient level. To see this, note that the marginal product of y in producing x (f_y) is smaller under the market allocation represented by Equation 19.15 than under the optimal allocation represented by Equation 19.12. More y is used to produce x in the market allocation (and hence more x is produced) than is optimal. Example 19.1 provides a quantitative example of this nonoptimality in a partial equilibrium context.

EXAMPLE 19.1 Production Externalities

As a partial equilibrium illustration of the losses from failure to consider production externalities, suppose two newsprint producers are located along a river. The upstream firm (x) has a production function of the form

$$x = 2{,}000 l_x^{1/2}, \tag{19.16}$$

where l_x is the number of workers hired per day and x is newsprint output in feet. The downstream firm (y) has a similar production function, but its output may be affected by the chemicals firm x pours into the river:

$$y = \begin{cases} 2{,}000 l_y^{1/2}(x - x_0)^{\alpha} & \text{for } x > x_0, \\ 2{,}000 l_y^{1/2} & \text{for } x \le x_0, \end{cases} \tag{19.17}$$

where x_0 represents the river's natural capacity for neutralizing pollutants. If $\alpha = 0$, then x's production process has no effect on firm y, but if $\alpha < 0$, an increase in x above x_0 causes y's output to decrease.

Assuming newsprint sells for $1 per foot and workers earn $50 per day, firm x will maximize profits by setting this wage equal to labor's marginal revenue product:

$$50 = p \cdot \frac{\partial x}{\partial l_x} = 1{,}000 l_x^{-1/2}. \tag{19.18}$$

The solution then is $l_x = 400$. If $\alpha = 0$ (there are no externalities), firm y will also hire 400 workers. Each firm will produce 40,000 feet of newsprint.

Effects of an externality. When firm x does have a negative externality ($\alpha < 0$), its profit-maximizing hiring decision is not affected—it will still hire $l_x = 400$ and produce $x = 40{,}000$. But for firm y, labor's marginal product will be lower because of this externality. If $\alpha = -0.1$ and $x_0 = 38{,}000$, for example, then profit maximization will require

$$\begin{aligned} 50 = p \cdot \frac{\partial y}{\partial l_y} &= 1{,}000 l_y^{-1/2}(x - 38{,}000)^{-0.1} \\ &= 1{,}000 l_y^{-1/2}(2{,}000)^{-0.1} \\ &= 468 l_y^{-1/2}. \end{aligned} \tag{19.19}$$

Solving this equation for l_y shows that firm y now hires only 87 workers because of this lowered productivity. Output of firm y will now be

$$y = 2{,}000(87)^{1/2}(2{,}000)^{-0.1} = 8{,}723. \tag{19.20}$$

Because of the externality ($\alpha = -0.1$), newsprint output will be lower than without the externality ($\alpha = 0$).

Inefficiency. We can demonstrate that decentralized profit maximization is inefficient in this situation by imagining that firms x and y merge and that the manager must decide how to allocate the combined workforce. If one worker is transferred from firm x to firm y, then x output becomes

$$\begin{aligned} x &= 2{,}000(399)^{1/2} \\ &= 39{,}950; \end{aligned} \tag{19.21}$$

for firm y,

$$\begin{aligned} y &= 2{,}000(88)^{1/2}(1{,}950)^{-0.1} \\ &= 8{,}796. \end{aligned} \tag{19.22}$$

Total output has increased by 23 feet of newsprint with no change in total labor input. The market-based allocation was inefficient because firm x did not take into account the negative effect of its hiring decisions on firm y.

Marginal productivity. This can be illustrated in another way by computing the true social marginal productivity of labor input to firm x. If that firm were to hire one more worker, its own output would increase to

$$x = 2{,}000(401)^{1/2} = 40{,}050. \tag{19.23}$$

As profit maximization requires, the (private) marginal value product of the 401st worker is equal to the wage. But increasing x's output now also has an effect on firm y—its output decreases by about 21 units. Hence the social marginal revenue product of labor to firm x actually amounts to only $29 ($50 − $21). That is why the manager of a merged firm would find it profitable to shift some workers from firm x to firm y.

QUERY: Suppose $\alpha = +0.1$. What would that imply about the relationship between the firms? How would such an externality affect the allocation of labor?

SOLUTIONS TO THE EXTERNALITY PROBLEM

Incentive-based solutions to the allocational harm of externalities start from the basic observation that output of the externality-producing activity is too high under a market-determined equilibrium. Perhaps the first economist to provide a complete analysis of this distortion was A. C. Pigou, who in the 1920s suggested that the most direct solution would simply be to tax the externality-creating entity.[4] All incentive-based solutions to the externality problem stem from this basic insight.[5]

A graphic analysis

Figure 19.1 provides the traditional illustration of an externality together with Pigou's taxation solution. The competitive supply curve for good x also represents that good's private marginal costs of production (MC). When the demand for x is given by DD, the market equilibrium will occur at x_1. The external costs involved in x production create a divergence between private marginal costs (MC) and overall social marginal costs (MC')—the vertical distance between the two curves represents the costs that x

FIGURE 19.1

Graphic Analysis of an Externality

The demand curve for good x is given by DD. The supply curve for x represents the private marginal costs (MC) involved in x production. If x production imposes external costs on third parties, social marginal costs (MC') will exceed MC by the extent of these costs. Market equilibrium occurs at x_1 and, at this output level, social marginal costs exceed what consumers pay for good x. A tax of amount t that reflects the costs of the externalities would achieve the efficient output of x—given by output level x_2.

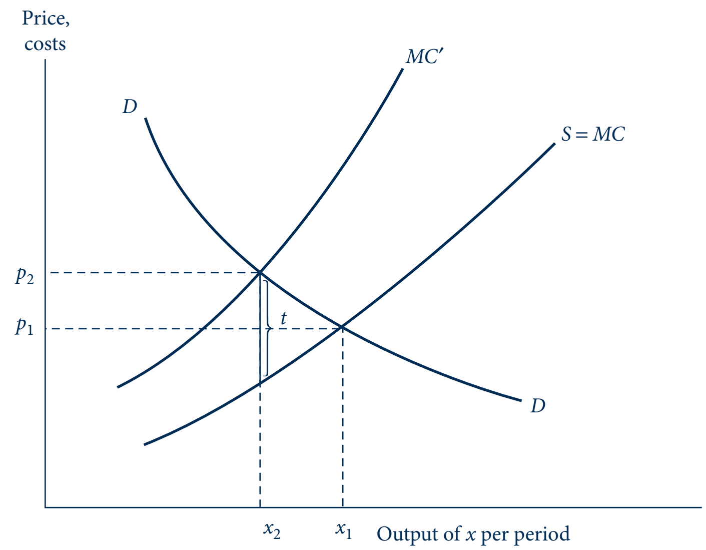

[4] A. C. Pigou, *The Economics of Welfare* (London: MacMillan, 1920). Pigou also recognized the importance of subsidizing goods that yield positive externalities.

[5] We do not discuss purely regulatory solutions here, although the study of such solutions forms an important part of most courses in environmental economics. See W. J. Baumol and W. E. Oates, *The Theory of Environmental Policy*, 2nd ed. (Cambridge: Cambridge University Press, 2005) and the Extensions to this chapter.

production poses for third parties (in our examples, only on firm y). Notice that the per-unit costs of these externalities need not be constant, independent of x output. In the figure, for example, the size of these external costs increases as x output expands (i.e., MC' and MC become further apart). At the market-determined output level x_1, the comprehensive social marginal cost exceeds the market price p_1, thereby indicating that the production of x has been pushed "too far." It is clear from the figure that the optimal output level is x_2, at which the market price p_2 paid for the good now reflects all costs.

As is the case for any tax, imposition of a Pigovian tax would create a vertical wedge between the demand and supply curves for good x. In Figure 19.1 this optimal tax is shown as t. Imposition of this tax serves to reduce output to x_2, the social optimum. Tax collections equal the precise amount of external harm that x production causes. These collections might be used to compensate firm y for these costs, but this is not crucial to the analysis. Notice here that the tax must be set at the level of harm prevailing at the optimum (i.e., at x_2), not at the level of harm at the original market equilibrium (x_1). This point is also made in the next example and more completely in the next section by returning to our simple general equilibrium model.

EXAMPLE 19.2 A Pigovian Tax on Newsprint

The inefficiency in Example 19.1 arises because the upstream newsprint producer (firm x) takes no account of the effect that its production has on firm y. A suitably chosen tax on firm x can cause it to reduce its hiring to a level at which the externality vanishes. Because the river can absorb the pollutants generated with an output of $x = 38{,}000$, we might consider imposing a tax (t) on the firm's output that encourages it to reduce output to this level. Because output will be 38,000 if $l_x = 361$, we can calculate t from the labor demand condition:

$$(1 - t)MP_L = (1 - t)1{,}000(361)^{-0.5} = 50, \tag{19.24}$$

or

$$t = 0.05. \tag{19.25}$$

Such a 5 percent tax would effectively reduce the price firm x receives for its newsprint to \$0.95 and provide it with an incentive to reduce its hiring by 39 workers. Now, because the river can handle all the pollutants that x produces, there is no externality in the production function of firm y. It will hire 400 workers and produce 40,000 feet of newsprint per day. Observe that total newsprint output is now 78,000, a significantly higher figure than would be produced in the untaxed situation. The taxation solution provides a considerable improvement in the efficiency of resource allocation.

QUERY: The tax rate proposed here (0.05) seems rather small given the significant output gains obtained relative to the situation in Example 19.1. Can you explain why? Would a merged firm opt for $x = 38{,}000$ even without a tax?

Taxation in the general equilibrium model

The optimal Pigovian tax in our general equilibrium model is to set $t = -p_y g_2$. That is, the per-unit tax on good x should reflect the marginal harm that x does in reducing y output, valued at the market price of good y. Notice again that this tax must be based on the value of this externality at the optimal solution; because g_2 will generally be a function of the level of x output, a tax based on some other output level would be inappropriate.

With the optimal tax, firm x now faces a net price for its output of $p_x - t$ and will choose y input according to

$$p_y = (p_x - t)f_y. \tag{19.26}$$

Hence the resulting allocation of resources will achieve

$$MRS = \frac{p_x}{p_y} = \frac{1}{f_y} + \frac{t}{p_y} = \frac{1}{f_y} - g_2, \tag{19.27}$$

which is precisely what is required for optimality (compare to the efficiency condition, Equation 19.12). The Pigovian taxation solution can be generalized in a variety of ways that provide insights about the conduct of policy toward externalities. For example, in an economy with many x-producers, the tax would convey information about the marginal impact that output from any one of these would have on y output. Hence the tax scheme mitigates the need for regulatory attention to the specifics of any particular firm. It does require that regulators have enough information to set taxes appropriately—that is, they must know firm y's production function.

Pollution rights

An innovation that would mitigate the informational requirements involved with Pigovian taxation is the creation of a market for "pollution rights." Suppose, for example, that firm x must purchase from firm y rights to pollute the river they share. In this case, x's decision to purchase these rights is identical to its decision to choose its output level, because it cannot produce without them. The net revenue x receives per unit is given by $p_x - r$, where r is the payment the firm must make for each unit it produces. Firm y must decide how many rights to sell to firm x. Because it will be paid r for each right, it must "choose" x output to maximize its profits:

$$\pi_y = p_y g(x_i, x_0) + r x_0; \tag{19.28}$$

the first-order condition for a maximum is

$$\frac{\partial \pi_y}{\partial x_0} = p_y g_2 + r = 0 \quad \text{or} \quad r = -p_y g_2. \tag{19.29}$$

Equation 19.29 makes clear that the equilibrium solution to pricing in the pollution rights market will be identical to the Pigovian tax equilibrium. From the point of view of firm x, it makes no difference whether a tax of amount t is paid to the government or a royalty r of the same amount is paid to firm y. So long as $t = r$ (a condition ensured by Equation 19.29), the same efficient equilibrium will result.

The Coase theorem

In a famous 1960 paper, Ronald Coase showed that the key feature of the pollution rights equilibrium is that these rights be well defined and tradable with zero transaction costs.[6] The initial assignment of rights is irrelevant because subsequent trading will always yield the same efficient equilibrium. In our example we initially assigned the rights to firm y, allowing that firm to trade them away to firm x for a per-unit fee r. If the rights had been assigned to firm x instead, that firm still would have to impute some cost to using these rights themselves rather than selling them to firm y. This calculation, in combination with firm y's decision about how many such rights to buy, will again yield an efficient result.

[6]R. Coase, "The Problem of Social Cost," *Journal of Law and Economics* 3 (October 1960): 1–44.

To illustrate the Coase result, assume that firm x is given x^T rights to produce (and to pollute). It can choose to use some of these to support its own production (x_0), or it may sell some to firm y (an amount given by $x^T - x_0$). Gross profits for x are given by

$$\pi_x = p_x x_0 + r(x^T - x_0) = (p_x - r)x_0 + rx^T = (p_x - r)f(y_i) + rx^T \qquad \textbf{(19.30)}$$

and for y by

$$\pi_y = p_y g(x_i, x_0) - r(x^T - x_0). \qquad \textbf{(19.31)}$$

Clearly, profit maximization in this situation will lead to precisely the same solution as in the case where firm y was assigned the rights. Because the overall total number of rights (x^T) is a constant, the first-order conditions for a maximum will be exactly the same in the two cases. This independence of initial rights assignment is usually referred to as the *Coase theorem*.

Although the results of the Coase theorem may seem counterintuitive (how can the level of pollution be independent of who initially owns the rights?), it is in reality nothing more than the assertion that, in the absence of impediments to making bargains, all mutually beneficial transactions will be completed. When transaction costs are high or when information is asymmetric, initial rights assignments *will* matter because the sorts of trading implied by the Coase theorem may not occur. Therefore, it is the limitations of the Coase theorem that provide the most interesting opportunities for further analysis. This analysis has been especially far reaching in the field of law and economics,[7] where the theorem has been applied to such topics as tort liability laws, contract law, and product safety legislation (see Problem 19.4).

ATTRIBUTES OF PUBLIC GOODS

We now turn our attention to a related set of problems about the relationship between competitive markets and the allocation of resources: those raised by the existence of public goods. We begin by providing a precise definition of this concept and then examine why such goods pose allocational problems. We then briefly discuss theoretical ways in which such problems might be mitigated before turning to examine how actual decisions on public goods are made through voting.

The most common definitions of public goods stress two attributes of such goods: nonexclusivity and nonrivalness. We now describe these attributes in detail.

Nonexclusivity

The first property that distinguishes public goods concerns whether individuals may be excluded from the benefits of consuming the good. For most private goods such exclusion is indeed possible: I can easily be excluded from consuming a hamburger if I don't pay for it. In some cases, however, such exclusion is either very costly or impossible. National defense is the standard example. Once a defense system is established, everyone in a country benefits from it whether they pay for it or not. Similar comments apply, on a more local level, to goods such as mosquito control or a program to inoculate against disease. In these cases, once the programs are implemented, no one in the community can be excluded from those benefits whether he or she pays for them or not. Hence we can divide goods into two categories according to the following definition.

[7]The classic text is R. A. Posner, *Economic Analysis of Law*, 4th ed. (Boston: Little, Brown, 1992). A more mathematical approach is T. J. Miceli, *Economics of the Law* (New York: Oxford University Press, 1997).

DEFINITION

Exclusive goods. A good is *exclusive* if it is relatively easy to exclude individuals from benefiting from the good once it is produced. A good is *nonexclusive* if it is impossible (or costly) to exclude individuals from benefiting from the good.

Nonrivalry

A second property that characterizes public goods is nonrivalry. A nonrival good is one for which additional units can be consumed at zero social marginal cost. For most goods, of course, consumption of additional amounts involves some marginal costs of production. Consumption of one more hot dog requires that various resources be devoted to its production. However, for certain goods this is not the case. Consider, for example, having one more automobile cross a highway bridge during an off-peak period. Because the bridge is already in place, having one more vehicle cross requires no additional resource use and does not reduce consumption elsewhere. Similarly, having one more viewer tune in to a television channel involves no additional cost, even though this action would result in additional consumption taking place. Therefore, we have developed the following definition.

DEFINITION

Nonrival goods. A good is *nonrival* if consumption of additional units of the good involves zero social marginal costs of production.

Typology of public goods

The concepts of nonexclusion and nonrivalry are in some ways related. Many nonexclusive goods are also nonrival. National defense and mosquito control are two examples of goods for which exclusion is not possible and additional consumption takes place at zero marginal cost. Many other instances might be suggested. The concepts, however, are not identical: some goods may possess one property but not the other. For example, it is impossible (or at least very costly) to exclude some fishing boats from ocean fisheries, yet the arrival of another boat clearly imposes social costs in the form of a reduced catch for all concerned. Similarly, use of a bridge during off-peak hours may be nonrival, but it is possible to exclude potential users by erecting toll booths. Table 19.1 presents a cross-classification of goods by their possibilities for exclusion and their rivalry. Several examples of goods that fit into each of the categories are provided. Many of the examples, other than those in the upper left corner of the table (exclusive and rival private goods), are often produced by governments. That is especially the case for nonexclusive goods because, as we shall see, it is difficult to develop ways of paying for such goods other than through compulsory taxation. Nonrival goods often are privately produced (there are, after all, private bridges, swimming pools, and highways that consumers must pay to use) as long as nonpayers can be excluded from consuming them.[8] Still, we will use the following stringent definition, which requires both conditions.

[8]Nonrival goods that permit imposition of an exclusion mechanism are sometimes referred to as *club goods*, because provision of such goods might be organized along the lines of private clubs. Such clubs might then charge a "membership" fee and permit unlimited use by members. The optimal size of a club is determined by the economies of scale present in the production process for the club good. For an analysis, see R. Cornes and T. Sandler, *The Theory of Externalities, Public Goods, and Club Goods* (Cambridge: Cambridge University Press, 1986).

TABLE 19.1 EXAMPLES SHOWING THE TYPOLOGY OF PUBLIC AND PRIVATE GOODS

		Exclusive	
		Yes	**No**
Rival	**Yes**	Hot dogs, automobiles, houses	Fishing grounds, public grazing land, clean air
	No	Bridges, swimming pools, satellite television transmission (scrambled)	National defense, mosquito control, justice

DEFINITION

Public good. A good is a (pure) *public good* if, once produced, no one can be excluded from benefiting from its availability and if the good is nonrival—the marginal cost of an additional consumer is zero.

PUBLIC GOODS AND RESOURCE ALLOCATION

To illustrate the allocational problems created by public goods, we again employ a simple general equilibrium model. In this model there are only two individuals—a single-person economy would not experience problems from public goods because he or she would incorporate all of the goods' benefits into consumption decisions. We denote these two individuals by A and B. There are also only two goods in this economy. Good y is an ordinary private good, and each person begins with an allocation of this good given by y^{A*} and y^{B*}, respectively. Each person may choose to consume some of his or her y directly or to devote some portion of it to the production of a single public good, x. The amounts contributed are given by y_s^A and y_s^B, and the public good is produced according to the production function

$$x = f(y_s^A + y_s^B). \qquad \textbf{(19.32)}$$

Resulting utilities for these two people in this society are given by

$$U^A(x, y^{A*} - y_s^A) \qquad \textbf{(19.33)}$$

and

$$U^B(x, y^{B*} - y_s^B) \qquad \textbf{(19.34)}$$

Notice here that the level of public good production, x, enters identically into each person's utility function. This is the way in which the nonexclusivity and nonrivalry characteristics of such goods are captured mathematically. Nonexclusivity is reflected by the fact that each person's consumption of x is the same and independent of what he or she contributes individually to its production. Nonrivalry is shown by the fact that the consumption of x by each person is identical to the total amount of x produced. Consumption of x benefits by A does not diminish what B can consume. These two characteristics of good x constitute the barriers to efficient production under most decentralized decision schemes, including competitive markets.

The necessary conditions for efficient resource allocation in this problem consist of choosing the levels of public goods subscriptions (y_s^A and y_s^B) that maximize, say, A's utility for any given level of B's utility. The Lagrangian expression for this problem is

$$\mathscr{L} = U^A(x, y^{A*} - y_s^A) + \lambda[U^B(x, y^{B*} - y_s^B) - K], \qquad \textbf{(19.35)}$$

where K is a constant level of B's utility. The first-order conditions for a maximum are

$$\frac{\partial \mathcal{L}}{\partial y_s^A} = U_1^A f' - U_2^A + \lambda U_1^B f' = 0, \tag{19.36}$$

$$\frac{\partial \mathcal{L}}{\partial y_s^B} = U_1^A f' - \lambda U_2^B + \lambda U_1^B f' = 0. \tag{19.37}$$

A comparison of these two equations yields the immediate result that

$$\lambda U_2^B = U_2^A. \tag{19.38}$$

As might have been expected here, optimality requires that the marginal utility of y consumption for A and B be equal except for the constant of proportionality, λ. This equation may now be combined with either Equation 19.36 or 19.37 to derive the optimality condition for producing the public good x. Using Equation 19.36, for example, gives

$$\frac{U_1^A}{U_2^A} + \frac{\lambda U_1^B}{\lambda U_2^B} = \frac{1}{f'} \tag{19.39}$$

or, more simply,

$$MRS^A + MRS^B = \frac{1}{f'}. \tag{19.40}$$

The intuition behind this condition, which was first articulated by P. A. Samuelson,[9] is that it is an adaptation of the efficiency conditions described in Chapter 13 to the case of public goods. For such goods, the MRS in consumption must reflect the amount of y that *all* consumers would be willing to give up to get one more x, because everyone will obtain the benefits of the extra x output. Hence it is the sum of each individual's MRS that should be equated to dy/dx in production (here given by $1/f'$).

Failure of a competitive market

Production of goods x and y in competitive markets will fail to achieve this allocational goal. With perfectly competitive prices p_x and p_y, each individual will equate his or her MRS to the price ratio p_x/p_y. A producer of good x would also set $1/f'$ to be equal to p_x/p_y, as would be required for profit maximization. This behavior would not achieve the optimality condition expressed in Equation 19.40. The price ratio p_x/p_y would be "too low" in that it would provide too little incentive to produce good x. In the private market, a consumer takes no account of how his or her spending on the public good benefits others, so that consumer will devote too few resources to such production.

The allocational failure in this situation can be ascribed to the way in which private markets sum individual demands. For any given quantity, the market demand curve reports the marginal valuation of a good. If one more unit were produced, it could then be consumed by someone who would value it at this market price. For public goods, the value of producing one more unit is in fact the sum of each consumer's valuation of that extra output, because all consumers will benefit from it. In this case, then, individual demand curves should be added vertically (as shown in Figure 19.2) rather than horizontally (as they are in competitive markets). The resulting price on such a public good

[9]P. A. Samuelson, "The Pure Theory of Public Expenditure," *Review of Economics and Statistics* (November 1954): 387–89.

FIGURE 19.2

Derivation of the Demand for a Public Good

For a public good, the price individuals are willing to pay for one more unit (their "marginal valuations") is equal to the sum of what each individual would pay. Hence, for public goods, the demand curve must be derived by a vertical summation rather than the horizontal summation used in the case of private goods.

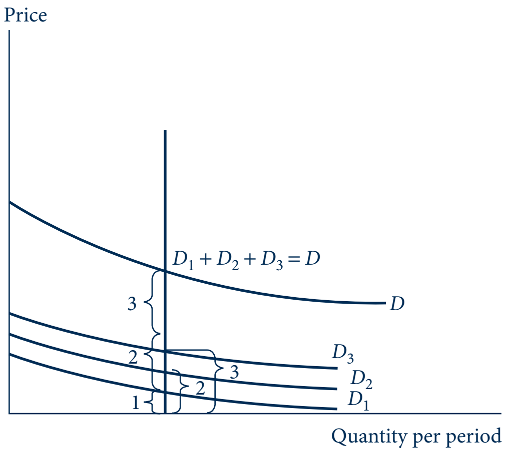

demand curve will then reflect, for any level of output, how much an extra unit of output would be valued by all consumers. But the usual market demand curve will not properly reflect this full marginal valuation.

Inefficiency of a Nash equilibrium

An alternative approach to the production of public goods in competitive markets might rely on individuals' voluntary contributions. Unfortunately, this also will yield inefficient results. Consider the situation of person A, who is thinking about contributing s_A of his or her initial y endowment to public goods production. The utility maximization problem for A is then

$$\text{choose } s_A \text{ to maximize } U^A[f(s_A + s_B), y^{A*} - s_A]. \tag{19.41}$$

The first-order condition for a maximum is

$$U_1^A f' - U_2^A = 0 \quad \text{or} \quad \frac{U_1^A}{U_2^A} = MRS^A = \frac{1}{f'}. \tag{19.42}$$

Because a similar logic will apply to person B, the efficiency condition of Equation 19.40 will once more fail to be satisfied. Again the problem is that each person considers only his or her benefit from investing in the public good, taking no account of the benefits provided to others. With many consumers, this direct benefit may be very small indeed. (For example, how much do one person's taxes contribute to national defense in the United States?) In this case, any one person may opt for $s_A = 0$ and become a pure "free rider," hoping to benefit from the expenditures of others. If every person adopts this strategy, then no resources will be subscribed to public goods. Example 19.3 illustrates the free-rider problem in a situation that may be all too familiar.

EXAMPLE 19.3 Purchasing a Public Good: The Roommates' Dilemma

To illustrate the nature of the public good problem numerically, suppose two roommates with identical preferences derive utility from the number of music compact disks (CDs, denoted by x) in their shared music collection and on the number of granola bars (y) eaten. The specific utility function for $i = 1, 2$ is given by

$$U_i(x, y_i) = x^{1/2} y_i^{1/2}. \qquad \textbf{(19.43)}$$

Utility for each roommate depends on the total number of CDs ($x = x_1 + x_2$) in their collection but only on the number of granola bars eaten by the individual. Hence in this problem a CD is a public good and a granola bar is a private good. (We could justify the classification of CDs as a public good by assuming that the purchaser of the CD cannot exclude his or her roommate from borrowing and playing it on their shared sound system. Playing the CD once does not diminish its value when played again, so there is nonrivalry in CD consumption.) Assume each roommate has \$300 to spend and that $p_x = \$10$ and $p_y = \$1$.

Nash equilibrium. We first consider the outcome if the roommates make their consumption decisions independently without coming to a more or less formal agreement about how many CDs to buy. Roommate 1's decision depends on how many CDs roommate 2 buys and vice versa. We are in a strategic situation for which we need the tools of game theory from Chapter 8 to analyze. We will look for the Nash equilibrium, in which both roommates are playing a best response.

To find roommate 1's best response, take as given the number x_2 of CDs purchased by roommate 2. Roommate 1 maximizes utility $(x_1 + x_2)^{1/2} y_i^{1/2}$ subject to the budget constraint

$$300 = 10x_1 + y_1, \qquad \textbf{(19.44)}$$

leading to the Lagrangian

$$\mathscr{L} = (x_1 + x_2)^{1/2} y_i^{1/2} + \lambda(300 - 10x_1 - y_1). \qquad \textbf{(19.45)}$$

The first-order conditions with respect to roommate 1 choice variables are

$$\begin{aligned} \frac{\partial \mathscr{L}}{\partial x_1} &= \frac{1}{2}(x_1 + x_2)^{-1/2} y_i^{1/2} - 10\lambda = 0 \\ \frac{\partial \mathscr{L}}{\partial y_1} &= \frac{1}{2}(x_1 + x_2)^{1/2} y_i^{-1/2} - \lambda = 0. \end{aligned} \qquad \textbf{(19.46)}$$

Solving Equations 19.46 in the usual way gives

$$y_1 = 10(x_1 + x_2), \qquad \textbf{(19.47)}$$

which, when substituted into 1's budget constraint and rearranged, gives the best-response function

$$x_1 = 15 - \frac{x_2}{2}. \qquad \textbf{(19.48)}$$

Because the problem is symmetric, roommate 2's best-response function will have the same form:

$$x_2 = 15 - \frac{x_1}{2}. \qquad \textbf{(19.49)}$$

These best-response functions reflect a free-rider problem in that the more CDs one roommate is expected to purchase, the fewer CDs the other wants to buy.

Solving Equations 19.48 and 19.49 simultaneously gives $x_1^* = x_2^* = 10$, and substituting this into Equation 19.47 gives $y_1^* = y_2^* = 200$. Nash equilibrium utilities are $U_1^* = U_2^* \approx 63.2$.

Efficient allocation. We saw that the efficient level of a public good can be calculated by setting the sum of each person's *MRS* equal to the good's price ratio. In this example, the *MRS* for roommate i is

$$MRS_i = \frac{\partial U_i/\partial x}{\partial U_i/\partial y_i} = \frac{y_i}{x}. \tag{19.50}$$

Hence the condition for efficiency is

$$MRS_1 + MRS_2 = \frac{y_1}{x} + \frac{y_2}{x} = \frac{p_x}{p_y} = \frac{10}{1}. \tag{19.51}$$

Consequently,

$$y_1 + y_2 = 10x, \tag{19.52}$$

which can be substituted into the combined budget constraint

$$600 = 10x + y_1 + y_2 \tag{19.53}$$

to obtain $x^{**} = 30$ and $y_1^{**} + y_2^{**} = 300$ (double stars distinguish efficient values from the Nash equilibrium ones with single stars). Assuming each roommate eats half (150) of the granola bars, the resulting utilities are $U_1^{**} = U_2^{**} \approx 67.1$.

Comparison. In the Nash equilibrium, too little of the public good (CDs) is purchased. The most efficient outcome has them purchasing five more CDs than they would on their own. It might be possible for them to come to a formal or informal agreement to buy more CDs, perhaps putting money in a pool and purchasing them together; the utility of both could simultaneously be increased this way. In the absence of such an agreement, the roommates face a similar dilemma as the players in the Prisoners' Dilemma: the Nash equilibrium (both fink) is Pareto dominated by another outcome (their utility is higher if both are silent).

QUERY: Solve the problem for three roommates. In what sense has the public good problem become worse with more players? How would an increase in the number of roommates affect their ability to enforce a cooperative agreement to buy more CDs?

LINDAHL PRICING OF PUBLIC GOODS

An important conceptual solution to the public goods problem was first suggested by the Swedish economist Erik Lindahl[10] in the 1920s. Lindahl's basic insight was that individuals might voluntarily consent to be taxed for beneficial public goods if they knew that others were also being taxed. Specifically, Lindahl assumed that each individual would be presented by the government with the proportion of a public good's cost he or she would be expected to pay and then reply (honestly) with the level of public good output he or she would prefer. In the notation of our simple general equilibrium model, individual A would be quoted a specific percentage (α^A) and then asked the level of public goods that he or she would want given the knowledge that this fraction of total cost would have to be paid. To answer that question (truthfully), this person would choose that overall level of public goods output, x, that maximizes

$$\text{utility} = U^A[x, y^{A*} - \alpha^A f^{-1}(x)]. \tag{19.54}$$

[10]Excerpts from Lindahl's writings are contained in R. A. Musgrave and A. T. Peacock, Eds., *Classics in the Theory of Public Finance* (London: Macmillan, 1958).

The first-order condition for this utility-maximizing choice of x is given by

$$U_1^A - \alpha U_2^B\left(\frac{1}{f'}\right) = 0 \quad \text{or} \quad MRS^A = \frac{\alpha^A}{f'}. \tag{19.55}$$

Individual B, presented with a similar choice, would opt for a level of public goods satisfying

$$MRS^B = \frac{\alpha^B}{f'}. \tag{19.56}$$

An equilibrium would then occur where $\alpha^A + \alpha^B = 1$—that is, where the level of public goods expenditure favored by the two individuals precisely generates enough in tax contributions to pay for it. For in that case

$$MRS^A + MRS^B = \frac{\alpha^A + \alpha^B}{f'} = \frac{1}{f'}, \tag{19.57}$$

and this equilibrium would be efficient (see Equation 19.40). Hence, at least on a conceptual level, the Lindahl approach solves the public good problem. Presenting each person with the equilibrium tax share "price" will lead him or her to opt for the efficient level of public goods production.

EXAMPLE 19.4 Lindahl Solution for the Roommates

Lindahl pricing provides a conceptual solution to the roommates' problem of buying CDs in Example 19.3. If "the government" (or perhaps social convention) suggests that each roommate will pay half of CD purchases, then each would face an effective price of CDs of $5. Since the utility functions for the roommates imply that half of each person's total income of $300 will be spent on CDs, it follows that each will be willing to spend $150 on such music and will, if each is honest, report that he or she would like to have 15 CDs. Hence the solution will be $x^{**} = 30$ and $y_1^{**} = y_2^{**} = 150$. This is indeed the efficient solution calculated in Example 19.3.

This solution works if the government knows enough about the roommates' preferences that it can set the payment shares in advance and stick to them. Knowing that the roommates have symmetric preferences in this example, it could set equal payment shares $\alpha_1 = \alpha_2 = 1/2$, and rest assured that both will honestly report the same demands for the public good, $x^{**} = 30$. If, however, the government does not know their preferences, it would have to tweak the payment shares based on their reports to make sure the reported demands end up being equal as required for the Lindahl solution to be "in equilibrium." Anticipating the effect of their reports on their payment shares, the roommates would have an incentive to underreport demand. In fact, this underreporting would lead to the same outcome as in the Nash equilibrium from Example 19.3.

QUERY: Although the 50–50 sharing in this example might arise from social custom, in fact the optimality of such a split is a special feature of this problem. What is it about this problem that leads to such a Lindahl outcome? Under what conditions would Lindahl prices result in other than a 50–50 sharing?

Shortcomings of the Lindahl solution

Unfortunately, Lindahl's solution is only a conceptual one. We have already seen in our examination of the Nash equilibrium for public goods production and in our roommates' example that the incentive to be a free rider in the public goods case is very strong. This fact makes it difficult to envision how the information necessary to compute equilibrium

Lindahl shares might be obtained. Because individuals know their tax shares will be based on their reported demands for public goods, they have a clear incentive to understate their true preferences—in so doing they hope that the "other guy" will pay. Hence, simply asking people about their demands for public goods should not be expected to reveal their true demands. We will discuss more sophisticated mechanisms for eliciting honest demand reports at the end of the chapter.

Local public goods

Some economists believe that demand revelation for public goods may be more tractable at the local level.[11] Because there are many communities in which individuals might reside, they can indicate their preferences for public goods (i.e., for their willingness to pay Lindahl tax shares) by choosing where to live. If a particular tax burden is not utility maximizing then people can, in principle, "vote with their feet" and move to a community that does provide optimality. Hence, with perfect information, zero costs of mobility, and enough communities, the Lindahl solution may be implemented at the local level. Similar arguments apply to other types of organizations (such as private clubs) that provide public goods to their members; given a sufficiently wide spectrum of club offerings, an efficient equilibrium might result. Of course, the assumptions that underlie the purported efficiency of such choices by individuals are quite strict. Even minor relaxation of these assumptions may yield inefficient results owing to the fragile nature of the way in which the demand for public goods is revealed.

EXAMPLE 19.5 The Relationship between Environmental Externalities and Public Goods Production

In recent years, economists have begun to study the relationship between the two issues we have been discussing in this chapter: externalities and public goods. The basic insight from this examination is that one must take a general equilibrium view of these problems in order to identify solutions that are efficient overall. Here we illustrate this point by returning to the computable general equilibrium model firms described in Chapter 13 (see Example 13.4). To simplify matters we will now assume that this economy includes only a single representative person whose utility function is given by

$$\text{utility} = U(x, y, l, g, c) = x^{0.5}y^{0.3}l^{0.2}g^{0.1}c^{0.2}, \tag{19.58}$$

where we have added terms for the utility provided by public goods (g), which are initially financed by a tax on labor, and by clean air (c). Production of the public good requires capital and labor input according to the production function $g = k^{0.5}l^{0.5}$; there is an externality in the production of good y, so that the quantity of clean air is given by $c = 10 - 0.2y$. The production functions for goods x and y remain as described in Example 13.4, as do the endowments of k and l. Hence our goal is to allocate resources in such a way that utility is maximized.

Base case: Optimal public goods production with no Pigovian tax. If no attempt is made to control the externality in this problem, then the optimal level of public goods production requires $g = 2.93$ and this is financed by a tax rate of 0.25 on labor. Output of good y in this case is 29.7, and the quantity of clean air is given by $c = 10 - 5.94 = 4.06$. Overall utility in this situation is $U = 19.34$. This is the highest utility that can be obtained in this situation without regulating the externality.

[11]The classic reference is C. M. Tiebout, "A Pure Theory of Local Expenditures," *Journal of Political Economy* (October 1956): 416–24.

A Pigovian tax. As suggested by Figure 19.1, a unit tax on the production of good y may improve matters in this situation. With a tax rate of 0.1, for example, output of good y is reduced to $y = 27.4$ ($c = 10 - 5.48 = 4.52$), and the revenue generated is used to expand public goods production to $g = 3.77$. Utility is increased to $U = 19.38$. By carefully specifying how the revenue generated by the Pigovian tax is used, a general equilibrium model permits a more complete statement of welfare effects.

The "double dividend" of environmental taxes. The solution just described is not optimal, however. Production of public goods is actually too high in this case, since the revenues from environmental taxes are also used to pay for public goods. In fact, simulations show that optimality can be achieved by reducing the labor tax to 0.20 and public goods production to $g = 3.31$. With these changes, utility expands even further to $U = 19.43$. This result is sometimes referred to as the "double dividend" of environmental taxation: not only do these taxes reduce externalities relative to the untaxed situation (now $c = 10 - 5.60 = 4.40$), but also the extra governmental revenue made available thereby may permit the reduction of other distorting taxes.

QUERY: Why does the quantity of clean air decrease slightly when the labor tax is reduced relative to the situation where it is maintained at 0.25? More generally, describe whether environmental taxes would be expected always to generate a double dividend.

VOTING AND RESOURCE ALLOCATION

Voting is used as a social decision process in many institutions. In some instances, individuals vote directly on policy questions. That is the case in some New England town meetings, many statewide referenda (for example, California's Proposition 13 in 1977), and for many of the national policies adopted in Switzerland. Direct voting also characterizes the social decision procedure used for many smaller groups and clubs such as farmers' cooperatives, university faculties, or the local Rotary Club. In other cases, however, societies have found it more convenient to use a representative form of government, in which individuals vote directly only for political representatives, who are then charged with making decisions on policy questions. For our study of public choice theory, we will begin with an analysis of direct voting. This is an important subject not only because such a procedure applies to many cases but also because elected representatives often engage in direct voting (in Congress, for example), and the theory we will illustrate applies to those instances as well.

Majority rule

Because so many elections are conducted on a majority rule basis, we often tend to regard that procedure as a natural and, perhaps, optimal one for making social choices. But even a cursory examination indicates that there is nothing particularly sacred about a rule requiring that a policy obtain 50 percent of the vote to be adopted. In the U.S. Constitution, for example, two thirds of the states must adopt an amendment before it becomes law. And 60 percent of the U.S. Senate must vote to limit debate on controversial issues. Indeed, in some institutions (Quaker meetings, for example), unanimity may be required for social decisions. Our discussion of the Lindahl equilibrium concept suggests there may exist a distribution of tax shares that would obtain unanimous support in voting for public goods. But arriving at such unanimous agreements is usually thwarted by emergence of the free-rider problem. Examining in detail the forces that lead societies to move

TABLE 19.2 PREFERENCES THAT PRODUCE THE PARADOX OF VOTING

Choices: A—Low Spending
B—Medium Spending
C—High Spending

Preferences	Smith	Jones	Fudd
	A	B	C
	B	C	A
	C	A	B

away from unanimity and to choose some other determining fraction would take us too far afield here. We instead will assume throughout our discussion of voting that decisions will be made by majority rule. Readers may wish to ponder for themselves what kinds of situations might call for a decisive proportion of other than 50 percent.

The paradox of voting

In the 1780s, the French social theorist M. de Condorcet observed an important peculiarity of majority rule voting systems—they may not arrive at an equilibrium but instead may cycle among alternative options. Condorcet's paradox is illustrated for a simple case in Table 19.2. Suppose there are three voters (Smith, Jones, and Fudd) choosing among three policy options. For our subsequent analysis we will assume the policy options represent three levels of spending (A low, B medium, or C high) on a particular public good, but Condorcet's paradox would arise even if the options being considered did not have this type of ordering associated with them. Preferences of Smith, Jones, and Fudd among the three policy options are indicated in Table 19.2. These preferences give rise to Condorcet's paradox.

Consider a vote between options A and B. Here option A would win, because it is favored by Smith and Fudd and opposed only by Jones. In a vote between options A and C, option C would win, again by 2 votes to 1. But in a vote of C versus B, B would win and we would be back where we started. Social choices would endlessly cycle among the three alternatives. In subsequent votes, any choice initially decided upon could be defeated by an alternative, and no equilibrium would ever be reached. In this situation, the option finally chosen will depend on such seemingly nongermane issues as when the balloting stops or how items are ordered on an agenda—rather than being derived in some rational way from the preferences of voters.

Single-peaked preferences and the median voter theorem

Condorcet's voting paradox arises because there is a degree of irreconcilability in the preferences of voters. Therefore, one might ask whether restrictions on the types of preferences allowed could yield situations where equilibrium voting outcomes are more likely. A fundamental result about this probability was discovered by Duncan Black in 1948.[12] Black showed that equilibrium voting outcomes always occur in cases where the issue being voted upon is one-dimensional (such as how much to spend on a public good) and where voters' preferences are "single peaked." To understand what the notion of single peaked means, consider again Condorcet's paradox. In Figure 19.3 we illustrate the

[12]D. Black, "On the Rationale of Group Decision Making," *Journal of Political Economy* (February 1948): 23–34.

FIGURE 19.3

Single-Peaked Preferences and the Median Voter Theorem

This figure illustrates the preferences in Table 19.2. Smith's and Jones's preferences are single peaked, but Fudd's have two local peaks and these yield the voting paradox. If Fudd's preferences had instead been single peaked (the dashed line), then option B would have been chosen as the preferred choice of the median voter (Jones).

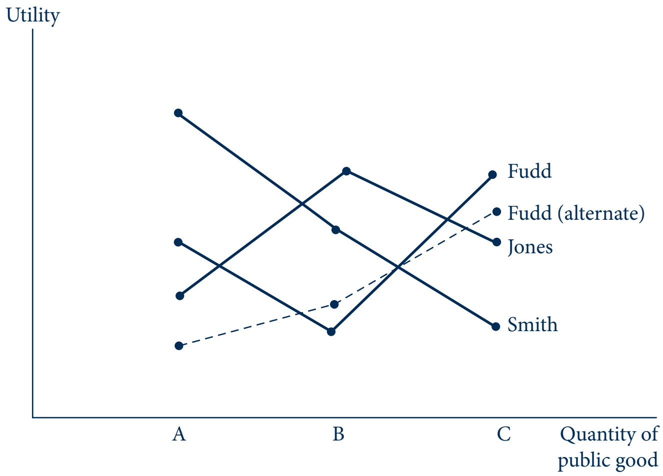

preferences that gave rise to the paradox by assigning hypothetical utility levels to options A, B, and C that are consistent with the preferences recorded in Table 19.2. For Smith and Jones, preferences are single peaked: as levels of public goods expenditures increase, there is only one local utility-maximizing choice (A for Smith, B for Jones). Fudd's preferences, on the other hand, have two local maxima (A and C). It is these preferences that produced the cyclical voting pattern. If instead Fudd had the preferences represented by the dashed line in Figure 19.3 (where now C is the only local utility maximum), then there would be no paradox. In this case, option B would be chosen because that option would defeat both A and C by votes of 2 to 1. Here B is the preferred choice of the "median" voter (Jones), whose preferences are "between" the preferences of Smith and the revised preferences of Fudd.

Black's result is quite general and applies to any number of voters. If choices are unidimensional[13] and if preferences are single peaked, then majority rule will result in the selection of the project that is most favored by the median voter. Hence, that voter's preferences will determine what public choices are made. This result is a key starting point for many models of the political process. In such models, the median voter's preferences dictate policy choices—either because that voter determines which policy gets a majority of votes in a direct election or because the median voter will dictate choices in competitive elections in which candidates must adopt policies that appeal to this voter.

A SIMPLE POLITICAL MODEL

To illustrate how the median voter theorem is applied in political models, suppose a community is characterized by a large number (n) of voters each with an income given by y_i.

[13]The result can be generalized a bit to deal with multidimensional policies if individuals can be characterized in their support for such policies along a single dimension.

The utility of each voter depends on his or her consumption of a private good (c_i) and of a public good (g) according to the additive utility function

$$\text{utility of person } i = U_i = c_i + f(g), \tag{19.59}$$

where $f_g > 0$ and $f_{gg} < 0$.

Each voter must pay income taxes to finance g. Taxes are proportional to income and are imposed at a rate t. Therefore, each person's budget constraint is given by

$$c_i = (1-t)y_i. \tag{19.60}$$

The government is also bound by a budget constraint:

$$g = \sum_{1}^{n} ty_i = tny^A, \tag{19.61}$$

where y^A denotes average income for all voters.

Given these constraints, the utility of person i can be written as a function of his or her choice of g only:

$$U_i(g) = \left(y^A - \frac{g}{n}\right)\frac{y^i}{y^A} + f(g). \tag{19.62}$$

Utility maximization for person i shows that his or her preferred level of expenditures on the public good satisfies

$$\frac{dU_i}{dg} = -\frac{y_i}{ny^A} + f_g(g) = 0 \quad \text{or} \quad g = f_g^{-1}\left(\frac{y_i}{ny^A}\right). \tag{19.63}$$

This shows that desired spending on g is inversely related to income. Because (in this model) the benefits of g are independent of income but taxes increase with income, high-income voters can expect to have smaller net gains (or even losses) from public spending than can low-income voters.

The median voter equilibrium

If g is determined here through majority rule, its level will be chosen to be that level favored by the "median voter." In this case, voters' preferences align exactly with incomes, so g will be set at that level preferred by the voter with median income (y^m). Any other level for g would not get 50 percent of the vote. Hence, equilibrium g is given by

$$g^* = f_g^{-1}\left(\frac{y^m}{ny^A}\right) = f_g^{-1}\left[\left(\frac{1}{n}\right)\left(\frac{y^m}{y^A}\right)\right]. \tag{19.64}$$

In general, the distribution of income is skewed to the right in practically every political jurisdiction in the world. With such an income distribution, $y^m < y^A$, and the difference between the two measures becomes larger the more skewed is the income distribution. Hence Equation 19.64 suggests that, ceteris paribus, the more unequal is the income distribution in a democracy, the higher will be tax rates and the greater will be spending on public goods. Similarly, laws that extend the vote to increasingly poor segments of the population can also be expected to increase such spending.

Optimality of the median voter result

Although the median voter theorem permits a number of interesting positive predictions about the outcome of voting, the normative significance of these results is more difficult to pinpoint. In this example, it is clear that the result does not replicate the Lindahl voluntary equilibrium—high-income voters would not voluntarily agree to the taxes imposed.[14] The result also does not necessarily correspond to any simple criterion for social welfare. For example, under a "utilitarian" social welfare criterion, g would be chosen so as to maximize the sum of utilities:

$$SW = \sum_{i=1}^{n} U_i = \sum_{i=1}^{n} \left[\left(y^A - \frac{g}{n} \right) \frac{y_i}{y^A} + f(g) \right] = ny^A - g + nf(g). \quad \textbf{(19.65)}$$

The optimal choice for g is then found by differentiation:

$$\frac{dSW}{dg} = -1 + nf_g = 0,$$

or

$$g^* = f_g^{-1}\left(\frac{1}{n}\right) = f_g^{-1}\left[\left(\frac{1}{n}\right)\left(\frac{y^A}{y^A}\right)\right], \quad \textbf{(19.66)}$$

which shows that a utilitarian choice would opt for the level of g favored by the voter with *average* income. That output of g would be smaller than that favored by the median voter because $y^m < y^A$. In Example 19.6 we take this analysis a bit further by showing how it might apply to governmental transfer policy.

EXAMPLE 19.6 Voting for Redistributive Taxation

Suppose voters were considering adoption of a lump-sum transfer to be paid to every person and financed through proportional taxation. If we denote the per-person transfer by b, then each individual's utility is now given by

$$U_i = c_i + b \quad \textbf{(19.67)}$$

and the government budget constraint is

$$nb = tny^A \quad \text{or} \quad b = ty^A. \quad \textbf{(19.68)}$$

For a voter whose income is greater than average, utility would be maximized by choosing $b = 0$, because such a voter would pay more in taxes than he or she would receive from the transfer. Any voter with less than average income will gain from the transfer no matter what the tax rate is. Hence such voters (including the decisive median voter) will opt for $t = 1$ and $b = y^A$. That is, they would vote to fully equalize incomes through the tax system. Of course, such a tax scheme is unrealistic—primarily because a 100 percent tax rate would undoubtedly create negative work incentives that reduce average income.

To capture such incentive effects, assume[15] that each person's income has two components, one responsive to tax rates $[y_i(t)]$ and one not responsive (n_i). Assume also that the average value of n_i is 0 but that its distribution is skewed to the right, so $n_m < 0$. Now utility is given by

$$U_i = (1 - t)[y_i(t) + n_i] + b. \quad \textbf{(19.69)}$$

[14] Although they might if the benefits of g were also proportional to income.

[15] What follows represents a much simplified version of a model first developed by T. Romer in "Individual Welfare, Majority Voting, and the Properties of a Linear Income Tax," *Journal of Public Economics* (December 1978): 163–68.

Assuming that each person first optimizes over those variables (such as labor supply) that affect $y_i(t)$, the first-order condition[16] for a maximum in his or her political decisions about t and b then becomes (using the government budget constraint in Equation 19.68)

$$\frac{dU_i}{dt} = -n_i + t\frac{dy^A}{dt} = 0. \tag{19.70}$$

Hence for voter i the optimal redistributive tax rate is given by

$$t_i = \frac{n_i}{dy^A/dt}. \tag{19.71}$$

Assuming political competition under majority rule voting will opt for that policy favored by the median voter, the equilibrium rate of taxation will be

$$t^* = \frac{n_m}{dy^A/dt}. \tag{19.72}$$

Because both n_m and dy^A/dt are negative, this rate of taxation will be positive. The optimal tax will be greater the farther n_m is from its average value (i.e., the more unequally income is distributed). Similarly, the larger are distortionary effects from the tax, the smaller the optimal tax. This model then poses some rather strong testable hypotheses about redistribution in the real world.

QUERY: Would progressive taxation be more likely to raise or lower t^* in this model?

VOTING MECHANISMS

The problems involved in majority rule voting arise in part because such voting is simply not informative enough to provide accurate appraisals of how people value public goods. This situation is in some ways similar to some of the models of asymmetric information examined in the previous chapter. Here voters are more informed than is the government about the value they place on various tax-spending packages. Resource allocation would be improved if mechanisms could be developed that encourage people to be more accurate in what they reveal about these values. In this section we examine two such mechanisms. Both are based on the basic insight from Vickrey second-price auctions (see Chapter 18) that incorporating information about other bidders' valuations into decision-makers' calculations can yield a greater likelihood of revealing truthful valuations.

The Groves mechanism

In a 1973 paper, T. Groves proposed a way to incorporate the Vickrey insight into a method for encouraging people to reveal their demands for a public good.[17] To illustrate this mechanism, suppose that there are n individuals in a group and each has a private (and unobservable) net valuation v_i for a proposed taxation–expenditure project. In seeking information about these valuations, the government states that, should the project be undertaken, each person will receive a transfer given by

$$t_i = \sum_{j \neq i} \tilde{v}_j \tag{19.73}$$

[16]Equation 19.70 can be derived from 19.69 through differentiation and by recognizing that $dy_i/dt = 0$ because of the assumption of individual optimization.

[17]T. Groves, "Incentives in Teams," *Econometrica* (July 1973): 617–31.

where $\tilde{v}_j$ represents the valuation reported by person j and the summation is taken over all individuals other than person i. If the project is not undertaken, then no transfers are made.

Given this setup, the problem for voter i is to choose his or her reported net valuation so as to maximize utility, which is given by

$$\text{utility} = v_i + t_i = v_i + \sum_{j \neq i} \tilde{v}_j. \qquad \textbf{(19.74)}$$

Since the project will be undertaken only if $\sum_{i=1}^{n} \tilde{v}_i$ and since each person will wish the project to be undertaken only if it increases utility (i.e., $v_i + \sum_{j \neq i} \tilde{v}_j > 0$), it follows that a utility-maximizing strategy is to set $\tilde{v}_i = v_i$. Hence, the Groves mechanism encourages each person to be truthful in his or her reporting of valuations for the project.

The Clarke mechanism

A similar mechanism was proposed by E. Clarke, also in the early 1970s.[18] This mechanism also envisions asking individuals about their net valuations for some public project, but it focuses mainly on "pivotal voters"—those whose reported valuations can change the overall evaluation from negative to positive or vice versa. For all other voters, there are no special transfers, on the presumption that reporting a nonpivotal valuation will not change either the decision or the (zero) payment, so he or she might as well report truthfully. For voters reporting pivotal valuations, however, the Clarke mechanism incorporates a Pigovian-like tax (or transfer) to encourage truth telling. To see how this works, suppose that the net valuations reported by all other voters are negative ($\sum_{j \neq i} \tilde{v}_j < 0$), but that a truthful statement of the valuation by person i would make the project acceptable ($v_i + \sum_{j \neq i} \tilde{v}_j > 0$). Here, as for the Groves mechanism, a transfer of $t_i = \sum_{j \neq i} \tilde{v}_j$ (which in this case would be negative—i.e., a tax) would encourage this pivotal voter to report $\tilde{v}_i = v_i$. Similarly, if all other individuals reported valuations favorable to a project ($\sum_{j \neq i} \tilde{v}_j > 0$) but inclusion of person i's evaluation of the project would make it unfavorable, then a transfer of $t_i = \sum_{j \neq i} \tilde{v}_j$ (which in this case is positive) would encourage this pivotal voter to choose $\tilde{v}_i = v_i$ also. Overall, then, the Clarke mechanism is also truth revealing. Notice that in this case the transfers play much the same role that Pigovian taxes did in our examination of externalities. If other voters view a project as unfavorable, then voter i must compensate them for accepting it. On the other hand, if other voters find the project acceptable, then voter i must be sufficiently against the project that he or she cannot be "bribed" by other voters into accepting it.

Generalizations

The voter mechanisms we have been describing are sometimes called *VCG mechanisms* after the three pioneering economists in this area of research (Vickrey, Clarke, and Groves). These mechanisms can be generalized to include multiple governmental projects, alternative concepts of voter equilibrium, or an infinite number of voters. One assumption behind the mechanisms that does not seem amenable to generalization is the quasi-linear utility functions that we have been using throughout. Whether this assumption provides a good approximation for modeling political decision making remains an open question, however.

[18] E. Clarke, "Multipart Pricing for Public Goods," *Public Choice* (Fall 1971): 19–33.