CHAPTER 36

PUBLIC GOODS

In Chapter 34 we argued that for certain kinds of externalities, it was not difficult to eliminate the inefficiencies. In the case of a consumption externality between two people, for example, all one had to do was to ensure that initial property rights were clearly specified. People could then trade the right to generate the externality in the normal way. In the case of production externalities, the market itself provided profit signals to sort out the property rights in the most efficient way. In the case of common property, assigning property rights to someone would eliminate the inefficiency.

Unfortunately, not all externalities can be handled in that manner. As soon as there are more than two economic agents involved things become much more difficult. Suppose, for example, that instead of the two roommates examined in the last chapter, we had *three* roommates—one smoker and two nonsmokers. Then the amount of smoke would be a negative externality for both of the nonsmokers.

Let's suppose that property rights are well defined—say the nonsmokers have the right to demand clean air. Just as before, although they have the *right* to clean air, they also have the right to trade some of that clean air away in return for appropriate compensation. But now there is a problem involved—the nonsmokers have to agree among themselves how much smoke should be allowed and what the compensation should be.

Perhaps one of the nonsmokers is much more sensitive than the other, or one of them is much richer than the other. They may have very different preferences and resources, and yet they both have to reach some kind of agreement to allow for an efficient allocation of smoke.

Instead of roommates, we can think of inhabitants of a whole country. How much pollution should be allowed in the country? If you think that reaching an agreement is difficult with only three roommates, imagine what it is like with millions of people!

The smoke externality with three people is an example of a **public good**—a good that must be provided in the same amount to all the affected consumers. In this case the amount of smoke generated will be the same for all consumers—each person may value it differently, but they all have to face the same amount.

Many public goods are provided by the government. For example, streets and sidewalks are provided by local municipalities. There are a certain number and quality of streets in a town, and everyone has that number available to use. National defense is another good example; there is one level of national defense provided for all the inhabitants of a country. Each citizen may value it differently—some may want more, some may want less—but they are all provided with the same amount.

Public goods are an example of a particular kind of consumption externality: everyone must consume the same amount of the good. They are a particularly troublesome kind of externality, for the decentralized market solutions that economists are fond of don't work very well in allocating public goods. People can't purchase different amounts of public defense; somehow they have to decide on a common amount.

The first issue to examine is what the ideal amount of the public good should be. Then we'll discuss some ways that might be used to make social decisions about public goods.

36.1 When to Provide a Public Good?

Let us start with a simple example. Suppose that there are two roommates, 1 and 2. They are trying to decide whether or not to purchase a TV. Given the size of their apartment, the TV will necessarily go in the living room, and both roommates will be able to watch it. Thus it will be a public good, rather than a private good. The question is, is it worth it for them to acquire the TV?

Let's use w_1 and w_2 to denote each person's initial wealth, g_1 and g_2 to denote each person's contribution to the TV, and x_1 and x_2 to denote each person's money left over to spend on private consumption. The budget constraints are given by

$$x_1 + g_1 = w_1$$
$$x_2 + g_2 = w_2.$$

We also suppose that the TV costs c dollars, so that in order to purchase it, the sum of the two contributions must be at least c:

$$g_1 + g_2 \geq c.$$

This equation summarizes the technology available to provide the public good: the roommates can acquire one TV if together they pay the cost c.

The utility function of person 1 will depend on his or her private consumption, x_1, and the availability of the TV—the public good. We'll write person 1's utility function as $u_1(x_1, G)$, where G will either be 0, indicating no TV, or 1, indicating that a TV is present. Person 2 will have utility function $u_2(x_2, G)$. Each person's private consumption has a subscript to indicate that the good is consumed by person 1 or person 2, but the public good has no subscript. It is "consumed" by both people. Of course, it isn't really consumed in the sense of being "used up"; rather, it is the *services* of the TV that are consumed by the two roommates.

The roommates may value the services of the TV quite differently. We can measure the value that each person places on the TV by asking how much each person would be willing to pay to have the TV available. To do this, we'll use the concept of the **reservation price**, introduced in Chapter 15.

The reservation price of person 1 is the maximum amount that person 1 would be willing to pay to have the TV present. That is, it is that price, r_1, such that person 1 is just indifferent between paying r_1 and having the TV available, and not having the TV at all. If person 1 pays the reservation price and gets the TV, he will have $w_1 - r_1$ available for private consumption. If he doesn't get the TV, he will have w_1 available for private consumption. If he is to be just indifferent between these two alternatives, we must have

$$u_1(w_1 - r_1, 1) = u_1(w_1, 0).$$

This equation defines the reservation price for person 1—the maximum amount that he would be willing to pay to have the TV present. A similar equation defines the reservation price for person 2. Note that in general the reservation price of each person will depend on that person's wealth: the maximum amount that an individual will be *willing* to pay will depend to some degree on how much that individual is *able* to pay.

Recall that an allocation is Pareto efficient if there is no way to make both people better off. An allocation is Pareto *inefficient* if there *is* some way to make both people better off; in this case, we say that a **Pareto improvement** is possible. In the TV problem there are only two sorts of allocations that are of interest. One is an allocation where the TV is not provided. This allocation takes the simple form $(w_1, w_2, 0)$; that is, each person spends his wealth only on his private consumption.

The other kind of allocation is the one where the public good is provided. This will be an allocation of the form $(x_1, x_2, 1)$, where

$$x_1 = w_1 - g_1$$

$$x_2 = w_2 - g_2.$$

These two equations come from rewriting the budget constraints. They say that each individual's private consumption is determined by the wealth that he has left over after making his contribution to the public good.

Under what conditions should the TV be provided? That is, when is there a payment scheme (g_1, g_2) such that both people will be better off having the TV and paying their share than not having the TV? In the language of economics, when will it be a Pareto improvement to provide the TV?

It will be a Pareto improvement to provide the allocation $(x_1, x_2, 1)$ if both people would be better off having the TV provided than not having it provided. This means

$$u_1(w_1, 0) < u_1(x_1, 1)$$
$$u_2(w_2, 0) < u_2(x_2, 1).$$

Now use the definition of the reservation prices r_1 and r_2 and the budget constraint to write

$$u_1(w_1 - r_1, 1) = u_1(w_1, 0) < u_1(x_1, 1) = u_1(w_1 - g_1, 1)$$
$$u_2(w_2 - r_2, 1) = u_2(w_2, 0) < u_2(x_2, 1) = u_2(w_2 - g_2, 1).$$

Looking at the left- and the right-hand sides of these inequalities, and remembering that more private consumption must increase utility, we can conclude that

$$w_1 - r_1 < w_1 - g_1$$
$$w_2 - r_2 < w_2 - g_2,$$

which in turn implies

$$r_1 > g_1$$
$$r_2 > g_2.$$

This is a condition that must be satisfied if an allocation $(w_1, w_2, 0)$ is Pareto inefficient: it must be that the contribution that each person is making to the TV is less than his willingness to pay for the TV. If a consumer can acquire the good for less than the maximum that he would be willing to pay, then the acquisition would be to his benefit. Thus the condition that the reservation price exceeds the cost share simply says that a Pareto improvement will result when each roommate can acquire the services of the TV for less than the maximum that he would be willing to pay for it. This is clearly a *necessary* condition for purchase of the TV to be a Pareto improvement.

If each roommate's willingness to pay exceeds his cost share, then the *sum* of the willingnesses to pay must be greater than the cost of the TV:

$$r_1 + r_2 > g_1 + g_2 = c. \tag{36.1}$$

This condition is a *sufficient* condition for it to be a Pareto improvement to provide the TV. If the condition is satisfied, then there will be some payment plan such that both people will be made better off by providing the public good. If $r_1 + r_2 \geq c$, then the total amount that the roommates will be willing to pay is at least as large as the cost of purchase, so they can easily find a payment plan (g_1, g_2) such that $r_1 \geq g_1$, $r_2 \geq g_2$, and $g_1 + g_2 = c$. This condition is so simple that you might wonder why we went through all the detail in deriving it. Well, there are a few subtleties involved.

First, it is important to note that the condition describing when provision of the public good will be a Pareto improvement only depends on each agent's *willingness* to pay and on the total cost. If the sum of the reservation prices exceeds the cost of the TV, then there will always *exist* a payment scheme such that both people will be better off having the public good than not having it.

Second, whether or not it is Pareto efficient to provide the public good will, in general, depend on the initial distribution of wealth (w_1, w_2). This is true because, in general, the reservation prices r_1 and r_2 will depend on the distribution of wealth. It perfectly possible that for some distributions of wealth $r_1 + r_2 > c$, and for other distributions of wealth $r_1 + r_2 < c$.

To see how this can be, imagine a situation where one roommate really loves the TV and the other roommate is nearly indifferent about acquiring it. Then if the TV-loving roommate had all of the wealth, he would be willing to pay more than the cost of the TV all by himself. Thus it would be a Pareto improvement to provide the TV. But if the indifferent roommate had all of the wealth, then the TV lover wouldn't have much money to contribute toward the TV, and it would be Pareto efficient *not* to provide the TV.

Thus, in general, whether or not the public good should be provided will depend on the distribution of wealth. But in specific cases the provision of the public good may be independent of the distribution of wealth. For example, suppose that the preferences of the two roommates were quasilinear. This means that the utility functions take the form

$$u_1(x_1, G) = x_1 + v_1(G)$$
$$u_2(x_2, G) = x_2 + v_2(G),$$

where G will be 0 or 1, depending on whether or not the public good is available. For simplicity, suppose that $v_1(0) = v_2(0) = 0$. This says that no TV provides zero utility from watching TV.[1]

In this case the definitions of the reservation prices become

$$u_1(w_1 - r_1, 1) = w_1 - r_1 + v_1(1) = u_1(w_1, 0) = w_1$$
$$u_2(w_2 - r_2, 1) = w_2 - r_2 + v_2(1) = u_2(w_2, 0) = w_2,$$

[1] Perhaps watching TV should be assigned a negative utility.

which implies that the reservation prices are given by

$$r_1 = v_1(1)$$
$$r_2 = v_2(1).$$

Thus the reservation prices are independent of the amount of wealth, and hence the optimal provision of the public good will be independent of wealth, at least over some range of wealths.[2]

36.2 Private Provision of the Public Good

We have seen above that acquiring the TV will be Pareto efficient for the two roommates if the sum of their willingnesses to pay exceeds the cost of providing the public good. This answers the question about efficient allocation of the good, but it does not necessarily follow that they will actually decide to acquire the TV. Whether they actually decide to acquire the TV depends on the particular method they adopt to make joint decisions.

If the two roommates cooperate and truthfully reveal how much they value the TV, then it should not be difficult for them to agree on whether or not they should buy the TV. But under some circumstances, they may not have incentives to tell the truth about their values.

For example, suppose that each person valued the TV equally, and that each person's reservation price was greater than the cost, so that $r_1 > c$ and $r_2 > c$. Then person 1 might think that if he said he had 0 value for the TV, the other person would acquire it anyway. But person 2 could reason the same way! One can imagine other situations where both people would refuse to contribute in the hopes that the other person would go out and unilaterally purchase the TV.

In this kind of situation, economists say that the people are attempting to **free ride** on each other: each person hopes that the other person will purchase the public good on his own. Since each person will have full use of the services of the TV if it is acquired, each person has an incentive to try to pay as little as possible toward the provision of the TV.

36.3 Free Riding

Free riding is similar, but not identical, to the prisoner's dilemma that we examined in Chapter 28. To see this, let us construct a numerical example of the TV problem described above. Suppose that each person has a wealth of \$500, that each person values the TV at \$100, and that the cost of the

[2] Even this will only be true for some ranges of wealth, since we must always require that $r_1 \leq w_1$ and $r_2 \leq w_2$—i.e., the willingness to pay is less than the ability to pay.

TV is \$150. Since the sum of the reservation prices exceeds the cost, it is Pareto efficient to buy the TV.

Let us suppose that there is no way for one of the roommates to exclude the other one from watching the TV and that each roommate will decide independently whether or not to buy the TV. Consider the decision of one of the roommates, Player A. If he buys the TV, he gets benefits of \$100 and pays a cost of \$150, leaving him with net benefits of -50. However, if Player A buys the TV, Player B gets to watch it for free, which gives B a benefit of \$100. The payoffs to the game are depicted in Table 36.1.

Free riding game matrix.

		Player B	
		Buy	Don't buy
Player A	Buy	-50, -50	-50, 100
	Don't buy	100, -50	0, 0

The dominant strategy equilibrium for this game is for neither player to buy the TV. If player A decides to buy the TV, then it is in player B's interest to free ride: to watch the TV but not contribute anything to paying for it. If player A decides not to buy, then it is in player B's interest not to buy the TV either. This is similar to the prisoners' dilemma, but not exactly the same. In the prisoners' dilemma, the strategy that maximizes the sum of the players' utilities is for each player to make the *same* choice. Here the strategy that maximizes the sum of the utilities is for just one of the players to buy the TV (and both players to watch it).

If Player A buys the TV and both players watch it, we can construct a Pareto improvement simply by having Player B make a "sidepayment" to Player A. For example, if Player B gives Player A \$51, then both players will be made better off when Player A buys the TV. More generally, any payment between \$50 and \$100 will result in a Pareto improvement for this example.

In fact, this is probably what would happen in practice: each player would contribute some fraction of the cost of the TV. This public goods problem is relatively easy to solve, but more difficult free riding problems can arise in the sharing of other household public goods. For example, what about cleaning the living room? Each person may prefer to see the living room clean and is willing to do his part. But each may also be tempted to free ride on the other—so that neither one ends up cleaning the room, with the usual untidy results.

The situation becomes even worse if there are more than just two people involved—since there are more people on whom to free ride! Letting the other guy do it may be optimal from an *individual* point of view, but it is Pareto inefficient from the viewpoint of society as a whole.

36.4 Different Levels of the Public Good

In the above example, we had an either/or decision: either provide the TV or not. But the same kind of phenomena occurs when there is a choice of *how much* of the public good to provide. Suppose, for example, that the two roommates have to decide how much money to spend on the TV. The more money they decide to spend, the better the TV they can get.

As before we'll let x_1 and x_2 measure the private consumption of each person and g_1 and g_2 be their contributions to the TV. Let G now measure the "quality" of the TV they buy, and let the cost function for quality be given by $c(G)$. This means that if the two roommates want to purchase a TV of quality G, they have to spend $c(G)$ dollars to do so.

The constraint facing the roommates is that the total amount that they spend on their public and private consumption has to add up to how much money they have:

$$x_1 + x_2 + c(G) = w_1 + w_2.$$

A Pareto efficient allocation is one where consumer 1 is as well-off as possible given consumer 2's level of utility. If we fix the utility of consumer 2 at $\overline{u}_2$, we can write this problem as

$$\max_{x_1,x_2,G} u_1(x_1, G)$$

$$\text{such that } u_2(x_2, G) = \overline{u}_2$$

$$x_1 + x_2 + c(G) = w_1 + w_2.$$

It turns out that the appropriate optimality condition for this problem is that the *sum* of the absolute values of the marginal rates of substitution between the private good and the public good for the two consumers equals the marginal cost of providing an extra unit of the public good:

$$|MRS_1| + |MRS_2| = MC(G)$$

or, spelling out the definitions of the marginal rates of substitution,

$$\left|\frac{\Delta x_1}{\Delta G}\right| + \left|\frac{\Delta x_2}{\Delta G}\right| = \frac{MU_G}{MU_{x_1}} + \frac{MU_G}{MU_{x_2}} = MC(G).$$

In order to see why this must be the right efficiency condition, let us apply the usual trick and think about what would be the case if it were violated. Suppose, for example, that the sum of the marginal rates of substitution were less than the marginal cost: say $MC = 1$, $|MRS_1| = 1/4$, and $|MRS_2| = 1/2$. We need to show that there is some way to make both people better off.

Given his marginal rate of substitution, we know that person 1 would be willing to accept 1/4 more dollars of the private good for the loss of 1 dollar of the public good (since both goods cost \$1 per unit). Similarly, person 2 would accept 1/2 more dollars of the private good for a 1-dollar decrease in the public good. Suppose we reduce the amount of the public good and offer to compensate both individuals. When we reduce the public good by one unit we save a dollar. After we pay each individual the amount he requires to allow this change $(3/4 = 1/4 + 1/2)$, we find that we still have 1/4 of a dollar left over. This remaining money could be shared between the two individuals, thereby making them both better off.

Similarly, if the sum of the marginal rates of substitution were greater than 1, we could increase the amount of the public good to make them both better off. If $|MRS_1| = 2/3$ and $|MRS_2| = 1/2$, say, this means that person 1 would give up 2/3 of a dollar of private consumption to get 1 unit more of the public good and person 2 would give up 1/2 of a dollar of private consumption to get 1 unit more of the public good. But if person 1 gave up his 2/3 units, and person 2 gave up his 1/2 unit, we would have more than enough to produce the extra unit of the public good, since the marginal cost of providing the public good is 1. Thus we could give the left-over amount back to both people, thereby making them both better off.

What does the condition for Pareto efficiency mean? One way to interpret it is to think of the marginal rate of substitution as measuring the *marginal* willingness to pay for an extra unit of the public good. Then the efficiency condition just says that the *sum* of the marginal willingnesses to pay must equal the marginal cost of providing an extra unit of the public good.

In the case of a discrete good that was either provided or not provided, we said that the efficiency condition was that the sum of the willingnesses to pay should be at least as large as the cost. In the case we're considering here, where the public good can be provided at different levels, the efficiency condition is that the sum of the *marginal* willingnesses to pay should *equal* the marginal cost at the optimal amount of the public good. For whenever the sum of the marginal willingnesses to pay for the public good exceeds the marginal cost, it is appropriate to provide more of the public good.

It is worthwhile comparing the efficiency condition for a public good to the efficiency condition for a private good. For a private good, each person's marginal rate of substitution must equal the marginal cost; for a public good, the *sum* of the marginal rates of substitution must equal the marginal cost. In the case of a private good, each person can consume a different amount of the private good, but they all must value it the same at the margin—otherwise they would want to trade. In the case of a public good, each person must consume the same amount of the public good, but they can all value it differently at the margin.

We can illustrate the public good efficiency condition in Figure 36.1. We simply draw each person's MRS curve and then add them vertically to get the sum of the MRS curves. The efficient allocation of the public good will occur where the sum of the MRSs equals the marginal cost, as illustrated in Figure 36.1.

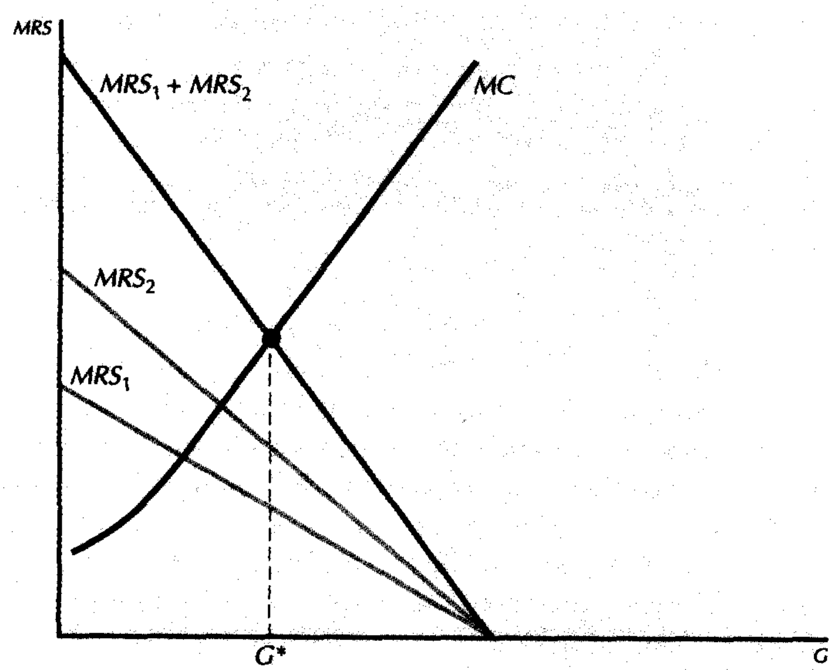

Determining the efficient amount of a public good. The sum of the marginal rates of substitution must equal the marginal cost.

36.5 Quasilinear Preferences and Public Goods

In general, the optimal amount of the public good will be different at different allocations of the private good. But if the consumers have quasilinear preferences it turns out that there will be a unique amount of the public good supplied at every efficient allocation. The easiest way to see this is to think about the kind of utility function that represents quasilinear preferences.

As we saw in Chapter 4, quasilinear preferences have a utility representation of the form: $u_i(x_i, G) = x_i + v_i(G)$. This means that the marginal utility of the private good is always 1, and thus the marginal rate of substitution between the private and the public good—the ratio of the marginal utilities—will depend only on G. In particular:

$$
\begin{aligned}
|MRS_1| &= \frac{\Delta u_1(x_1, G)/\Delta G}{\Delta u_1/\Delta x_1} = \frac{\Delta v_1(G)}{\Delta G} \\
|MRS_2| &= \frac{\Delta u_2(x_2, G)/\Delta G}{\Delta u_2/\Delta x_2} = \frac{\Delta v_2(G)}{\Delta G}.
\end{aligned}
$$

We already know that a Pareto efficient level of the public good must satisfy the condition

$$|MRS_1| + |MRS_2| = MC(G).$$

Using the special form of the MRSs in the case of quasilinear utility, we can write this condition as

$$\frac{\Delta v_1(G)}{\Delta G} + \frac{\Delta v_2(G)}{\Delta G} = MC(G).$$

Note that this equation determines G without any reference to x_1 or x_2. Thus there is a unique efficient level of provision of the public good.

Another way to see this is to think about the behavior of the indifference curves. In the case of quasilinear preferences, all of the indifference curves are just shifted versions of each other. This means, in particular, that the slope of the indifference curves—the marginal rate of substitution—doesn't change as we change the amount of the private good. Suppose that we find one efficient allocation of the public and private goods, where the sum of the absolute value of the MRSs equals $MC(G)$. Now if we take some amount of the private good away from one person and give it to another, the slopes of both indifference curves stay the same, so the sum of the absolute value of the MRSs is still equal to $MC(G)$ and we have another Pareto efficient allocation.

In the case of quasilinear preferences, all Pareto efficient allocations are found by just redistributing the private good. The amount of the public good stays fixed at the efficient level.

EXAMPLE: Pollution Revisited

Recall the model of the steel firm and the fishery described in Chapter 34. There we argued that the efficient provision of pollution was one which internalized the pollution costs borne by the steel firm and the fishery. Suppose now that there are two fisheries, and that the amount of pollution produced by the steel firm is a public good. (Or, perhaps more appropriately, is a public bad!)

Then the efficient provision of pollution will involve maximizing the sum of the profits of all three firms—that is, minimizing the total social cost of the pollution. Formally, let $c_s(s, x)$ be the cost to the steel firm of producing s units of steel and x units of pollution, and write $c_f^1(f_1, x)$ for the costs for firm 1 to catch f_1 fish when the pollution level is x, and $c_f^2(f_2, x)$ as the analogous expression for firm 2. Then to compute the Pareto efficient amount of pollution, we maximize the sum of the three firms' profits:

$$\max_{s, f_1, f_2, x} \; p_s s + p_f f_1 + p_f f_2 - c_s(s, x) - c_f^1(f_1, x) - c_f^2(f_2, x).$$

The interesting effect for our purposes is the effect on aggregate profits of increasing pollution. Increasing pollution lowers the cost of producing steel but raises the costs of producing fish for each of the fisheries. The appropriate optimality condition from the profit-maximization problem is

$$\frac{\Delta c_s(\hat{s}, \hat{x})}{\Delta x} + \frac{\Delta c_f^1(\hat{f}_1, \hat{x})}{\Delta x} + \frac{\Delta c_f^2(\hat{f}_2, \hat{x})}{\Delta x} = 0,$$

which simply says that the *sum* of the marginal costs of pollution over the three firms should equal zero. Just as in the case of a public consumption good, it is the *sum* of the marginal benefits or costs over the economic agents that is relevant for determining the Pareto efficient provision of a public good.

36.6 The Free Rider Problem

Now that we know what the Pareto efficient allocations of public goods are, we can turn our attention to asking how to get there. In the case of private goods with no externalities we saw that the market mechanism will generate an efficient allocation. Will the market work in the case of public goods?

We can think of each person as having some endowment of a private good, w_i. Each person can spend some fraction of this private good on his own private consumption, or he or she can contribute some of it to purchase the public good. Let's use x_1 for 1's private consumption, and let g_1 denote the amount of the public good he buys, and similarly for person 2. Suppose for simplicity that $c(G) \equiv G$, which implies that the marginal cost of providing a unit of the public good is constant at 1. The total amount of the public good provided will be $G = g_1 + g_2$. Since each person cares about the *total* amount of the public good provided, the utility function of person i will have the form $u_i(x_i, g_1 + g_2) = u_i(x_i, G)$.

In order for person 1 to decide how much he should contribute to the public good, he has to have some forecast of how much person 2 will contribute. The simplest thing to do here is to adopt the Nash equilibrium model described in Chapter 28, and suppose that person 2 will make some contribution $\overline{g}_2$. We assume that person 2 also makes a guess about person 1's contribution, and we look for an equilibrium where each person is making an optimal contribution given the other person's behavior.

Thus person 1's maximization problem takes the form

$$\max_{x_1, g_1} u_1(x_1, g_1 + \overline{g}_2)$$

$$\text{such that } x_1 + g_1 = w_1.$$

This is just like an ordinary consumer maximization problem. The optimization condition is therefore the same: if both people purchase both goods the marginal rate of substitution between the public and the private goods should be 1 for each consumer:

$$|MRS_1| = 1$$

$$|MRS_2| = 1.$$

However, we have to be careful here. It is true that if person 2 purchases any amount of the public good at all, he will purchase it until the marginal rate of substitution equals one. But it can easily happen that person 2 decides that the amount already contributed by person 1 is sufficient and that it would therefore be unnecessary for him to contribute anything toward the public good at all.

Formally, we are assuming that the individuals can only make positive contributions to the public good—they can put money into the collection plate, but they can't take money out. Thus there is an extra constraint on each person's contributions, namely, that $g_1 \geq 0$ and $g_2 \geq 0$. Each person can only decide whether or not he wants to *increase* the amount of the public good. But then it may well be that one person decides that the amount provided by the other is just fine and would prefer to make no contribution at all.

A case like this is depicted in Figure 36.2. Here we have illustrated each person's private consumption on the horizontal axis and his or her public consumption on the vertical axis. The "endowment" of each person consists of his or her wealth, w_i, along with the amount of the public good contribution of the *other* person—since this is how much of the public good will be available if the person in question decides not to contribute. Figure 36.2A shows a case where person 1 is the only contributor to the public good, so that $g_1 = G$. If person 1 contributes G units to the public good, then person 2's endowment will consist of her private wealth, w_2, and the amount of the public good G—since person 2 gets to consume the public good whether or not she contributes to it. Since person 2 cannot reduce the amount of the public good, but can only increase it, her budget constraint is the bold line in Figure 36.2B. Given the shape of 2's indifference curve,

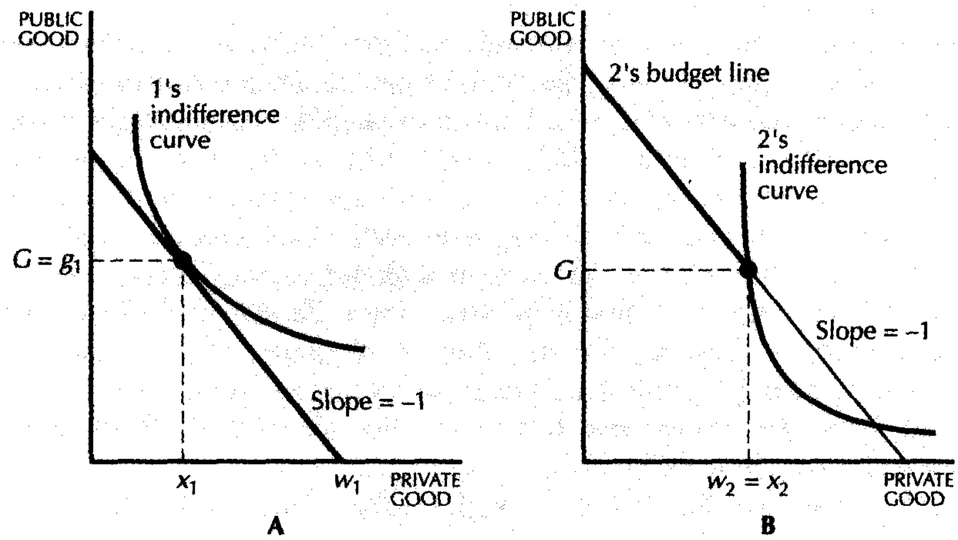

The free rider problem. Person 1 contributes while person 2 free rides.

it is optimal from her point of view to free ride on 1's contribution and simply consume her endowment, as depicted.

This is an example where person 2 is free riding on person 1's contribution to the public good. Since a public good is a good that everyone must consume in the same amount, the provision of a public good by any one person will tend to reduce the other peoples' provision. Thus in general there will be too little of the public good supplied in a voluntary equilibrium, relative to an efficient provision of the public good.

36.7 Comparison to Private Goods

In our discussion of private goods, we were able to show that a particular social institution—the competitive market—was capable of achieving a Pareto efficient allocation of private goods. Each consumer deciding for himself or herself how much to purchase of various goods would result in a pattern of consumption that was Pareto efficient. A major assumption in this analysis was that an individual's consumption did not affect other people's utility—that is, that there were no consumption externalities. Thus each person optimizing with respect to his or her own consumption was sufficient to achieve a kind of social optimum.

The situation is radically different with respect to public goods. In this case, the utilities of the individuals are inexorably linked since everyone is required to consume the same amount of the public good. In this case the market provision of *public* goods would be very unlikely to result in a Pareto efficient provision.

Indeed, for the most part we use *different* social institutions to determine the provision of public goods. Sometimes people use a **command mechanism**, where one person or small group of people determines the amount of various public goods that will be provided by the populace. Other times people use a **voting system** where individuals vote on the provision of public goods. One can well ask the same sorts of questions about voting, or other social mechanisms for decision making, that we asked about the private market: are they capable of achieving a Pareto efficient allocation of public goods? Can any Pareto efficient allocation of public goods be achieved by such mechanisms? A complete analysis of these questions is beyond the scope of this book, but we will be able to shed a little light on how some methods work below.

36.8 Voting

Private provision of a public good doesn't work very well, but there are several other mechanisms for social choice. One of the most common mechanisms in democratic countries is **voting**. Let's examine how well it works for the provision of public goods.

Voting isn't very interesting in the case of two consumers, so we will suppose that we have n consumers. Furthermore, so as not to worry about ties, we'll suppose that n is an odd number. Let's imagine that the consumers are voting about the size of some public good—say the magnitude of expenditures on public defense. Each consumer has a most-preferred level of expenditure, and his valuation of other levels of expenditure depends on how close they are to his preferred level of expenditure.

The first problem with voting as a way of determining social outcomes has already been examined in Chapter 33. Suppose that we are considering three levels of expenditure, A, B, and C. It is perfectly possible that there is a majority of the consumers who prefer A to B, a majority who prefer B to C ... and a majority who prefer C to A!

Using the terminology of Chapter 33, the social preferences generated by these consumers are not transitive. This means that the outcome of voting on the level of public good may not be well defined—there is always a level of expenditure that beats every expenditure. If a society is allowed to vote many times on an issue, this means that it may "cycle" around various choices. Or if a society votes only once on an issue, the outcome depends on the order in which the choices are presented.

If first you vote on A versus B and then on A versus C, C will be the outcome. But if you vote on C versus A and then C versus B, B will be the outcome. You can get any of the three outcomes by choosing how the alternatives are presented!

The "paradox of voting" described above is disturbing. One natural thing to do is to ask what restrictions on preferences will allow us to rule it out; that is, what form must preferences have so as to ensure that the kinds of cycles described above cannot happen?

Let us depict the preferences of consumer i by a graph like those in Figure 36.3, where the height of the graph illustrates the value or the net utility for different levels of the expenditure on the public good. The term "net utility" is appropriate since each person cares both about the level of the public good, and the amount that he has to contribute to it. Higher levels of expenditure mean more public goods but also higher taxes in order to pay for those public goods. Thus it is reasonable to assume that the net utility of expenditure on the public good rises at first due to the benefits of the public good but then eventually falls, due to the costs of providing it.

One restriction on preferences of this sort is that they be **single-peaked**. This means that preferences must have the shape depicted in Figure 36.3A rather than that depicted in Figure 36.3B. With single-peaked preferences, the net utility of different levels of expenditure rises until the most-preferred point and then falls, as it does in Figure 36.3A; it never goes up, down, and then up again, as it does in Figure 36.3B.

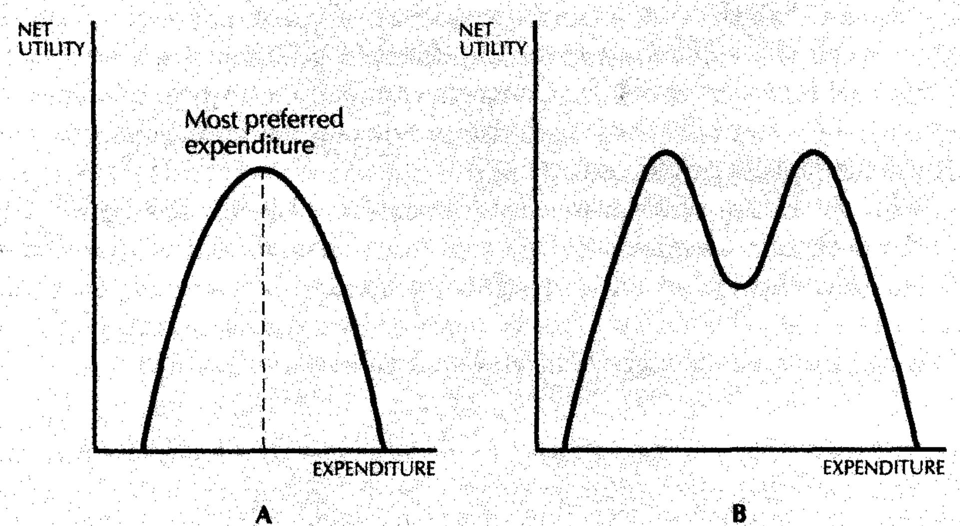

Shapes of preferences. Single-peaked preferences are shown in panel A and multiple peaked preferences in panel B.

If each individual has single-peaked preferences, then it can be shown that the social preferences revealed by majority vote will never exhibit the kind of intransitivity we described above. Accepting this result for the moment, we can ask which level of expenditure will be chosen if everyone has single-peaked preferences. The answer turns out to be the **median expenditure**—that expenditure such that one-half of the population wants to spend more, and one-half wants to spend less. This result is reasonably intuitive: if more than one-half wanted more expenditure on the public good, they would vote for more, so the only possible equilibrium voting outcome is when the votes for increasing and decreasing expenditure on the public good are just balanced.

Will this be an efficient level of the public good? In general, the answer is no. The median outcome just means that half the population wants more and half wants less; it doesn't say anything about *how much more* they want of the public good. Since efficiency takes this kind of information into account, voting will not in general lead to an efficient outcome.

Furthermore, even if peoples' true preferences are single-peaked, so that voting may lead to a reasonable outcome, individuals may choose to misrepresent their true preferences when they vote. Thus people will have an incentive to vote differently than their true preferences would indicate in order to manipulate the final outcome.

EXAMPLE: Agenda Manipulation

We have seen that the outcome of a sequence of votes may depend on the order in which the votes are taken. Experienced politicians are well aware of this possibility. In the U.S. Congress, amendments to a bill must be voted on before the bill itself, and this provides a commonly used way to influence the legislative process.

In 1956 the House of Representatives considered a bill calling for Federal aid to school construction. One representative offered an amendment requiring that the bill would only provide Federal aid to states with integrated schools. There were three more-or-less equally sized groups of representatives with strongly held views on this issue.

- Republicans. They were opposed to Federal aid to education, but preferred the amended bill to the original. Their ranking of the alternatives was no bill, amended bill, original bill.

- Northern Democrats. They wanted Federal aid to education and supported integrated schools, so they ranked the alternatives amended bill, original bill, no bill.

- Southern Democrats. This group wanted Federal aid to education, but would not get any aid under the amended bill due to the segregated schools in the South. Their ranking was original bill, no bill, amended bill.

In the vote on the amendment, the Republicans and the Northern Democrats were in the majority, thereby substituting the amended bill for the original. In the vote on the amended bill, the Republicans and the Southern Democrats were in the majority, and the amended bill was defeated. However, before being amended the original bill had a majority of the votes!

36.9 Demand Revelation

We have seen above that majority voting, even if it leads to a well-defined outcome, will not necessarily provide the correct incentives for people to honestly reveal their true preferences. In general, there will be an incentive to misrepresent preferences in order to manipulate the voting outcome.

This observation leads to the issue of what other methods there might be that would ensure that individuals have the proper incentives to correctly reveal their true preferences about a public good. Are there any procedures that provide the right incentives to tell the truth about the value of a public good?

It turns out that there is a way to ensure that people will correctly reveal their true value for a public good using a kind of market or "auction" process. Unfortunately, this method also requires a special restriction on preferences, namely, that they be quasilinear. As we've seen earlier, quasilinear preferences imply that we will have a unique optimal amount of the public good, and the issue is to discover what it is. In order to keep things simple, we'll just consider the case where there is one level of the public good to be supplied, and the question is whether to supply it or not.

We can think of a neighborhood association that is considering putting up a street light. The cost of providing the street light is known; say it is \$100. Each person i places some value on the street light, which we denote by v_i. From our analysis of the public goods problem, we know that it is efficient to provide the street light if the *sum* of the values is greater than or equal to the cost:

$$\sum_{i=1}^{n} v_i \geq \$100.$$

One way to decide whether or not to put up the light is to ask each person how much they value the light, with the understanding that their share of the cost will be proportional to their stated value, if the street light gets built. The trouble with this mechanism is that people will have the incentive to free ride: if each person thinks that the other people are willing to pay enough to provide the street light, why should he contribute? It can easily happen that the street light may not get provided even though it would have been efficient to do so.

The problem with this mechanism is that the declaration of how much a person values the good affects how much he or she has to pay, so there is a natural incentive to shade the true value. Let's try to devise a scheme that doesn't have this fault. Suppose that we decide in advance that if the street light is built, everyone will pay a predetermined amount towards its construction, c_i. Then each person will announce his or her value and we'll see whether the sum of the values exceeds the cost. It is convenient to define the term net value, n_i, as the difference between person i's value, v_i, and his or her cost, c_i:

$$n_i = v_i - c_i.$$

Using this definition, we can think of each person announcing his or her net value, and then we simply sum up the net values to see if the total is positive.

The problem with *this* decision mechanism is that it contains an incentive to exaggerate the statements of the true values. If you value the street light only a little more than your cost, you might as well say that you value it a million dollars more—that won't affect what you have to pay, and it will help to ensure that the sum of the values exceeds the cost. Similarly, if you value the light less than your cost, you might as well say it is worth zero to you. Again, that doesn't affect your payment, and it helps to ensure that the street light *won't* get built.

The problem with both of these schemes is that there is no cost for deviating from the truth. And without some incentive to tell the truth about your true value of the public good, there is an incentive to understate or overstate your true value.

Let's think about a way to correct this. The first important idea is that exaggeration doesn't matter if it doesn't affect the social decision. If the sum of the values of everyone else already exceeds the cost, it doesn't matter if you state an exaggerated value. Similarly, if the sum of the values is already less than the cost, it doesn't matter what value you state, as long as the sum of everyone's values remains below the cost.

The only individuals that matter are those who *change* the sum of the values to be greater or less than the cost of the public good. These are called the **pivotal agents**. There might be no pivotal agents, or everyone might be pivotal. The importance of the pivotal agents is that they are the ones who have to have the right incentives to tell the truth; the nonpivotal people don't really matter. Of course, anybody *might* be pivotal, so in ensuring that the pivotal people have the right incentives, we are making sure that everyone has the right incentives in deciding whether or not to tell the truth.

So let's consider the situation of a pivotal person, one who changes the social decision. When the social decision is changed, there will be some harm imposed on the other agents. If the other agents wanted the street light, and this particular pivotal person voted it down, then the other agents have been made worse off by this agent's decision. Similarly, if the other agents didn't want the street light, and this agent cast the dollar vote that provided it, the other agents have been made worse off.

How much worse off are they? Well, if the sum of the net values was positive without person j, say, and person j made the sum go negative, then person j has imposed a total harm of

$$H_j = \sum_{i \neq j} n_i > 0$$

on the other people. This is because the other people wanted the street light, and person j has ensured that they won't get it.

Similarly, if everyone else didn't want the light on the average, so that the sum of their net values was negative, and j made it go positive, then the harm that j imposed is given by

$$H_j = -\sum_{i \neq j} n_i > 0.$$

In order to give person j the right incentives to decide whether or not to be pivotal, we'll just impose this social cost on him. By doing this we make sure that he faces the true social cost of his decision—namely, the harm that he imposes on the other people. This is much like the Pigovian taxes we considered in regulating externalities; in the case of public good provision, the kind of tax is known as a **Groves-Clarke tax**, or the **Clarke tax**, after the economists who first investigated it.

We can now describe the Groves-Clarke mechanism for making public goods decisions.

1. Assign a cost to each agent, c_i, which the individual will have to pay if it is decided that the public good will be provided.

2. Have each agent state a net value s_i. (This may or may not be his or her *true* net value n_i.)

3. If the sum of the stated net values is positive, the public good will be provided; if it is negative, it won't be.

4. Each pivotal person will be required to pay a tax. If person j changes the decision from provision to no provision, the tax on that person will be

$$H_j = \sum_{i \neq j} s_i.$$

If person j changes the decision from no provision to provision, the tax will be

$$H_j = -\sum_{i \neq j} s_i.$$

The tax is *not* paid to the other agents—it is paid to the state. It doesn't matter where the money goes, as long as it doesn't influence anybody else's decision; all that matters is that it be paid by the pivotal people so that they face the proper incentives to tell the truth.

EXAMPLE: An Example of the Clarke Tax

It is convenient to consider a numerical example to see just how the Clarke tax works. Suppose that we have three roommates who have to decide whether or not to acquire a TV that costs \$300. They agree in advance that if they jointly decide to get the TV, then they will each contribute \$100 towards the cost. Persons A and B are willing to pay \$50 each to have the TV present, while person C is willing to pay \$250. This information is summarized in Table 36.2.

Example of the Clarke tax.

Person	Cost share	Value	Net value	Clarke tax
A	100	50	−50	0
B	100	50	−50	0
C	100	250	150	100

Note that the TV provides a positive net value only to person C. Thus if the roommates voted on whether or not to purchase the TV, a majority would be opposed. Nevertheless, it is Pareto efficient to provide the TV since the *sum* of the values (\$350) exceeds the cost (\$300).

Let us consider how the Clarke tax works in this example. Consider person A. The sum of the net values *excluding* person A is 100, and person A's net value is −50. Thus person A is not pivotal. Since person A is made worse off in the net by the provision of the public good, he might have a temptation to exaggerate his bid downward. In order to ensure that the public good is *not* provided, A would have to bid −100 or below. But if he did this, then A would become pivotal, and he would have to pay a Clarke tax equal to the amount the other two people bid: $-50+150=100$. Thus reducing his bid saves him \$50 in net value, but costs him \$100 in taxes—leaving him with a net loss of \$50.

The same thing goes for person B. What about person C? In the example, person C is pivotal—without his bid the public good would not be supplied, and with his bid the good will be supplied. He receives a net value from the public good of \$150, but pays a \$100 tax, leaving him with a total value of his actions of \$50. Would it be worth it for him to increase his bid above his true value? No, because that doesn't change any of his payoffs. Would it be worth it to reduce his bid? No, because that lowers the chance that the public good will be supplied and doesn't change the amount of tax he has to pay. Thus it is in the interest of each of the parties to truthfully reveal his or her net value of the public good. Honesty is the best policy—at least in a situation involving a Clarke tax![3]

36.10 Problems with the Clarke Tax

Despite the nice features of the Clarke tax it does have some problems. The first problem is that it only works with quasilinear preferences. This is because we can't have the amount that you have to pay influence your demand for the public good. It is important that there is a unique optimal level of the public good.

The second problem is that the Clarke tax doesn't really generate a Pareto efficient outcome. The level of the public good will be optimal, but the private consumption could be greater. This is because of the tax collection. Remember that in order to have the correct incentives, the pivotal people must actually pay some taxes that reflect the harm that they do to the other people. And these taxes cannot go to anybody else involved in the decision process, since that might affect their decisions. The taxes have to disappear from the system. And that's the problem—if the taxes actually have to be paid, the private consumption will end up being lower than it could be otherwise, and therefore be Pareto inefficient.

However, the taxes only have to be paid if someone is pivotal. If there are many people involved in the decision, the probability that any one person is pivotal may not be very large; thus the tax collections might typically be expected to be rather small.

The final problem concerns the equity and efficiency tradeoff inherent in the Clarke tax. Since the payment scheme must be fixed in advance, there will generally be situations where some people will be made worse off by providing the public good, even though the Pareto efficient *amount* of the public good will be provided. To say that it is Pareto preferred to provide the public good is to say that there is *some* payment scheme for which everyone is better off having the public good provided than not having it. But this doesn't mean that for an *arbitrary* payment scheme everyone will be better off. The Clarke tax ensures that if everyone *could* be better off

[3] For a more detailed discussion of the Clarke tax, see N. Tideman and G. Tullock, "A New and Superior Process for Making Social Choices," *Journal of Political Economy*, 84, December 1976, pp. 1145–59.

having the good provided, then it will be provided. But that doesn't imply that everyone will actually be better off.

It would be nice if there were a scheme that not only determined whether or not to provide the public good, but also the Pareto efficient way to pay for it—that is, a payment plan that makes everyone better off. However, it does not appear that such a general plan is available.

Summary

1. Public goods are goods for which everyone must "consume" the same amount, such as national defense, air pollution, and so on.

2. If a public good is to be provided in some fixed amount or not provided at all, then a necessary and sufficient condition for provision to be Pareto efficient is that the sum of the willingnesses to pay (the reservation prices) exceeds the cost of the public good.

3. If a public good can be provided in a variable amount, then the necessary condition for a given amount to be Pareto efficient is that the sum of the marginal willingnesses to pay (the marginal rates of substitution) should equal the marginal cost.

4. The free rider problem refers to the temptation of individuals to let others provide the public goods. In general, purely individualistic mechanisms will not generate the optimal amount of a public good because of the free rider problem.

5. Various collective decision methods have been proposed to determine the supply of a public good. Such methods include the command mechanism, voting, and the Clarke tax.

REVIEW QUESTIONS

1. Consider an auction in which people will bid in turn, where each bid has to be at least a dollar higher than the previous bid, and the item is sold to the person who bids the highest. If the value of the good to person i is v_i, what will be the winning bid? Which person will get the good?

2. Consider a sealed bid auction among n people for some good. Let v_i be the value of the good to person i. Prove that if the good is sold to the highest bidder at the *second* highest price bid, it will be in each player's interest to tell the truth.

3. Suppose that 10 people live on a street and that each of them is willing to pay \$2 for each extra streetlight, regardless of the number of streetlights provided. If the cost of providing x streetlights is given by $c(x) = x^2$, what is the Pareto efficient number of streetlights to provide?

APPENDIX

Let's solve the maximization problem that determines the Pareto efficient allocations of the public good:

$$\max_{x_1, x_2, G} u_1(x_1, G)$$

$$\text{such that } u_2(x_2, G) = \overline{u}_2$$
$$x_1 + x_2 + c(G) = w_1 + w_2.$$

We set up the Lagrangian:

$$L = u_1(x_1, G) - \lambda[u_2(x_2, G) - \overline{u}_2] - \mu[x_1 + x_2 + c(G) - w_1 - w_2]$$

and differentiate with respect to x_1, x_2, and G to get

$$\begin{aligned}
\frac{\partial L}{\partial x_1} &= \frac{\partial u_1(x_1, G)}{\partial x_1} - \mu = 0 \\
\frac{\partial L}{\partial x_2} &= -\lambda \frac{\partial u_2(x_2, G)}{\partial x_2} - \mu = 0 \\
\frac{\partial L}{\partial G} &= \frac{\partial u_1(x_1, G)}{\partial G} - \lambda \frac{\partial u_2(x_2, G)}{\partial G} - \mu \frac{\partial c(G)}{\partial G} = 0.
\end{aligned}$$

If we divide the third equation by μ and rearrange, we get

$$\frac{1}{\mu} \frac{\partial u_1(x_1, G)}{\partial G} - \frac{\lambda}{\mu} \frac{\partial u_2(x_2, G)}{\partial G} = \frac{\partial c(G)}{\partial G}. \tag{36.2}$$

Now solve the first equation for μ to get

$$\mu = \frac{\partial u_1(x_1, G)}{\partial x_1},$$

and solve the second equation for μ/λ to get

$$\frac{\mu}{\lambda} = -\frac{\partial u_2(x_2, G)}{\partial x_2}.$$

Substitute these two equations into equation (36.2) to find

$$\frac{\partial u_1(x_1, G)/\partial G}{\partial u_1(x_1, G)/\partial x_1} + \frac{\partial u_2(x_2, G)/\partial G}{\partial u_2(x_2, G)/\partial x_2} = \frac{\partial c(G)}{\partial G},$$

which is just

$$MRS_1 + MRS_2 = MC(G)$$

as given in the text.