# CHAPTER 10

# CONSUMERS' SURPLUS

When the economic environment changes a consumer may be made better off or worse off. Economists often want to measure how consumers are affected by changes in the economic environment, and have developed several tools to enable them to do this.

The classical measure of welfare change examined in elementary courses is consumer's surplus. However, consumer's surplus is an exact measure of welfare change only in special circumstances. In this chapter we describe some more general methods for measuring welfare change. These more general methods will include consumer's surplus as a special case.

## 10.1 Compensating and equivalent variations

Let us first consider what an "ideal" measure of welfare change may be. At the most fundamental level, we would like to have a measure of the change in utility resulting from some policy. Suppose that we have two budgets, $(\mathbf{p}^0, m^0)$ and $(\mathbf{p}', m')$, that measure the prices and incomes that a given consumer would face under two different policy regimes. It is convenient to think of $(\mathbf{p}^0, m^0)$ as being the status quo and $(\mathbf{p}', m')$ as being a proposed change, although this is not the only interpretation.

Then the obvious measure of the welfare change involved in moving from $(\mathbf{p}^0, m^0)$ to $(\mathbf{p}', m')$ is just the difference in indirect utility:

$$v(\mathbf{p}', m') - v(\mathbf{p}^0, m^0).$$

If this utility difference is positive, then the policy change is worth doing, at least as far as this consumer is concerned; and if it is negative, the policy change is not worth doing.

This is about the best we can do in general; utility theory is purely ordinal in nature and there is no unambiguously right way to quantify utility changes. However, for some purposes it is convenient to have monetary measure of changes in consumer welfare. Perhaps the policy analyst wants to have some rough idea of the magnitude of the welfare change for purposes of establishing priorities. Or perhaps the policy analyst wants to compare the benefits and costs accruing to different consumers. In circumstances such as these, it is convenient to choose a "standard" measure of utility differences. A reasonable measure to adopt is the (indirect) money metric utility function described in Chapter 7, page 109.

Recall that $\mu(\mathbf{q}; \mathbf{p}, m)$ measures how much income the consumer would need at prices $\mathbf{q}$ to be as well off as he or she would be facing prices $\mathbf{p}$ and having income $m$. That is, $\mu(\mathbf{q}; \mathbf{p}, m)$ is defined to be $e(\mathbf{q}, v(\mathbf{p}, m))$. If we adopt this measure of utility, we find that the above utility difference becomes

$$\mu(\mathbf{q}; \mathbf{p}', m') - \mu(\mathbf{q}; \mathbf{p}^0, m^0).$$

It remains to choose the base prices $\mathbf{q}$. There are two obvious choices: we may set $\mathbf{q}$ equal to $\mathbf{p}^0$ or to $\mathbf{p}'$. This leads to the following two measures for the utility difference:

$$\begin{aligned} EV &= \mu(\mathbf{p}^0; \mathbf{p}', m') - \mu(\mathbf{p}^0; \mathbf{p}^0, m^0) = \mu(\mathbf{p}^0; \mathbf{p}', m') - m^0 \\ CV &= \mu(\mathbf{p}'; \mathbf{p}', m') - \mu(\mathbf{p}'; \mathbf{p}^0, m^0) = m' - \mu(\mathbf{p}'; \mathbf{p}^0, m^0). \end{aligned} \quad (10.1)$$

The first measure is known as the **equivalent variation**. It uses the current prices as the base and asks what income change at the current prices would be equivalent to the proposed change in terms of its impact on utility. The second measure is known as the **compensating variation**. It uses the new prices as the base and asks what income change would be necessary to compensate the consumer for the price change. (Compensation takes place after some change, so the compensating variation uses the after-change prices.)

Both of these numbers are reasonable measures of the welfare effect of a price change. Their magnitudes will generally differ because the value of a dollar will depend on what the relevant prices are. However, their sign will always be the same since they both measure the same utility differences, just using a different utility function. Figure 10.1 depicts an example of the equivalent and compensating variations in the two-good case.

Which measure is the most appropriate depends on the circumstances involved and what question you are trying to answer. If you are trying to arrange for some compensation scheme at the new prices, then the compensating variation seems reasonable. However, if you are simply trying to get a reasonable measure of "willingness to pay," the equivalent variation is probably better. This is so for two reasons. First, the equivalent variation measures the income change at current prices, and it is much easier for decision makers to judge the value of a dollar at current prices than at some hypothetical prices. Second, if we are comparing more than one proposed policy change, the compensating variation uses different base prices for each new policy while the equivalent variation keeps the base prices fixed at the status quo. Thus, the equivalent variation is more suitable for comparisons among a variety of projects.

Given, then, that we accept the compensating and equivalent variations as reasonable indicators of utility change, how can we measure them in practice? This is equivalent to the question: how we can measure $\mu(\mathbf{q}; \mathbf{p}, m)$ in practice?

We have already answered this question in our study of integrability theory in Chapter 8. There we investigated how to recover the preferences represented by $\mu(\mathbf{q}; \mathbf{p}, m)$ by observing the demand behavior $\mathbf{x}(\mathbf{p}, m)$. Given any observed demand behavior one can solve the integrability equations, at least in principle, and derive the associated money metric utility function.

We have seen in Chapter 8 how to derive the money metric utility function for several common functional forms for demand function including linear, log-linear, semilog, and so on. In principle, we can do similar calculations for any demand function that satisfies the integrability conditions.

However, in practice it is usually simpler to make the parametric specification in the *other* direction: first specify a functional form for the indirect utility function and then derive the form of the demand functions by Roy's identity. After all, it is usually a lot easier to differentiate a function than to solve a system of partial differential equations!

If we specify a parametric form for the indirect utility function, and then derive the associated demand equations, then estimating the parameters of the demand function immediately gives us the parameters of the underlying utility function. We can derive the money metric utility function—and the compensating and equivalent variations—either algebraically or numerically without much difficulty once we have the relevant parameters. See Chapter 12 for a more detailed description of this approach.

Of course this approach only makes sense if the estimated parameters satisfy the various restrictions implied by the optimization model. We may want to test these restrictions, to see if they are plausible in our particular empirical example, and, if so, estimate the parameters subject to these

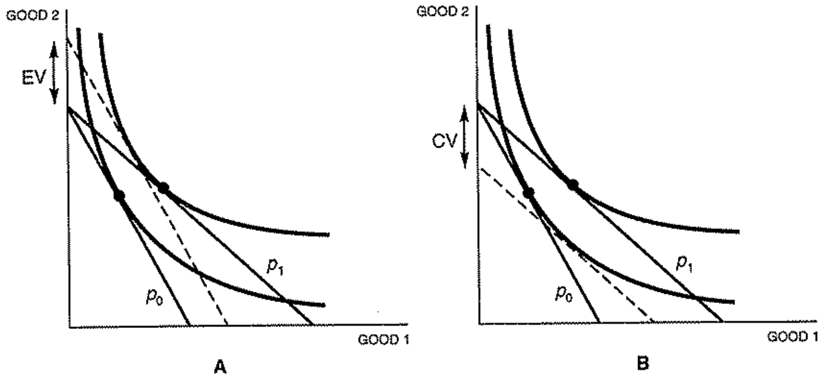


**Figure 10.1** **Equivalent variation and compensating variation.** In this diagram the price of good 1 decreases from $p_0$ to $p_1$. Panel $A$ depicts the equivalent variation in income—how much additional money is needed at the original price $p_0$ to make the consumer as well off as she would be facing $p_1$. Panel $B$ depicts the compensating variation in income—how much money should be taken away from the consumer to leave him as well off as he was facing price $p_0$.

restrictions.

In summary: the compensating and equivalent variations are in fact observable if the demand functions are observable and if the demand functions satisfy the conditions implied by utility maximization. The observed demand behavior can be used to construct a measure of welfare change, which can then be used to analyze policy alternatives.

## 10.2 Consumer's surplus

The classic tool for measuring welfare changes is **consumer's surplus.** If $x(p)$ is the demand for some good as a function of its price, then the consumer's surplus associated with a price movement from $p^0$ to $p'$ is

$$CS = \int_{p^0}^{p'} x(t)\,dt.$$

This is simply the area to the left of the demand curve between $p^0$ and $p^1$. It turns out that when the consumer's preferences can be represented by a quasilinear utility function, consumer's surplus is an exact measure of welfare change. More precisely, when utility is quasilinear, the compensating variation equals the equivalent variation, and both are equal to the consumer's surplus integral. For more general forms of the utility function, the compensating variation will be different from the equivalent variation and consumer's surplus will not be an exact measure of welfare change. However, even when utility is not quasilinear, consumer's surplus may be a reasonable approximation to more exact measures. We investigate these ideas further below.

## 10.3 Quasilinear utility

Suppose that there exists a monotonic transformation of utility that has the form

$$U(x_0, x_1, \ldots, x_k) = x_0 + u(x_1, \ldots, x_k).$$

Note that the utility function is linear in one of the goods, but (possibly) nonlinear in the other goods. For this reason we call this a **quasilinear utility function.**

In this section we will focus on the special case where $k = 1$, so that the utility function takes the form $x_0 + u(x_1)$, although everything that we say will work if there are an arbitrary number of goods. We will assume that $u(x_1)$ is a strictly concave function.

Let us consider the utility maximization problem for this form of utility:

$$\max_{x_0, x_1} x_0 + u(x_1)$$
$$\text{such that } x_0 + p_1 x_1 = m.$$

It is tempting to substitute into the objective function and reduce this problem to the unconstrained maximization problem

$$\max_{x_1} \ u(x_1) + m - p_1 x_1.$$

This has the obvious first-order condition

$$u'(x_1) = p_1,$$

which simply requires that the marginal utility of consumption of good 1 be equal to its price.

By inspection of the first-order condition, the demand for good 1 is only a function of the price of good 1, so we can write the demand function as $x_1(p_1)$. The demand for good 0 is then determined from the budget constraint, $x_0 = m - p_1 x_1(p_1)$. Substituting these demand functions into the utility function gives us the indirect utility function

$$V(p_1, m) = u(x_1(p_1)) + m - p_1 x_1(p_1) = v(p_1) + m,$$

where $v(p_1) = u(x_1(p_1)) - p_1 x_1(p_1)$.

This approach is perfectly fine, but it hides a potential problem. Upon reflection, it is clear that demand for good 1 can't be independent of income for *all* prices and income levels. If income is small enough, the demand for good 1 must be constrained by income.

Suppose that we write the utility maximization problem in a way that explicitly recognizes the nonnegativity constraint on $x_0$:

$$\max_{x_0, x_1} u(x_1) + x_0$$
$$\text{such that } p_1 x_1 + x_0 = m$$
$$x_0 \geq 0.$$

Now we see that we will get two classes of solutions, depending on whether $x_0 > 0$ or $x_0 = 0$. If $x_0 > 0$, we have the solution that we described above—the demand for good 1 depends only on the price of good 1 and is independent of income. If $x_0 = 0$, then indirect utility will just be given by $u(m/p_1)$.

Think of starting the consumer at $m = 0$ and increasing income by a small amount. The increment in utility is then $u'(m/p_1)/p_1$. If this is larger than 1, then the consumer is better off spending the first dollar of income on good 1 rather than good 0. We continue to spend on good 1 until the marginal utility of an extra dollar spent on that good just equals 1; that is, until the marginal utility of consumption equals price. All additional income will then be spent on the $x_0$ good.

The quasilinear utility function is often used in applied welfare economics since it has such a simple demand structure. Demand only depends on price—at least for large enough levels of income—and there are no income effects to worry about. This turns out to simplify the analysis of market equilibrium. You should think of this as modeling a situation where the demand for a good isn't very sensitive to income. Think of your demand for paper or pencils: how much would your demand change as your income changes? Most likely, any increases in income would go into consumption of other goods.

Furthermore, with quasilinear utility the integrability problem is very simple. Since the inverse demand function is given by $p_1(x_1) = u'(x_1)$, it follows that the utility associated with a particular level of consumption of good 1 can be recovered from the inverse demand curve by a simple integration:

$$u(x_1) - u(0) = \int_0^{x_1} u'(t)\, dt = \int_0^{x_1} p(t)\, dt.$$

The total utility from choosing to consume $x_1$ will consist of the utility from the consumption good 1, plus the utility from the consumption of good 0:

$$u(x_1(p_1)) + m - px_1(p_1) = \int_0^{x_1} p(t)\, dt + m - px_1(p_1).$$

If we disregard the constant $m$, the expression on the right-hand side of this equation is simply the area *under* the demand curve for good 1 minus the expenditure on good 1. Alternatively, this is the area to the left of the demand curve.

Another way to see this is to start with the indirect utility function, $v(p_1) + m$. By Roy's law, $x_1(p_1) = -v'(p_1)$. Integrating this equation, we have

$$v(p_1) + m = \int_{p_1}^{\infty} x_1(t)\, dt + m.$$

This is the area to the left of the demand curve down to the price $p_1$, which is just another way of describing the same area as described in the last paragraph.

## 10.4 Quasilinear utility and money metric utility

Suppose that utility takes the quasilinear form $u(x_1) + x_0$. We have seen that for such a utility function the demand function $x_1(p_1)$ will be independent of income. We saw above that we could recover an indirect utility function consistent with this demand function simply by integrating with respect to $p_1$.

Of course, any monotonic transformation of this indirect utility function is also an indirect utility function that describes the consumer's behavior. If the consumer makes choices that maximize consumer's surplus, then he also maximizes the square of consumer's surplus.

We saw above that the money metric utility function was a particularly convenient utility function for many purposes. It turns out that for quasilinear utility function, the integral of demand is essentially the money metric utility function.

This follows simply by writing down the integrability equations and verifying that consumer's surplus is the solution to these equations. If $x_1(p_1)$ is the demand function, the integrability equation is

$$\frac{d\mu(t; q, m)}{dt} = x_1(t)$$
$$\mu(q; q, m) = m.$$

It can be verified by direct calculation that the solution to these equations is given by

$$\mu(p; q, m) = \int_{p}^{q} x_1(t)\, dt + m.$$

The expression on the right is simply the consumer's surplus associated with a price change from $p$ to $q$.

For this form of the money metric utility function the compensating and equivalent variations take the form

$$EV = \mu(p^0; p', m') - \mu(p^0; p^0, m^0) = A(p^0, p') + m' - m^0$$
$$CV = \mu(p'; p', m') - \mu(p'; p^0, m^0) = A(p^0, p') + m' - m^0.$$

In this special case the compensating and equivalent variations coincide. It is not hard to see the intuition behind this result. Since the compensation function is linear in income the value of an extra dollar—the marginal utility of income—is independent of price. Hence the value of a compensating or equivalent change in income is independent of the prices at which the value is measured.

## 10.5 Consumer's surplus as an approximation

We have seen that consumer's surplus is an exact measure of the compensating and equivalent variation only when the utility function is quasilinear. However, it may be a reasonable approximation in more general circumstances.

For example, consider a situation where only the price of good 1 changes from $p^0$ to $p'$ and income is fixed at $m = m^0 = m'$. In this case, we can use the equation (10.1) and the fact that $\mu(\mathbf{p}; \mathbf{p}, m) \equiv m$ to write

$$EV = \mu(p^0; p', m) - \mu(p^0; p^0, m) = \mu(p^0; p', m) - \mu(p'; p', m)$$
$$CV = \mu(p'; p', m) - \mu(p'; p^0, m) = \mu(p^0; p^0, m) - \mu(p'; p^0, m).$$

We have written these expressions as a function of $p$ alone, since all other prices are assumed to be fixed. Letting $u^0 = v(p^0, m)$ and $u' = v(p', m)$ and using the definition of the money metric utility function given in Chapter 7, page 109, we have

$$EV = e(p^0, u') - e(p', u')$$
$$CV = e(p^0, u^0) - e(p', u^0).$$

Finally, using the fact that the Hicksian demand function is the derivative of the expenditure function, so that $h(p, u) \equiv \partial e / \partial p$, we can write these expressions as

$$EV = e(p^0, u') - e(p', u') = \int_{p'}^{p^0} h(p, u')\, dp$$
$$CV = e(p^0, u^0) - e(p', u^0) = \int_{p'}^{p^0} h(p, u^0)\, dp. \quad (10.2)$$

It follows from these expressions that the compensating variation is the integral of the *Hicksian* demand curve associated with the initial level of utility, and the equivalent variation is the integral of the Hicksian demand curve associated with the final level of utility. The correct measure of welfare *is* an integral of a demand curve—but you have to use the Hicksian demand curve rather than the Marshallian demand curve.

However, we can use (10.2) to derive a useful bound. The Slutsky equation tells us that

$$\frac{\partial h(p,u)}{\partial p} = \frac{\partial x(p,m)}{\partial p} + \frac{\partial x(p,m)}{\partial m} x(p,m).$$

If the good in question is a normal good, the derivative of the Hicksian demand curve will be larger than the derivative of the Marshallian demand curve, as depicted in Figure 10.2.

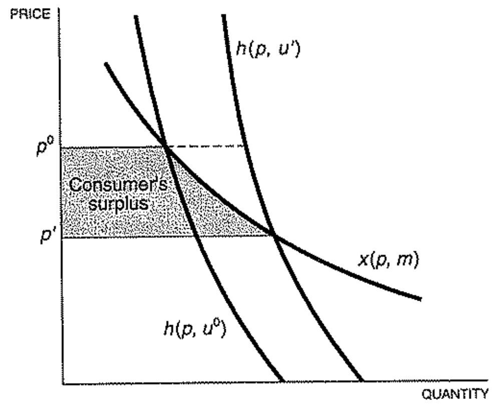


**Figure 10.2** **Bounds on consumer's surplus.** For a normal good, the Hicksian demand curves are steeper than the Marshallian demand curve. Hence, the area to the left of the Marshallian demand curve is bounded by the areas under the Hicksian demand curves.

It follows that the area to the left of the Hicksian demand curves will bound the area to the left of the Marshallian demand curve. In the case depicted, $p^0 > p^1$ so all of the areas are positive. It follows that $EV >$ consumer's surplus $> CV$.

## 10.6 Aggregation

The above relationships among compensating variation, equivalent variation, and consumer's surplus all hold for a single consumer. Here we investigate some issues involving many consumers.

We have seen in Chapter 9, page 153, that aggregate demand for a good will be a function of price and aggregate income only when the indirect utility function for agent $i$ has the Gorman form

$$v_i(\mathbf{p}, m_i) = a_i(\mathbf{p}) + b(\mathbf{p})m_i.$$

In this case the aggregate demand function for each good will be derived from an aggregate indirect utility function that has the form

$$V(\mathbf{p}, M) = \sum_{i=1}^{n} a_i(\mathbf{p}) + b(\mathbf{p})M,$$

where $M = \sum_{i=1}^{n} m_i$.

We saw above that the indirect utility function associated with quasilinear preferences has a form

$$v_i(\mathbf{p}) + m_i.$$

This is clearly a special case of the Gorman form with $b(\mathbf{p}) \equiv 1$. Hence, the aggregate indirect utility function that will generate aggregate demand is simply $V(\mathbf{p}) + M = \sum_{i=1}^{n} v_i(\mathbf{p}) + \sum_{i=1}^{n} m_i$.

How does this relate to aggregate consumers' surplus? Roy's law shows that the function $v_i(p)$ is given by

$$v_i(p) = \int_p^\infty x_i(t)\,dt.$$

It follows that

$$V(\mathbf{p}) = \sum_{i=1}^{n} v_i(p) = \sum_{i=1}^{n} \int_p^\infty x_i(t)\,dt = \int_p^\infty \sum_{i=1}^{n} x_i(t)\,dt.$$

That is, the utility function that generates the aggregate demand function is simply the integral of the aggregate demand function.

If all consumers have quasilinear utility functions, then the aggregate demand function will appear to maximize aggregate consumer's surplus. However, it is not entirely obvious that aggregate consumer's surplus is appropriate for welfare comparisons. Why should the unweighted sum of a particular representation of utility be a useful welfare measure? We examine this issue in Chapter 13, page 225. As it turns out, aggregate consumers' surplus is the appropriate welfare measure for quasilinear utility, but this case is rather special. In general, aggregate consumers' surplus will not be an exact welfare measure. However, it is often used as an approximate measure of consumer welfare in applied work.

## 10.7 Nonparametric bounds

We've seen how Roy's identity can be used to calculate the demand function given a parametric form for indirect utility. Integrability theory can be used to calculate a parametric form for the money metric utility function if we are given a parametric form for the demand function. However, each of these operations requires that we specify a parametric form for either the demand function or the indirect utility function.

It is of interest to ask how far we can go without having to specify a parametric form. As it turns out it is possible to derive tight nonparametric bounds on the money metric utility function in an entirely nonparametric way.

We've seen in the discussion of recoverability in Chapter 8 that it is possible to construct sets of consumption bundles that are "revealed preferred" or "revealed worse" than a given consumption bundle. These sets can be thought of as inner and outer bounds to the consumer's preferred set.

Let $NRW(x_0)$ be the set of points "not revealed worse" than $x_0$. This is just the complement of the set $RW(x_0)$. We know from Chapter 8 that the true preferred set associated with $x_0$, $P(x_0)$, must contain $RP(x_0)$ and be contained in the set of points $NRW(x_0)$.

We illustrate this situation in Figure 10.3. In order not to clutter the diagram, we've left out many of the budget lines and observed choices and have only depicted $RP(x_0)$ and $RW(x_0)$. We've also shown the "true" indifference curve through $x_0$. By definition, the money metric utility of $x_0$ is defined by

$$m(\mathbf{p}, x_0) = \min_{\mathbf{x}} \mathbf{px}$$
$$\text{such that } u(\mathbf{x}) \geq u(\mathbf{x}_0).$$

This is the same problem as

$$m(\mathbf{p}, \mathbf{x}_0) = \min_{\mathbf{x}} \mathbf{px}$$
$$\text{such that } x \text{ in } P(\mathbf{x}_0).$$

Define $m^+(\mathbf{p}, \mathbf{x}_0)$ and $m^-(\mathbf{p}, \mathbf{x}_0)$ by

$$m^+(\mathbf{p}, \mathbf{x}_0) = \min_{\mathbf{x}} \mathbf{px}$$
$$\text{such that } x \text{ in } NRW(\mathbf{x}_0),$$

and

$$m^-(\mathbf{p}, \mathbf{x}_0) = \min_{\mathbf{x}} \mathbf{px}$$
$$\text{such that } x \text{ in } RP(\mathbf{x}_0).$$

Since $NRW(x_0) \supset P(x_0) \supset RP(x_0)$, it follows from the standard sort of argument that $m^+(\mathbf{p}, \mathbf{x}_0) \geq m(\mathbf{p}, \mathbf{x}_0) \geq m^-(\mathbf{p}, \mathbf{x}_0)$. Hence, the **overcompensation function**, $m^+(\mathbf{p}, \mathbf{x}_0)$, and the **undercompensation function**, $m^-(\mathbf{p}, \mathbf{x}_0)$, bound the true compensation function, $m(\mathbf{p}, \mathbf{x}_0)$.

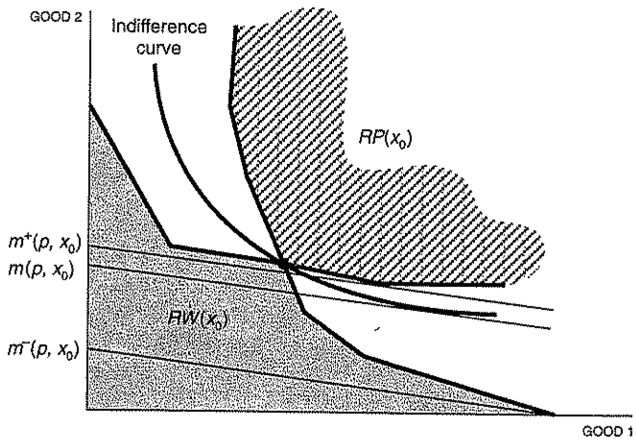


**Figure 10.3** **Bounds on the money metric utility.** The true preferred set, $P(\mathbf{x}_0)$, contains $RP(\mathbf{x}_0)$ and is contained in $NRW(\mathbf{x}_0)$. Hence the minimum expenditure over $P(\mathbf{x}_0)$ lies between the two bounds, as illustrated.

## Notes

The concepts of compensating and equivalent variation and their relationship to consumer's surplus is due to Hicks (1956). See Willig (1976) for tighter bounds on consumer's surplus. The nonparametric bounds on the money metric utility function are due to Varian (1982a).

## Exercises

10.1. Suppose that utility is quasilinear. Show that the indirect utility function is a convex function of prices.

10.2. Ellsworth's utility function is $U(x, y) = \min\{x, y\}$. Ellsworth has \$150 and the price of $x$ and the price of $y$ are both 1. Ellsworth's boss is thinking of sending him to another town where the price of $x$ is 1 and the price of $y$ is 2. The boss offers no raise in pay. Ellsworth, who understands compensating and equivalent variation perfectly, complains bitterly. He says that although he doesn't mind moving for its own sake and the new town is just as pleasant as the old, having to move is as bad as a cut in pay of \$$A$. He also says he wouldn't mind moving if when he moved he got a raise of \$$B$. What are $A$ and $B$ equal to?

# CHAPTER 24

# EXTERNALITIES

When the actions of one agent directly affect the environment of another agent, we will say that there is an **externality**. In a **consumption externality** the utility of one consumer is directly affected by the actions of another consumer. For example, some consumers may be affected by other agents' consumption of tobacco, alcohol, loud music, and so on. Consumers might also be adversely affected by firms who produce pollution or noise.

In **production externality** the production set of one firm is directly affected by the actions of another agent. For example, the production of smoke by a steel mill may directly affect the production of clean clothes by a laundry, or the production of honey by a beekeeper might directly affect the level of output of an apple orchard next door.

In this chapter we explore the economics of externalities. We find that in general market equilibria will be inefficient in the presence of externalities. This naturally leads to an examination of various suggestions for alternative ways to allocate resources that lead to efficient outcomes.

The First Theorem of Welfare Economics does not hold in the presence of externalities. The reason is that there are things that people care about that are not priced. Achieving an efficient allocation in the presense of externalities essentially involves making sure that agents face the correct prices for their actions.

## 24.1 An example of a production externality

Suppose that we have two firms. Firm 1 produces an output $x$ which it sells in a competitive market. However, the production of $x$ imposes a cost $e(x)$ on firm 2. For example, suppose the technology is such that $x$ units of output can only be produced by generating $x$ units of pollution, and this pollution harms firm 2.

Letting $p$ be the price of output, the profits of the two firms are given by

$$\pi_1 = \max_x \; px - c(x)$$
$$\pi_2 = -e(x).$$

We assume that both cost functions are increasing and convex as usual. (It may be that firm 2 receives profits from some production activity, but we ignore this for simplicity.)

The equilibrium amount of output, $x_q$, is given by $p = c'(x_q)$. However, this output is too large from a social point of view. The first firm takes account of the **private costs**—the costs that it imposes on itself—but it ignores the **social costs**—the private cost plus cost that it imposes on the other firm.

In order to determine the efficient amount of output, we ask what would happen if the two firms merged so as to **internalize** the externality. In this case the merged firm would maximize total profits

$$\pi = \max_x \; px - c(x) - e(x),$$

and this problem has first-order condition

$$p = c'(x_e) + e'(x_e). \tag{24.1}$$

The output $x_e$ is an efficient amount of output; it is characterized by price being equal to marginal *social* cost.

## 24.2 Solutions to the externalities problem

There have been several solutions proposed to solve the inefficiency of externalities.

### Pigovian taxes

According to this view, firm 1 simply faces the wrong price for its action, and a corrective tax can be imposed that will lead to efficient resource allocation. Corrective taxes of this sort are known as **Pigovian taxes**.

Suppose, for example, that the firm faced a tax on its output in amount $t$. Then the first-order condition for profit maximization becomes

$$p = c'(x) + t.$$

Under our assumption of a convex cost function we can set $t = e'(x_e)$, which leads the firm to choose $x = x_e$, as determined in equation (24.1). Even if the cost function were not convex, we could simply impose a nonlinear tax of $e(x)$ on firm 1, thus leading it to internalize the cost of the externality.

The problem with this solution is that it requires that the taxing authority know the externality cost function $e(x)$. But if the taxing authority knows this cost function, it might as well just tell the firm how much to produce in the first place.

## Missing markets

According to this view, the problem is that firm 2 cares about the pollution generated by firm 1 but has no way to influence it. Adding a market for firm 2 to express its demand for pollution—or for a reduction of pollution—will provide a mechanism for efficient allocation.

In our model when $x$ units of output are produced, $x$ units of pollution are unavoidably produced. If the market price of pollution is $r$, then firm 1 can decide how much pollution it wants to sell, $x_1$, and firm 2 can decide how much it wants to buy, $x_2$. The profit maximization problems become

$$\pi_1 = \max_{x_1} \; px_1 + rx_1 - c(x_1)$$
$$\pi_2 = \max_{x_2} \; -rx_2 - e(x_2).$$

The first-order conditions are

$$p + r = c'(x_1)$$
$$-r = e'(x_2).$$

When demand for pollution equals the supply of pollution, we have $x_1 = x_2$, and these first-order conditions become equivalent to those given in (24.1). Note that the equilibrium price of pollution, $r$, will be a negative number. This is natural, since pollution is a "bad" not a good.

More generally, suppose that pollution and output are not necessarily produced in a one-to-one ratio. If firm 1 produces $x$ units of output and $y$ units of pollution, then it pays a cost of $c(x, y)$. Presumably increasing $y$ from zero lowers the cost of production of $x$; otherwise, there wouldn't be any problem.

In the absence of any mechanism to control pollution, the profit maximization problem of firm 1 is

$$\max_{x,y} \; px - c(x,y),$$

which has first-order conditions

$$p = \frac{\partial c(x,y)}{\partial x}$$
$$0 = \frac{\partial c(x,y)}{\partial y}.$$

Firm 1 will equate the price of pollution to its marginal cost. In this case the price of pollution is zero, so firm 1 will pollute up to the point where the costs of production are minimized.

Now we add a market for pollution. Again, let $r$ be the cost per unit of pollution and $y_1$ and $y_2$ the supply and demand by firms 1 and 2. The maximization problems are

$$\pi_1 = \max_{x,y_1} \; px + ry_1 - c(x,y_1)$$
$$\pi_2 = \max_{y_2} -ry_2 - e(y_2).$$

The first-order conditions are

$$p = \frac{\partial c(x,y_1)}{\partial x}$$
$$r = \frac{\partial c(x,y_1)}{\partial y_1}$$
$$-r = \frac{\partial e(y_2)}{\partial y_2}.$$

Equating supply and demand, so $y_1 = y_2$, we have the first-order conditions for an efficient level of $x$ and $y$.

The problem with this solution is that the markets for pollution may be very thin. In the case depicted there are only two firms. There is no particular reason to think that such a market will behave competitively.

## Property rights

According to this view, the basic problem is that property rights are not conducive to full efficiency. If both technologies are operated by one firm, we have seen that there is no problem. However, we will see that there is a market signal that will encourage the agents to determine an efficient pattern of property rights.

If the externality of one firm adversely affects the operation of another, it always will pay one firm to buy out the other. It is clear that by coordinating the actions of both firms one can always produce more profits than by acting separately. Therefore, one firm could afford to pay the other firm its market value (in the presence of the externality) since its value when the externality is optimally adjusted would exceed this current market value. This argument shows the market mechanism itself provides signals to adjust property rights to internalize externalities.

We have already established this claim in some generality in the proof of the First Welfare Theorem in Chapter 18, page 345. The argument there shows that if an allocation is not Pareto efficient, then there is some way that aggregate profits can be increased. A careful examination of theorem shows that all that is necessary is that all goods that *consumers* care about are priced, or, equivalently, that consumers' preferences depend only on their own consumption bundles. There can be arbitrary sorts of production externalities and the proof still goes through up to the last line, where we show that aggregate profits at the Pareto dominating allocation exceed aggregate profits at the original allocation. If there are no production externalities, this is a contradiction. If production externalities are present, then this argument shows that there is some alternative production plan that increases aggregate profits—hence there is a market incentive for one firm to buy out the others, coordinate their production plans, and internalize the externality.

Essentially, the firm grows until it internalizes all relevant production externalities. This works well for some sorts of externalities, but not for all. For example, it doesn't deal very well with the case of consumption externalities, or the case of externalities that are public goods.

## 24.3 The compensation mechanism

We argued above that Pigovian taxes were not adequate in general to solve externalities due to the information problem: the taxing authority in general can't be expected to know the costs imposed by the externalities. However, it may be that the agents who generate the externalities have a reasonably good idea of the costs they impose. If so, there is a relatively simple scheme to internalize the externalities.

The scheme involves setting up a market for the externality, but it does so in a way that encourages the firms to correctly reveal the costs they impose on the other. Here is how the method works.

**Announcement stage.** Firm $i = 1, 2$ names a Pigovian tax $t_i$ which may or may not be the efficient level of such a tax.

**Choice stage.** If firm 1 produces $x$ units of output, then it has to pay a tax $t_2 x$, and firm 2 receives compensation in the amount of $t_1 x$. In addition, each firm pays penalty depending on the *difference* between their two announced tax rates.

The exact form of the penalty is irrelevant for our purposes; all that matters is that it is zero when $t_1 = t_2$ and positive otherwise. For purposes of exposition, we choose a quadratic penalty. In this case, the final payoffs to firm 1 and firm 2 are given by

$$\pi_1 = \max_x \; px - c(x) - t_2 x - (t_1 - t_2)^2$$
$$\pi_2 = t_1 x - e(x) - (t_2 - t_1)^2.$$

We want to show that the equilibrium outcome to this game involves an efficient level of production of the externality. In order to do this, we have to think a bit about what constitutes a reasonable equilibrium notion for this game. Since the game has two stages, it is reasonable to demand a **subgame perfect** equilibrium—that is, an equilibrium in which each firm takes into account the repercussions of its first-stage choices on the outcomes in the second stage. See Chapter 15, page 275.

As usual, we solve this game by looking at the second stage first. Consider the output choice in the second stage. Firm 1 will choose $x$ to satisfy the condition

$$p = c'(x) + t_2 \qquad (24.2).$$

For each choice of $t_2$, there will be some optimal choice of $x(t_2)$. If $c''(x) > 0$, then it is straightforward to show that $x'(t_2) < 0$.

In the first stage, each firm will choose tax rates so as to maximize their profits. For firm 1, the choice is simple: if firm 2 chooses $t_2$, then firm 1 also wants to choose

$$t_1 = t_2. \qquad (24.3)$$

To check this, just differentiate firm 1's profit function with respect to $t_1$.

Things are a little trickier for firm 2, since it has to recognize that its choice of $t_2$ affects firm 1's output through the function $x(t_2)$. Differentiating firm 2's profit function, taking account of this influence, we have

$$\pi_2'(t_2) = (t_1 - e'(x))x'(t_2) - (t_2 - t_1) = 0. \qquad (24.4)$$

Putting (24.2), (24.3), and (24.4) together, we find

$$p = c'(x) + e'(x),$$

which is the condition for efficiency.

This method works by setting opposing incentives for the two agents. It is clear from (24.3) that agent 1 always has an incentive to match the announcement of agent 2. But consider agent 2's incentive. If agent 2 thinks that agent 1 will propose a large compensation rate $t_1$ for him, then he wants agent 1 to be taxed as little as possible—so that agent 1 will produce as much as possible. On the other hand, if agent 2 thinks that 1 will propose a small compensation rate for him, then agent 2 wants agent 1 to be taxed as much as possible. The only point where agent 2 is indifferent about the production level of agent 1 is where agent 2 is exactly compensated, on the margin, for the costs of the externality.

## 24.4 Efficiency conditions in the presence of externalities

Here we derive general efficiency conditions in the presence of externalities. Suppose that there are two goods, an x-good and a y-good, and two agents. Each agent cares about the other agent's consumption of the x-good, but neither agent cares about the other agent's consumption of the y-good. Initially, there are $\overline{x}$ units of the x-good available and $\overline{y}$ units of the y-good.

According to Chapter 17, page 332, a Pareto efficient allocation maximizes the sum of the utilities subject to the resource constraint

$$\max_{x_i, y_i} a_1 u_1(x_1, x_2, y_1) + a_2 u_2(x_1, x_2, y_2)$$
$$\text{such that } x_1 + x_2 = \overline{x}$$
$$y_1 + y_2 = \overline{y}.$$

The first-order conditions are

$$a_1 \frac{\partial u_1}{\partial x_1} + a_2 \frac{\partial u_2}{\partial x_1} = \lambda$$
$$a_1 \frac{\partial u_1}{\partial x_2} + a_2 \frac{\partial u_2}{\partial x_2} = \lambda$$
$$a_1 \frac{\partial u_1}{\partial y_1} = \mu$$
$$a_2 \frac{\partial u_2}{\partial y_2} = \mu.$$

After some manipulation, these conditions can be written as

$$\frac{\frac{\partial u_1}{\partial x_1}}{\frac{\partial u_1}{\partial y_1}} + \frac{\frac{\partial u_2}{\partial x_1}}{\frac{\partial u_2}{\partial y_2}} = \frac{\lambda}{\mu}$$
$$\frac{\frac{\partial u_1}{\partial x_2}}{\frac{\partial u_1}{\partial y_1}} + \frac{\frac{\partial u_2}{\partial x_2}}{\frac{\partial u_2}{\partial y_2}} = \frac{\lambda}{\mu}.$$

The efficiency condition is that the *sum* of the marginal rates of substitution equals a constant. When determining whether or not it is a good idea whether agent 1 should increase his consumption of good 1, we have to take into account not how much he is willing to pay for this additional consumption, but how much agent 2 is willing to pay. These are essentially the same conditions as the efficiency conditions for a public good.

It is clear from these conditions how to internalize the externality. We simply regard $x_1$ and $x_2$ as different goods. The price of $x_1$ is $p_1 = \partial u_2/\partial x_1$, and the price of $x_2$ is $p_2 = \partial u_1/\partial x_2$. If each agent faces the appropriate price for his actions, the market equilibrium will lead to an efficient outcome.

## Notes

Pigou (1920) and Coase (1960) are classic works on externalities. The compensation mechanism is examined further in Varian (1989b).

## Exercises

24.1. Suppose that two agents are deciding how fast to drive their cars. Agent $i$ chooses speed $x_i$ and gets utility $u_i(x_i)$ from this choice; we assume that $u_i'(x_i) > 0$. However, the faster the agents drive, the more likely it is that they are involved in a mutual accident. Let $p(x_1, x_2)$ be the probability of an accident, assumed to be increasing in each argument, and let $c_i > 0$ be the cost that the accident imposes on agent $i$. Assume that each agent's utility is linear in money.

*(a)* Show that each agent has an incentive to drive too fast from the social point of view.

*(b)* If agent $i$ is fined an amount $t_i$ in the case of an accident, how large should $t_i$ be to internalize the externality?

*(c)* If the optimal fines are being used, what are the total costs, including fines, paid by the agents? How does this compare to the total cost of the accident?

*(d)* Suppose now that agent $i$ gets utility $u_i(x)$ only if there is no accident. What is the appropriate fine in this case?