CHAPTER 14

CONSUMER'S SURPLUS

In the preceding chapters we have seen how to derive a consumer's demand function from the underlying preferences or utility function. But in practice we are usually concerned with the reverse problem—how to estimate preferences or utility from observed demand behavior.

We have already examined this problem in two other contexts. In Chapter 5 we showed how one could estimate the parameters of a utility function from observing demand behavior. In the Cobb-Douglas example used in that chapter, we were able to estimate a utility function that described the observed choice behavior simply by calculating the average expenditure share of each good. The resulting utility function could then be used to evaluate changes in consumption.

In Chapter 7 we described how to use revealed preference analysis to recover estimates of the underlying preferences that may have generated some observed choices. These estimated indifference curves can also be used to evaluate changes in consumption.

In this chapter we will consider some more approaches to the problem of estimating utility from observing demand behavior. Although some of the methods we will examine are less general than the two methods we examined previously, they will turn out to be useful in several applications that we will discuss later in the book.

We will start by reviewing a special case of demand behavior for which it is very easy to recover an estimate of utility. Later we will consider more general cases of preferences and demand behavior.

14.1 Demand for a Discrete Good

Let us start by reviewing demand for a discrete good with quasilinear utility, as described in Chapter 6. Suppose that the utility function takes the form $v(x) + y$ and that the x-good is only available in integer amounts. Let us think of the y-good as money to be spent on other goods and set its price to 1. Let p be the price of the x-good.

We saw in Chapter 6 that in this case consumer behavior can be described in terms of the reservation prices, $r_1 = v(1) - v(0)$, $r_2 = v(2) - v(1)$, and so on. The relationship between reservation prices and demand was very simple: if n units of the discrete good are demanded, then $r_n \geq p \geq r_{n+1}$.

To verify this, let's look at an example. Suppose that the consumer chooses to consume 6 units of the x-good when its price is p. Then the utility of consuming $(6, m - 6p)$ must be at least as large as the utility of consuming any other bundle $(x, m - px)$:

$$v(6) + m - 6p \geq v(x) + m - px. \tag{14.1}$$

In particular this inequality must hold for $x = 5$, which gives us

$$v(6) + m - 6p \geq v(5) + m - 5p.$$

Rearranging, we have $v(6) - v(5) = r_6 \geq p$.

Equation (14.1) must also hold for $x = 7$. This gives us

$$v(6) + m - 6p \geq v(7) + m - 7p,$$

which can be rearranged to yield

$$p \geq v(7) - v(6) = r_7.$$

This argument shows that if 6 units of the x-good is demanded, then the price of the x-good must lie between r_6 and r_7. In general, if n units of the x-good are demanded at price p, then $r_n \geq p \geq r_{n+1}$, as we wanted to show. The list of reservation prices contains all the information necessary to describe the demand behavior. The graph of the reservation prices forms a "staircase" as shown in Figure 14.1. This staircase is precisely the demand curve for the discrete good.

14.2 Constructing Utility from Demand

We have just seen how to construct the demand curve given the reservation prices or the utility function. But we can also do the same operation in reverse. If we are given the demand curve, we can construct the utility function—at least in the special case of quasilinear utility.

At one level, this is just a trivial operation of arithmetic. The reservation prices are defined to be the difference in utility:

$$\begin{aligned} r_1 &= v(1) - v(0) \\ r_2 &= v(2) - v(1) \\ r_3 &= v(3) - v(2) \\ &\vdots \end{aligned}$$

If we want to calculate $v(3)$, for example, we simply add up both sides of this list of equations to find

$$r_1 + r_2 + r_3 = v(3) - v(0).$$

It is convenient to set the utility from consuming zero units of the good equal to zero, so that $v(0) = 0$, and therefore $v(n)$ is just the sum of the first n reservation prices.

This construction has a nice geometrical interpretation that is illustrated in Figure 14.1A. The utility from consuming n units of the discrete good is just the area of the first n bars which make up the demand function. This is true because the height of each bar is the reservation price associated with that level of demand and the width of each bar is 1. This area is sometimes called the **gross benefit** or the **gross consumer's surplus** associated with the consumption of the good.

Note that this is only the utility associated with the consumption of good 1. The final utility of consumption depends on the how much the consumer consumes of good 1 *and* good 2. If the consumer chooses n units of the discrete good, then he will have $m - pn$ dollars left over to purchase other things. This leaves him with a total utility of

$$v(n) + m - pn.$$

This utility also has an interpretation as an area: we just take the area depicted in Figure 14.1A, subtract off the expenditure on the discrete good, and add m.

The term $v(n) - pn$ is called **consumer's surplus** or the **net consumer's surplus**. It measures the net benefits from consuming n units of the discrete good: the utility $v(n)$ minus the reduction in the expenditure on consumption of the other good. The consumer's surplus is depicted in Figure 14.1B.

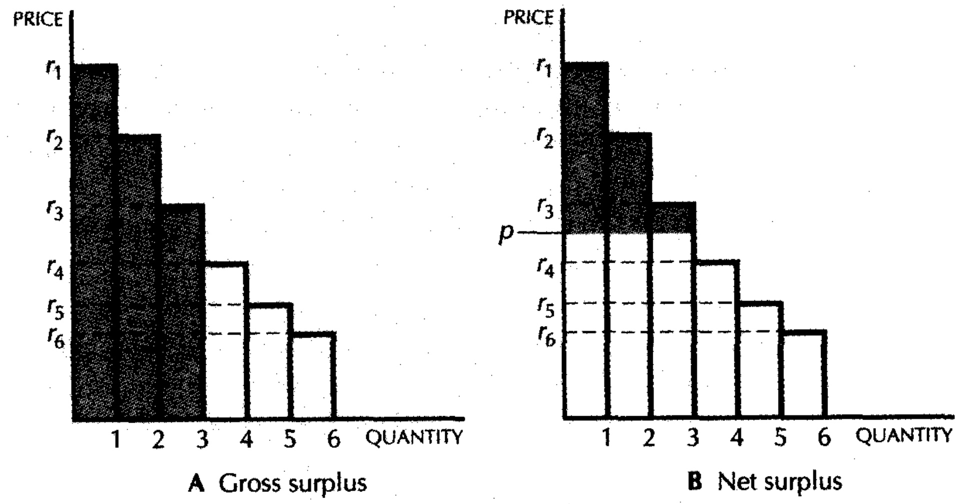

Reservation prices and consumer's surplus. The gross benefit in panel A is the area under the demand curve. This measures the utility from consuming the x-good. The consumer's surplus is depicted in panel B. It measures the utility from consuming both goods when the first good has to be purchased at a constant price p.

14.3 Other Interpretations of Consumer's Surplus

There are some other ways to think about consumer's surplus. Suppose that the price of the discrete good is p. Then the value that the consumer places on the first unit of consumption of that good is r_1, but he only has to pay p for it. This gives him a "surplus" of $r_1 - p$ on the first unit of consumption. He values the second unit of consumption at r_2, but again he only has to pay p for it. This gives him a surplus of $r_2 - p$ on that unit. If we add this up over all n units the consumer chooses, we get his total consumer's surplus:

$$CS = r_1 - p + r_2 - p + \cdots + r_n - p = r_1 + \cdots + r_n - np.$$

Since the sum of the reservation prices just gives us the utility of consumption of good 1, we can also write this as

$$CS = v(n) - pn.$$

We can interpret consumer's surplus in yet another way. Suppose that a consumer is consuming n units of the discrete good and paying pn dollars to do so. How much money would he need to induce him to give up his entire consumption of this good? Let R be the required amount of money. Then R must satisfy the equation

$$v(0) + m + R = v(n) + m - pn.$$

Since $v(0) = 0$ by definition, this equation reduces to

$$R = v(n) - pn,$$

which is just consumer's surplus. Hence the consumer's surplus measures how much a consumer would need to be paid to give up his entire consumption of some good.

14.4 From Consumer's Surplus to Consumers' Surplus

Up until now we have been considering the case of a single consumer. If several consumers are involved we can add up each consumer's surplus across all the consumers to create an aggregate measure of the **consumers' surplus**. Note carefully the distinction between the two concepts: consumer's surplus refers to the surplus of a single consumer; consumers' surplus refers to the sum of the surpluses across a number of consumers.

Consumers' surplus serves as a convenient measure of the aggregate gains from trade, just as consumer's surplus serves as a measure of the individual gains from trade.

14.5 Approximating a Continuous Demand

We have seen that the area underneath the demand curve for a discrete good measures the utility of consumption of that good. We can extend this to the case of a good available in continuous quantities by approximating the continuous demand curve by a staircase demand curve. The area under the continuous demand curve is then approximately equal to the area under the staircase demand.

See Figure 14.2 for an example. In the Appendix to this chapter we show how to use calculus to calculate the exact area under a demand curve.

14.6 Quasilinear Utility

It is worth thinking about the role that quasilinear utility plays in this analysis. In general the price at which a consumer is willing to purchase

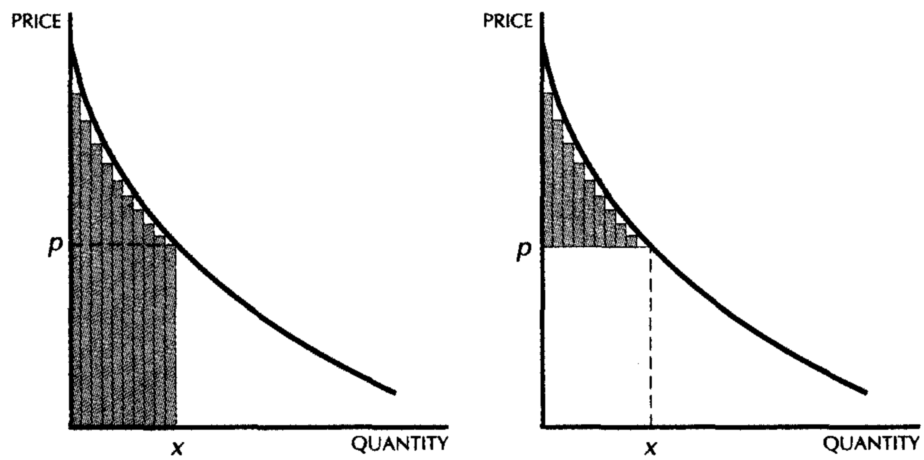

A Approximation to gross surplus **B** Approximation to net surplus

Approximating a continuous demand. The consumer's surplus associated with a continuous demand curve can be approximated by the consumer's surplus associated with a discrete approximation to it.

some amount of good 1 will depend on how much money he has for consuming other goods. This means that in general the reservation prices for good 1 will depend on how much good 2 is being consumed.

But in the special case of quasilinear utility the reservation prices are independent of the amount of money the consumer has to spend on other goods. Economists say that with quasilinear utility there is "no income effect" since changes in income don't affect demand. This is what allows us to calculate utility in such a simple way. Using the area under the demand curve to measure utility will only be *exactly* correct when the utility function is quasilinear.

But it may often be a good approximation. If the demand for a good doesn't change very much when income changes, then the income effects won't matter very much, and the change in consumer's surplus will be a reasonable approximation to the change in the consumer's utility.[1]

14.7 Interpreting the Change in Consumer's Surplus

We are usually not terribly interested in the absolute level of consumer's surplus. We are generally more interested in the change in consumer's

[1] Of course, the change in consumer's surplus is only one way to represent a change in utility—the change in the square root of consumer's surplus would be just as good. But it is standard to use consumer's surplus as a standard measure of utility.

surplus that results from some policy change. For example, suppose the price of a good changes from p' to p''. How does the consumer's surplus change?

In Figure 14.3 we have illustrated the change in consumer's surplus associated with a change in price. The change in consumer's surplus is the difference between two roughly triangular regions and will therefore have a roughly trapezoidal shape. The trapezoid is further composed of two subregions, the rectangle indicated by R and the roughly triangular region indicated by T.

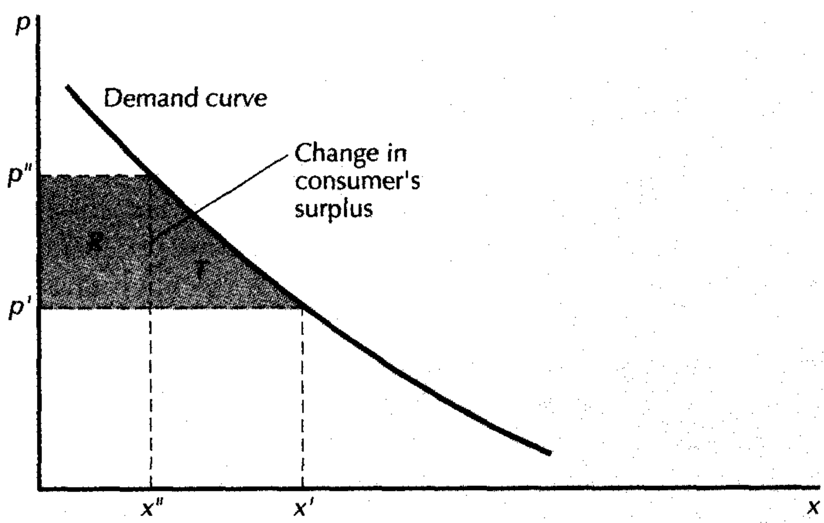

Change in consumer's surplus. The change in consumer's surplus will be the difference between two roughly triangular areas, and thus will have a roughly trapezoidal shape.

The rectangle measures the loss in surplus due to the fact that the consumer is now paying more for all the units he continues to consume. After the price increases the consumer continues to consume x'' units of the good, and each unit of the good is now more expensive by $p'' - p'$. This means he has to spend $(p'' - p')x''$ more money than he did before just to consume x'' units of the good.

But this is not the entire welfare loss. Due to the increase in the price of the x-good, the consumer has decided to consume less of it than he was before. The triangle T measures the value of the *lost* consumption of the x-good. The total loss to the consumer is the sum of these two effects: R measures the loss from having to pay more for the units he continues to consume, and T measures the loss from the reduced consumption.

EXAMPLE: The Change in Consumer's Surplus

Question: Consider the linear demand curve $D(p) = 20 - 2p$. When the price changes from 2 to 3 what is the associated change in consumer's surplus?

Answer: When $p = 2$, $D(2) = 16$, and when $p = 3$, $D(3) = 14$. Thus we want to compute the area of a trapezoid with a height of 1 and bases of 14 and 16. This is equivalent to a rectangle with height 1 and base 14 (having an area of 14), plus a triangle of height 1 and base 2 (having an area of 1). The total area will therefore be 15.

14.8 Compensating and Equivalent Variation

The theory of consumer's surplus is very tidy in the case of quasilinear utility. Even if utility is not quasilinear, consumer's surplus may still be a reasonable measure of consumer's welfare in many applications. Usually the errors in measuring demand curves outweigh the approximation errors from using consumer's surplus.

But it may be that for some applications an approximation may not be good enough. In this section we'll outline a way to measure "utility changes" without using consumer's surplus. There are really two separate issues involved. The first has to do with how to estimate utility when we can observe a number of consumer choices. The second has to do with how we can measure utility in monetary units.

We've already investigated the estimation problem. We gave an example of how to estimate a Cobb-Douglas utility function in Chapter 6. In that example we noticed that expenditure shares were relatively constant and that we could use the average expenditure share as estimates of the Cobb-Douglas parameters. If the demand behavior didn't exhibit this particular feature, we would have to choose a more complicated utility function, but the principle would be just the same: if we have enough observations on demand behavior and that behavior is consistent with maximizing something, then we will generally be able to estimate the function that is being maximized.

Once we have an estimate of the utility function that describes some observed choice behavior we can use this function to evaluate the impact of proposed changes in prices and consumption levels. At the most fundamental level of analysis, this is the best we can hope for. All that matters are the consumer's preferences; any utility function that describes the consumer's preferences is as good as any other.

However, in some applications it may be convenient to use certain monetary measures of utility. For example, we could ask how much money we would have to give a consumer to compensate him for a change in his consumption patterns. A measure of this type essentially measures a change in utility, but it measures it in monetary units. What are convenient ways to do this?

Suppose that we consider the situation depicted in Figure 14.4. Here the consumer initially faces some prices $(p_1^*, 1)$ and consumes some bundle (x_1^*, x_2^*). The price of good 1 then increases from p_1^* to $\hat{p}_1$, and the consumer changes his consumption to $(\hat{x}_1, \hat{x}_2)$. How much does this price change hurt the consumer?

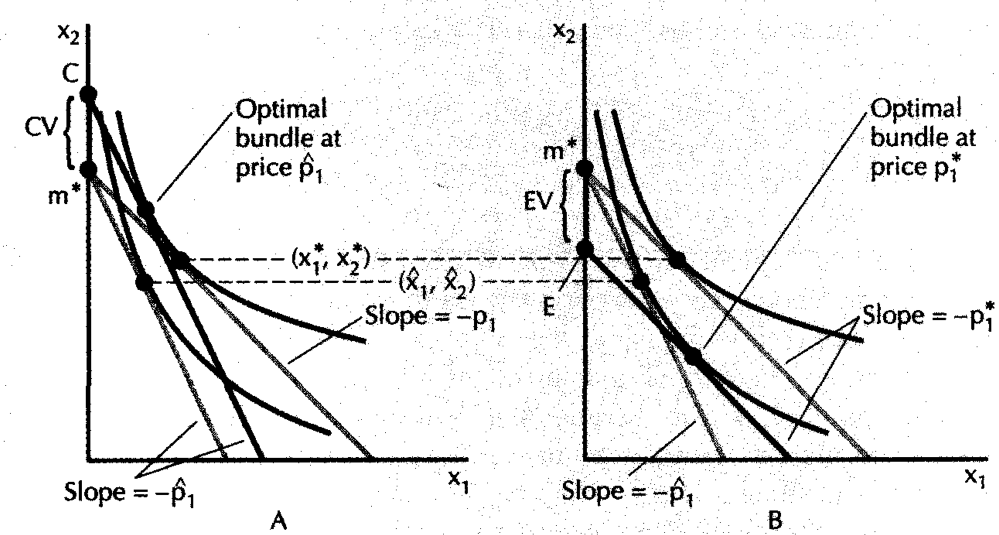

The compensating and the equivalent variations. Panel A shows the compensating variation (CV), and panel B shows the equivalent variation (EV).

One way to answer this question is to ask how much money we would have to give the consumer *after* the price change to make him just as well off as he was *before* the price change. In terms of the diagram, we ask how far up we would have to shift the new budget line to make it tangent to the indifference curve that passes through the original consumption point (x_1^*, x_2^*). The change in income necessary to restore the consumer to his original indifference curve is called the **compensating variation** in income, since it is the change in income that will just compensate the consumer for the price change. The compensating variation measures how much extra money the government would have to give the consumer if it wanted to exactly compensate the consumer for the price change.

Another way to measure the impact of a price change in monetary terms is to ask how much money would have to be taken away from the consumer *before* the price change to leave him as well off as he would be *after* the price change. This is called the **equivalent variation** in income since it is the income change that is equivalent to the price change in terms of the change in utility. In Figure 14.4 we ask how far down we must shift the original budget line to just touch the indifference curve that passes through the new consumption bundle. The equivalent variation measures the maximum amount of income that the consumer would be willing to pay to avoid the price change.

In general the amount of money that the consumer would be willing to pay to avoid a price change would be different from the amount of money that the consumer would have to be paid to compensate him for a price change. After all, at different sets of prices a dollar is worth a different amount to a consumer since it will purchase different amounts of consumption.

In geometric terms, the compensating and equivalent variations are just two different ways to measure "how far apart" two indifference curves are. In each case we are measuring the distance between two indifference curves by seeing how far apart their tangent lines are. In general this measure of distance will depend on the slope of the tangent lines—that is, on the prices that we choose to determine the budget lines.

However, the compensating and equivalent variation are the same in one important case—the case of quasilinear utility. In this case the indifference curves are parallel, so the distance between any two indifference curves is the same no matter where it is measured, as depicted in Figure 14.5. In the case of quasilinear utility the compensating variation, the equivalent variation, and the change in consumer's surplus all give the same measure of the monetary value of a price change.

EXAMPLE: Compensating and Equivalent Variations

Suppose that a consumer has a utility function $u(x_1, x_2) = x_1^{\frac{1}{2}} x_2^{\frac{1}{2}}$. He originally faces prices $(1, 1)$ and has income 100. Then the price of good 1 increases to 2. What are the compensating and equivalent variations?

We know that the demand functions for this Cobb-Douglas utility function are given by

$$x_1 = \frac{m}{2p_1}$$
$$x_2 = \frac{m}{2p_2}.$$

Using this formula, we see that the consumer's demands change from $(x_1^*, x_2^*) = (50, 50)$ to $(\hat{x}_1, \hat{x}_2) = (25, 50)$.

To calculate the compensating variation we ask how much money would be necessary at prices (2,1) to make the consumer as well off as he was consuming the bundle (50,50)? If the prices were (2,1) and the consumer

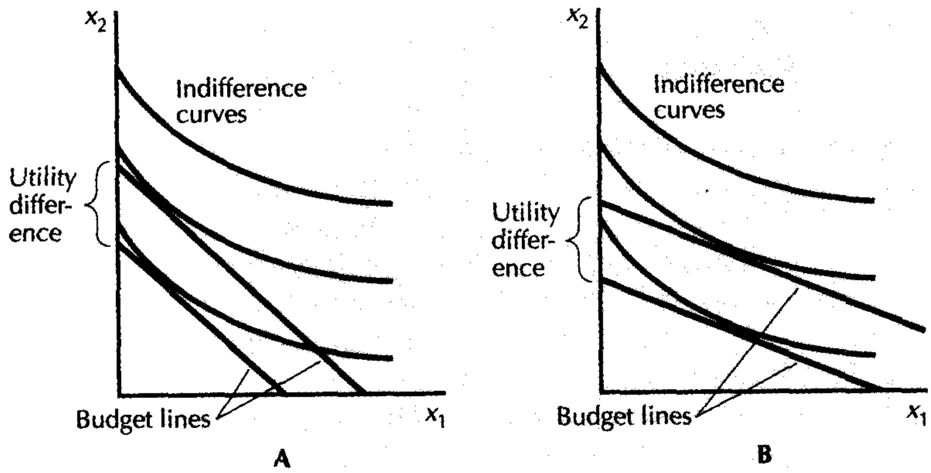

Quasilinear preferences. With quasilinear preferences, the distance between two indifference curves is independent of the slope of the budget lines.

had income m, we can substitute into the demand functions to find that the consumer would optimally choose the bundle $(m/4, m/2)$. Setting the utility of this bundle equal to the utility of the bundle $(50, 50)$ we have

$$\left(\frac{m}{4}\right)^{\frac{1}{2}} \left(\frac{m}{2}\right)^{\frac{1}{2}} = 50^{\frac{1}{2}} 50^{\frac{1}{2}}.$$

Solving for m gives us

$$m = 100\sqrt{2} \approx 141.$$

Hence the consumer would need about $141 - 100 = \$41$ of additional money after the price change to make him as well off as he was before the price change.

In order to calculate the equivalent variation we ask how much money would be necessary at the prices (1,1) to make the consumer as well off as he would be consuming the bundle (25,50). Letting m stand for this amount of money and following the same logic as before,

$$\left(\frac{m}{2}\right)^{\frac{1}{2}} \left(\frac{m}{2}\right)^{\frac{1}{2}} = 25^{\frac{1}{2}} 50^{\frac{1}{2}}.$$

Solving for m gives us

$$m = 50\sqrt{2} \approx 70.$$

Thus if the consumer had an income of \$70 at the original prices, he would be just as well off as he would be facing the new prices and having an income of \$100. The equivalent variation in income is therefore about $100 - 70 = \$30$.

EXAMPLE: Compensating and Equivalent Variation for Quasilinear Preferences

Suppose that the consumer has a quasilinear utility function $v(x_1) + x_2$. We know that in this case the demand for good 1 will depend only on the price of good 1, so we write it as $x_1(p_1)$. Suppose that the price changes from p_1^* to $\hat{p}_1$. What are the compensating and equivalent variations?

At the price p_1^*, the consumer chooses $x_1^* = x_1(p_1^*)$ and has a utility of $v(x_1^*) + m - p_1^* x_1^*$. At the price $\hat{p}_1$, the consumer choose $\hat{x}_1 = x_1(\hat{p}_1)$ and has a utility of $v(\hat{x}_1) + m - \hat{p}_1\hat{x}_1$.

Let C be the compensating variation. This is the amount of extra money the consumer would need after the price change to make him as well off as he would be before the price change. Setting these utilities equal we have

$$v(\hat{x}_1) + m + C - \hat{p}_1\hat{x}_1 = v(x_1^*) + m - p_1^* x_1^*.$$

Solving for C we have

$$C = v(x_1^*) - v(\hat{x}_1) + \hat{p}_1\hat{x}_1 - p_1^* x_1^*.$$

Let E be the equivalent variation. This is the amount of money that you could take away from the consumer before the price change that would leave him with the same utility that he would have after the price change. Thus it satisfies the equation

$$v(x_1^*) + m - E - p_1^* x_1^* = v(\hat{x}_1) + m - \hat{p}_1\hat{x}_1.$$

Solving for E, we have

$$E = v(x_1^*) - v(\hat{x}_1) + \hat{p}_1\hat{x}_1 - p_1^* x_1^*.$$

Note that for the case of quasilinear utility the compensating and equivalent variation are the same. Furthermore, they are both equal to the change in (net) consumer's surplus:

$$\Delta CS = [v(x_1^*) - p_1^* x_1^*] - [v(\hat{x}_1) - \hat{p}_1\hat{x}_1].$$

14.9 Producer's Surplus

The demand curve measures the amount that will be demanded at each price; the **supply curve** measures the amount that will be supplied at each price. Just as the area *under* the demand curve measures the surplus enjoyed by the demanders of a good, the area *above* the supply curve measures the surplus enjoyed by the suppliers of a good.

We've referred to the area under the demand curve as consumer's surplus. By analogy, the area above the supply curve is known as **producer's surplus**. The terms consumer's surplus and producer's surplus are somewhat misleading, since who is doing the consuming and who is doing the producing really doesn't matter. It would be better to use the terms "demander's surplus" and "supplier's surplus," but we'll bow to tradition and use the standard terminology.

Suppose that we have a supply curve for a good. This simply measures the amount of a good that will be supplied at each possible price. The good could be supplied by an individual who owns the good in question, or it could be supplied by a firm that produces the good. We'll take the latter interpretation so as to stick with the traditional terminology and depict the producer's supply curve in Figure 14.6. If the producer is able to sell x^* units of her product in a market at a price p^*, what is the surplus she enjoys?

It is most convenient to conduct the analysis in terms of the producer's *inverse* supply curve, $p_s(x)$. This function measures what the price would have to be to get the producer to supply x units of the good.

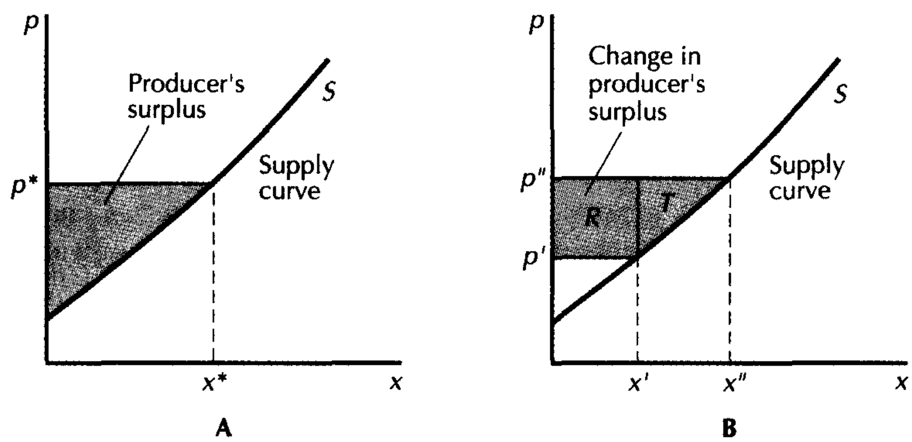

Producer's surplus. The net producer's surplus is the triangular area to the left of the supply curve in panel A, and the change in producer's surplus is the trapezoidal area in panel B.

Think about the inverse supply function for a discrete good. In this case the producer is willing to sell the first unit of the good at price $p_s(1)$, but she actually gets the market price p^* for it. Similarly, she is willing to sell the second unit for $p_s(2)$, but she gets p^* for it. Continuing in this way we see that the producer will be just willing to sell the last unit for $p_s(x^*) = p^*$.

The difference between the minimum amount she would be willing to sell the x^* units for and the amount she actually sells the units for is the **net producer's surplus**. It is the triangular area depicted in Figure 14.6A.

Just as in the case of consumer's surplus, we can ask how producer's surplus changes when the price increases from p' to p''. In general, the change in producer's surplus will be the difference between two triangular regions and will therefore generally have the roughly trapezoidal shape depicted in Figure 14.6B. As in the case of consumer's surplus, the roughly trapezoidal region will be composed of a rectangular region R and a roughly triangular region T. The rectangle measures the gain from selling the units previously sold anyway at p' at the higher price p''. The roughly triangular region measures the gain from selling the extra units at the price p''. This is analogous to the change in consumer's surplus considered earlier.

Although it is common to refer to this kind of change as an increase in producer's surplus, in a deeper sense it really represents an increase in consumer's surplus that accrues to the consumers who own the firm that generated the supply curve. Producer's surplus is closely related to the idea of profit, but we'll have to wait until we study firm behavior in more detail to spell out the relationship.

14.10 Benefit-Cost Analysis

We can use the consumer surplus apparatus we have developed to calculate the benefits and costs of various economic policies.

For example, let us examine the impact of a price ceiling. Consider the situation depicted in Figure 14.7. With no intervention, the price would be p_0 and the quantity sold would be q_0.

The authorities believe this price is too high and impose the price ceiling at p_c. This reduces the amount that suppliers are willing to supply to q_c which, in turn, reduces their producer surplus to the shaded area in the diagram.

Now that there is only q_c available for consumers, the question is who will get it?

One assumption is that the output will go to the consumers with the highest willingness to pay. Let p_e, the **effective price**, be the price that would induce consumers to demand q_e. If everyone who is willing to pay more than p_e gets the good, then the producer surplus will be the shaded area in the diagram.

Note that the lost consumer and producer surplus is given by the trapezoidal area in the middle of the diagram. This is the difference between

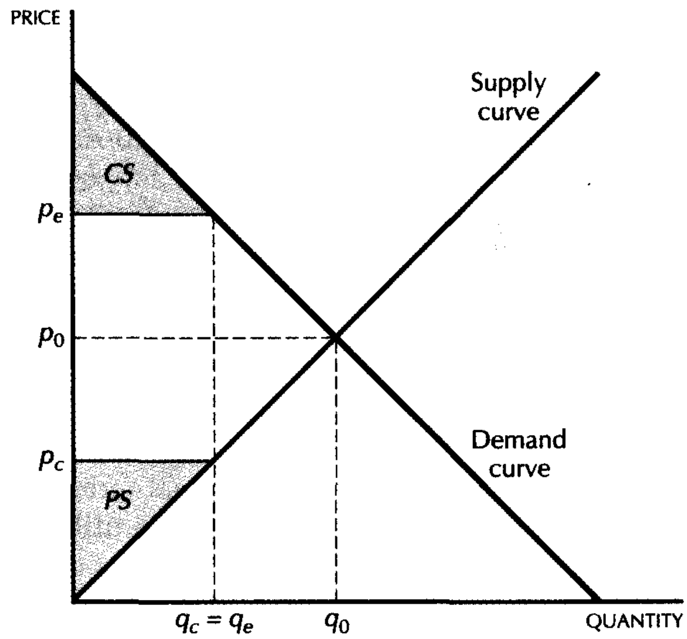

A price ceiling. The price ceiling at p_c reduces supply to q_e. It reduces consumer surplus to CS and producer surplus to PS. The effective price of the good, p_e, is the price that would clear the market. The diagram also shows what happens with rationing, in which case the price of a ration coupon would be $p_e - p_c$.

the consumer plus producer surplus in the competitive market and the difference in the market with the price ceiling.

Assuming that the quantity will go to consumers with the highest willingness to pay is overly optimistic in most situation. Hence, we we would generally expect that this trapezoidal area is a lower bound on the lost consumer plus producer surplus in the case of a price ceiling.

Rationing

The diagram we have just examined can also be used to describe the social losses due to rationing. Instead of fixing a price ceiling of p_c, suppose that the authorities issue ration coupons that allow for only q_c units to be purchased. In order to purchase one unit of the good, a consumer needs to pay p_c to the seller and produce a ration coupon.

If the ration coupons are marketable, then they would sell for a price of $p_e - p_c$. This would make the the total price of the purchase equal to p_e, which is the price that clears the market for the good being sold.

14.11 Calculating Gains and Losses

If we have estimates of the market demand and supply curves for a good, it is not difficult in principle to calculate the loss in consumers' surplus due to changes in government policies. For example, suppose the government decides to change its tax treatment of some good. This will result in a change in the prices that consumers face and therefore a change in the amount of the good that they will choose to consume. We can calculate the consumers' surplus associated with different tax proposals and see which tax reforms generate the smallest loss.

This is often useful information for judging various methods of taxation, but it suffers from two defects. First, as we've indicated earlier, the consumer's surplus calculation is only valid for special forms of preferences—namely, preferences representable by a quasilinear utility function. We argued earlier that this kind of utility function may be a reasonable approximation for goods for which changes in income lead to small changes in demand, but for goods whose consumption is closely related to income, the use of consumer surplus may be inappropriate.

Second, the calculation of this loss effectively lumps together all the consumers and producers and generates an estimate of the "cost" of a social policy only for some mythical "representative consumer." In many cases it is desirable to know not only the average cost across the population, but who bears the costs. The political success or failure of policies often depends more on the *distribution* of gains and losses than on the average gain or loss.

Consumer's surplus may be easy to calculate, but we've seen that it is not that much more difficult to calculate the true compensating or equivalent variation associated with a price change. If we have estimates of the demand functions of each household—or at least the demand functions for a sample of representative households—we can calculate the impact of a policy change on each household in terms of the compensating or equivalent variation. Thus we will have a measure of the "benefits" or "costs" imposed on each household by the proposed policy change.

Mervyn King, an economist at the London School of Economics, has described a nice example of this approach to analyzing the implications of reforming the tax treatment of housing in Britain in his paper "Welfare Analysis of Tax Reforms Using Household Data," *Journal of Public Economics*, 21 (1983), 183–214.

King first examined the housing expenditures of 5,895 households and estimated a demand function that best described their purchases of housing services. Next, he used this demand function to determine a utility function for each household. Finally, he used the estimated utility function to calculate how much each household would gain or lose under certain changes in the taxation of housing in Britain. The measure that he used was similar to the equivalent variation described earlier in this chapter. The basic nature of the tax reform he studied was to eliminate tax concessions to owner-occupied housing and to raise rents in public housing. The revenues generated by these changes would be handed back to the households in the form of transfers proportional to household income.

King found that 4,888 of the 5,895 households would benefit from this kind of reform. More importantly he could identify explicitly those households that would have significant losses from the tax reform. King found, for example, that 94 percent of the highest income households gained from the reform, while only 58 percent of the lowest income households gained. This kind of information would allow special measures to be undertaken which might help in designing the tax reform in a way that could satisfy distributional objectives.

Summary

1. In the case of a discrete good and quasilinear utility, the utility associated with the consumption of n units of the discrete good is just the sum of the first n reservation prices.

2. This sum is the gross benefit of consuming the good. If we subtract the amount spent on the purchase of the good, we get the consumer's surplus.

3. The change in consumer's surplus associated with a price change has a roughly trapezoidal shape. It can be interpreted as the change in utility associated with the price change.

4. In general, we can use the compensating variation and the equivalent variation in income to measure the monetary impact of a price change.

5. If utility is quasilinear, the compensating variation, the equivalent variation, and the change in consumer's surplus are all equal. Even if utility is not quasilinear, the change in consumer's surplus may serve as a good approximation of the impact of the price change on a consumer's utility.

6. In the case of supply behavior we can define a producer's surplus that measures the net benefits to the supplier from producing a given amount of output.

REVIEW QUESTIONS

1. A good can be produced in a competitive industry at a cost of \$10 per unit. There are 100 consumers are each willing to pay \$12 each to consume

a single unit of the good (additional units have no value to them.) What is the equilibrium price and quantity sold? The government imposes a tax of \$1 on the good. What is the deadweight loss of this tax?

2. Suppose that the demand curve is given by $D(p) = 10 - p$. What is the gross benefit from consuming 6 units of the good?

3. In the above example, if the price changes from 4 to 6, what is the change in consumer's surplus?

4. Suppose that a consumer is consuming 10 units of a discrete good and the price increases from \$5 per unit to \$6. However, after the price change the consumer continues to consume 10 units of the discrete good. What is the loss in the consumer's surplus from this price change?

APPENDIX

Let's use some calculus to treat consumer's surplus rigorously. Start with the problem of maximizing quasilinear utility:

$$\max_{x,y} v(x) + y$$
$$\text{such that } px + y = m.$$

Substituting from the budget constraint we have

$$\max_x v(x) + m - px.$$

The first-order condition for this problem is

$$v'(x) = p.$$

This means that the inverse demand function $p(x)$ is defined by

$$p(x) = v'(x). \qquad (14\ 2$$

Note the analogy with the discrete-good framework described in the text: the price at which the consumer is just willing to consume x units is equal to the marginal utility.

But since the inverse demand curve measures the derivative of utility, we can simply integrate under the inverse demand function to find the utility function

Carrying out the integration we have:

$$v(x) = v(x) - v(0) = \int_0^x v'(t)\,dt = \int_0^x p(t)\,dt.$$

Hence utility associated with the consumption of the x-good is just the area unde the demand curve.

Comparison of CV, CS, and EV.

p_1	CV	CS	EV
1	0.00	0.00	0.00
2	7.18	6.93	6.70
3	11.61	10.99	10.40
4	14.87	13.86	12.94
5	17.46	16.09	14.87

EXAMPLE: A Few Demand Functions

Suppose that the demand function is linear, so that $x(p) = a - bp$. Then the change in consumer's surplus when the price moves from p to q is given by

$$\int_p^q (a - bt)\, dt = at - b\frac{t^2}{2}\Big]_p^q = a(q-p) - b\frac{q^2 - p^2}{2}.$$

Another commonly used demand function, which we examine in more detail in the next chapter, has the form $x(p) = Ap^\epsilon$, where $\epsilon < 0$ and A is some positive constant. When the price changes from p to q, the associated change in consumer's surplus is

$$\int_p^q At^\epsilon\, dt = A\frac{t^{\epsilon+1}}{\epsilon+1}\Big]_p^q = A\frac{q^{\epsilon+1} - p^{\epsilon+1}}{\epsilon+1},$$

for $\epsilon \neq -1$.

When $\epsilon = -1$, this demand function is $x(p) = A/p$, which is closely related to our old friend the Cobb-Douglas demand, $x(p) = am/p$. The change in consumer's surplus for the Cobb-Douglas demand is

$$\int_p^q \frac{am}{t}\, dt = am \ln t\Big]_p^q = am(\ln q - \ln p).$$

EXAMPLE: CV, EV, and Consumer's Surplus

In the text we calculated the compensating and equivalent variations for the Cobb-Douglas utility function. In the preceding example we calculated the change in consumer's surplus for the Cobb-Douglas utility function. Here we compare these three monetary measures of the impact on utility of a price change.

Suppose that the price of good 1 changes from 1 to $2, 3\ldots$ while the price of good 2 stays fixed at 1 and income stays fixed at 100. Table 14.1 shows the equivalent variation (EV), compensating variation (CV), and the change in consumer's surplus (CS) for the Cobb-Douglas utility function $u(x_1, x_2) = x_1^{\frac{1}{10}} x_2^{\frac{9}{10}}$.

Note that the change in consumer's surplus always lies between the CV and the EV and that the difference between the three numbers is relatively small. It is possible to show that both of these facts are true in reasonably general circumstances. See Robert Willig, "Consumer's Surplus without Apology," *American Economic Review*, 66 (1976), 589–597.