Chapter 5

CONSUMER CHOICE: INDIFFERENCE THEORY

In this chapter we look more closely at the determinants of consumer demand. In particular, we discuss the concept of utility and how we can use this to gain insights into how consumers choose to allocate their spending. We first explain some key insights that were achieved by thinking about utility as if it can be measured. We then outline the approach that does not require utility to be measurable but that yields many similar insights into the determinants of demand. In particular, you will learn that:

- Consumers will maximize their overall satisfaction when the marginal utility per pound spent is equal for all products purchased.
- A theory of demand can be built by focusing on bundles of goods between which the consumer is indifferent.
- Indifference curves show combinations of goods that give the same level of satisfaction.
- A budget constraint shows what the consumer could buy with a given income.
- A consumer optimizes by moving to the highest indifference curve that is available with a given budget constraint.
- The response to a price change can be decomposed into an income and a substitution effect.
- For a good to have a negatively sloped demand curve, it is necessary (but not sufficient) that it be an inferior good.

In this chapter we will first explain some important insights that come from early analyses of utility in the nineteenth century. We then explain how modern economics uses indifference curves to develop a theory of consumer choice. We show how indifference curves can be used to describe consumers' tastes, and then introduce a budget line to describe the consumption possibilities open to a consumer who has a given income. After that, we show how consumers reach equilibrium by consuming the bundle that allows them to reach the highest possible levels of satisfaction. We can then see how any consumer alters behaviour when either income or prices change, and we go on to derive the negative slope of the demand curve.

This approach to consumer behaviour has two great advantages. First, it allows us to distinguish between two effects of a change in price, called the income effect and the substitution effects; this distinction has important practical applications. Second, it allows us to understand the rare but interesting exception to the prediction that all demand curves are negatively sloped, which arises with a so-called Giffen good.

All of the theories in this chapter use the basic assumption that consumers are motivated to make themselves as well off as they can—or, as economists like to put it, to maximize their satisfactions.

Early insights

All units of the same product are identical; for example, one tin of Heinz baked beans is the same as another tin of Heinz baked beans. But the satisfaction that a consumer gets from each unit of a product is not the same. If you are hungry you will get great satisfaction from a good meal, but you will not get the same satisfaction from having a second identical meal immediately. This suggests that the satisfaction that people get from consuming a unit of any product varies according to how many of this product they have already.

Economists and philosophers thinking about consumer choice and satisfaction in the nineteenth century developed the concept of utility and so were sometimes called *utilitarians*.[1] But the big breakthrough for economics came in the 1870s with what is known as the *marginal revolution*, which gave birth to *neoclassical economics*.[2] For a long time it was thought that utility could not be measured and therefore that utility theory was based on unverifiable concepts. Recently, however, economists and psychologists have succeeded in measuring utility. They have used these measurements to verify, among other things, two of the basic assumptions of utility theory: (1) that the marginal utility of any one good declines when more to it is consumed, with the consumption of all other goods held constant, and (2) that the marginal utility of income declines as people earn more of it.[3]

Marginal and total utility

What we want to think about first is how an individual consumer's satisfaction changes as he or she alters the amount consumed of a single product. The satisfaction a consumer receives from consuming that product is called *utility*. **Total utility** refers to the *total satisfaction* derived from all the units of that product consumed. **Marginal utility** refers to the *change in satisfaction* resulting from consuming one unit more or one unit less of that product. For example, the total utility of consuming 14 cups of coffee a week is the sum total satisfaction provided by all 14 cups of coffee. The marginal utility of the fourteenth cup of coffee consumed is the addition to total satisfaction provided by consuming that extra cup. Put another way, the marginal utility of the fourteenth cup is the addition to total utility gained from consuming 14 cups of coffee per week rather than 13.

Diminishing marginal utility

A basic assumption of utility theory, which is sometimes called the *law of diminishing marginal utility*, is as follows:

The marginal utility generated by additional units of any product diminishes as an individual consumes more of it, holding constant the consumption of all other products.

The way in which most of us use water provides a good example of diminishing marginal utility. We consume it in many forms: tap water, soft drinks, bottled water, or water flavoured with such things as tea leaves and coffee grounds. Whatever the form in which we consume it, we consume it: water is necessary to our very existence, and anyone denied water will not survive very long. So we value the minimum of water needed to sustain life as much as we value life itself. We would be willing, therefore, to pay quite a lot if this were the only way to obtain the amount of water needed to stay alive. The total utility of that much water is therefore extremely high, as is the marginal utility of the first few units drunk. More than this bare minimum will be drunk, but the marginal utility of successive amounts of water drunk over any period of time will decline steadily.

Furthermore, water has many uses other than for drinking. A fairly high marginal utility will be attached to some minimum quantity for bathing, but much more than this minimum will be used for more frequent baths or showers. The last weekly gallon used for bathing is likely to have a low marginal utility. Again, some small quantity of water is necessary for tooth brushing, but many people leave the water running while they brush. The water going down the drain between wetting and rinsing the brush surely has a low utility. When all the many uses of water by the modern consumer are considered (washing machines, dishwashers, lawn sprinklers, car washing, etc.), it is certain that the marginal utility of the last, say, 10 per cent of all units consumed is very low and falling, even though the total utility of all the units consumed is extremely high.

[1] Leading members of the utilitarian school were Jeremy Bentham (1748–1832), James Mill (1773–1836), and John Stuart Mill (1806–73).

[2] Key contributors to the marginal revolution were: the English economist William Stanley Jevons (1835–82), the Austrian Carl Menger (1841–1910), and the Swiss Leon Walras (1834–1910).

[3] See e.g. Layard (2005).

Maximizing utility

We can now ask: what does diminishing marginal utility imply for the way a consumer who has a given income will allocate spending in order to maximize total utility? How should a consumer allocate his or her income in order to get the greatest possible satisfaction, or total utility, from that spending?

If all products had the same price, the answer would be easy. A consumer should simply allocate spending so that the marginal utility of all products was the same. If the marginal utility of all products were not equal, then total utility could be increased by choosing a different spending pattern. For example, if one product had a higher marginal utility than the others, expenditure should be reallocated so as to buy more of this product, and less of all others that have lower marginal utilities. By buying more, the product's marginal utility would fall. Only when the last unit of all products bought gives the same satisfaction is the consumer getting the greatest possible total utility from his or her spending pattern.

How does this work if products have different prices? Again the same principles apply, but now the best a consumer can do is to rearrange spending until the last unit of satisfaction per pound spent on each product is the same. For example, suppose that a consumer is deciding to allocate income between going to football matches and going to the cinema, and that tickets to football cost £30 while a cinema ticket costs £10. If a consumer gets more than three times as much extra satisfaction from another football match as another movie, then off to more football matches he or she should go. This consumer will be maximizing total utility from her income only when the last match attended generates extra utility that is just three times that generated by the last film.

To maximize utility, consumers allocate spending between products so that equal utility is derived from the last unit of money spent on each.[4]

The conditions for maximizing utility can be stated more generally. Denote the marginal utility of the last unit of product X by MU_X and its price by p_X. Let MU_Y and p_Y refer respectively to the marginal utility of a second product, Y, and its price. The marginal utility per pound spent on X will be MU_X/p_X. For example, if the last unit adds 30 units to utility and costs £2, its marginal utility per pound is $30/2 = 15$.

The condition required for any consumer to maximize utility is that the following relationship should hold, for all pairs of products:

$$MU_X/p_X = MU_Y/p_Y \quad (1)$$

This merely says in symbols what we earlier said in words. Consumers who are maximizing their utility will allocate spending so that the utilities gained from the last £1 spent on both products are equal.

This is the fundamental equation of utility theory. Each consumer demands each good up to the point at which the marginal utility per pound spent on it is the same as the marginal utility of a pound spent on each other good. When this condition is met, the consumer cannot shift a pound of spending from one product to another and increase total utility.

Consumers choose quantities not prices

If we rearrange the terms in equation (1), we can gain additional insight into consumer behaviour:[5]

$$MU_X/MU_Y = p_X/p_Y \quad (2)$$

The right-hand side of this equation states the relative price of the two goods. This is determined by the market and is beyond the control of individual consumers, who react to the market prices but are powerless to change them. The left-hand side of the equation states the relative contribution of the two goods to add to satisfaction if a little more or a little less of either of them were consumed, a choice that is available.

If the two sides of equation (2) are not equal, the consumer can increase total satisfaction by changing their spending pattern. Assume, for example, that the price of a unit of X is twice the price of a unit of Y ($p_X/p_Y = 2$), while the marginal utility of a unit of X is three times that of a unit of Y ($MU_X/MU_Y = 3$). Under these conditions, it pays to buy more of X and less of Y. For example, reducing purchases of Y by two units frees enough purchasing power to buy a unit of X. Since one extra unit of X bought yields 1.5 times the satisfaction of two units of Y forgone, the switch is worth making. What about a further switch of X for Y? As the consumer buys more X and less Y, the marginal utility of X falls and the marginal utility of Y rises. In this example the comsuner will go on rearranging purchases—reducing Y consumption and increasing X consumption—until the marginal utility of X is only twice that of Y. At this point, total satisfaction cannot be further increased by rearranging purchases between the two products.

Think about what the utility-maximizing consumer is doing. She is faced with a set of prices that cannot be

[4] By the 'last unit' we do not mean money spent over successive time-periods: instead, we are talking about buying more or fewer units at one point in time, that is, alternative allocations of spending at a moment in time.

[5] This is done by multiplying both sides of the equation by p_X/MU_Y.

changed. She responds to these prices and maximizes satisfaction by adjusting the things that can be changed—the quantities of the various goods purchased—until equation (2) is satisfied for all pairs of products.

We see this sort of equation frequently in economics—one side representing the choices the outside world presents to decision-takers and the other side representing the effect of those choices. It shows the equilibrium position reached when decision-takers have made the best adjustment they can to the external forces that constrain their choices.

When they enter the market, all consumers face the same set of market prices. When they are fully adjusted to these prices, each one of them will have identical ratios of their marginal utilities for each pair of goods. Of course, a rich consumer may consume more of each product than a poor consumer and get more *total utility* from it. However, the rich and the poor consumer (and every other consumer who is maximizing utility) will adjust their relative purchases of each product so that the relative *marginal utilities* are the same for all. Thus, if the price of X is twice the price of Y, each consumer will purchase X and Y to the point at which their marginal utility of X is twice the marginal utility of Y. Consumers with different tastes, however, will derive different marginal utilities from their consumption of the various commodities, so they will consume differing relative quantities of products. But all will have declining marginal utilities for each commodity and hence, when they have maximized their utility, the ratios of their marginal utilities will be the same for all of them.

A very important insight can be derived from this analysis. It is that marginal, not average, values are what matter for maximization. We will return to this idea in subsequent chapters when we see that marginal values are also important for the profit-maximizing behaviour of firms. Box 5.1 reinforces just how important marginal utility is as a concept, in that it helps explain what used to be known as the 'paradox of value'. The key point to notice from this is that market prices reflect marginal utilities of various products and not total or average utilities. Hence market prices are not a measure of the total value to society of one good or service as compared with that of some other good or service.

Box 5.1 The paradox of value

Early thinkers about the economy struggled with the problem of what determines the relative prices of products. They encountered the *paradox of value*: many essential products without which we could not live, such as water, have relatively low prices. On the other hand, some luxury products, such as diamonds, have relatively high prices, even though we could easily survive without them. Does it not seem odd that water, which is so important to us, has such a low market value while diamonds, which are much less important, have a much higher market value? It took a long time to resolve this apparent paradox, so it is not surprising that even today similar confusions about the determinants of market values persist and cloud many policy discussions.

The key to resolving the 'paradox' lies in the distinction between total and marginal utility. We have already seen in this chapter that a utility-maximizing consumer will adjust his or her spending pattern so that the marginal utility per pound spent is equal for all products. It follows that the value consumers place on the last unit consumed of any product, i.e. its marginal utility, is equal in equilibrium to the product's price.

We will explain in Box 5.2 (on page XXX) that the area under the demand curve above market price represents the total benefit consumers get from consuming a product. We will call this benefit *consumers' surplus*, and we can think of consumer surplus as an indicator of the value of the total utility that consumers get from a product.

Now look at the total amount spent to purchase the product—the price paid for it multiplied by the quantity bought and sold—which we can call its total market value or sale value. The figure shows the markets for two goods, one for which total market value is a very small fraction of its total utility and another for which total market value is a much higher fraction of total utility.

The resolution of the paradox of value is that a good that is very plentiful, such as water, will have a low price. It will be consumed, therefore, to the point where all purchasers place a low value on the last unit consumed, whether or not they place a high value on their total consumption of the product; in other words, marginal utility will be low whatever the value of total utility. On the other hand, a product that is relatively scarce will have a high market price. Consumption will, therefore, stop at a point at which consumers place a high value on the last unit consumed whatever value they place on their total consumption of the good; that is, marginal utility will be high whatever the value of total utility.

These analysis leads to an important conclusion:

The market price of a product depends on demand and supply. Hence no paradox is involved when a product on which consumers place a high total utility sells for a low price, and hence has only a low total market value (i.e. a low amount spent on it).

→

Box 5.1 *continued*

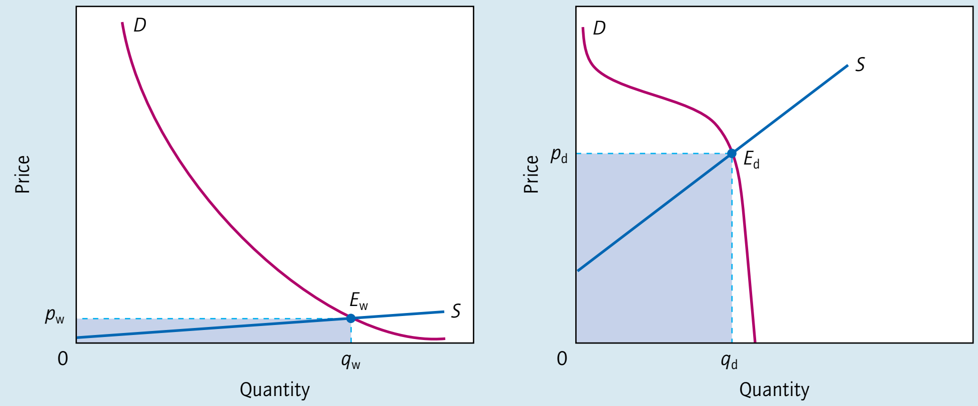

(i) Water

(ii) Diamonds

Total utility versus market value

The market value of the amount of some commodity bears no necessary relation to the total utility that consumers derive from that amount. The total utility that consumers derive from water, as shown by the area under the demand curve in part (i), is great—indeed, we cannot possibly show the curve for very small quantities, because people would pay all they had rather than be deprived completely of water. The total utility that consumers derive from diamonds is shown by the area under the demand curve in part (ii). This is less than the total utility derived from water. The supply curve of diamonds makes diamonds scarce and keeps their price high. Thus, when equilibrium is at E_d, the total market value of diamonds sold, indicated by the dark blue area of $p_d q_d$, is high. The supply curve of water makes water plentiful and makes water low in price. Thus, when equilibrium is at E_w, the total market value of water consumed, indicated by the dark blue area of $p_w q_w$, is low.

Implication of marginal utility theory for demand curves

The assumption that all products exhibit diminishing marginal utility has a simple implication for demand curves: they are all negatively sloped. This is because, if consumers were already maximizing utility and the price of one product fell, then in order to restore equation (2) above, consumers would have to buy more of the product whose price had fallen and less of all other products.

In the twentieth century economists moved away from relying upon the assumption of diminishing marginal utility as a key building block in their theory of demand. The reason is that it was harder to take the theory much further without being able to measure of utility, which seemed impossible at the time. However, considerable progress was made without having to measure utility. All that was needed was to assume that consumers could rank alternative bundles of products in order of preference without necessarily being able to say by *how much* they preferred one to another.

We now outline this modern approach to consumer choice. In it the two key insights that we have just discussed remain valid:

1. Marginal comparisons are what matter for consumer choice, and equations (1) and (2) above remain valid as optimization conditions for consumers whether or not utility is assumed to be measurable.

2. Market prices are determined by marginal utilities and not by total or average utilities.

Box 5.2 outlines a concept known as 'consumers' surplus', which is related to diminishing marginal utility.

Box 5.2 Consumers' surplus

The negative slope of the demand curve has an interesting consequence:

All consumers pay less than they would be willing to pay for the total amount of any product that they consume.

The difference between what they would be willing to pay—which is the value of the total utility that they derive from consuming the product—and what they do pay—which is their total spending on that product—is called **consumers' surplus**.

This concept is important and deserves further elaboration. The table gives hypothetical data for the weekly consumption of milk by one consumer, Ms Green. The second column, labelled 'Total utility', gives the total value she places on consumption of so many glasses per week (when the alternative is zero). Column (3), labelled 'Marginal utility', gives the amount she would pay to add the last glass indicated to weekly consumption. Thus, for example, the marginal utility of £0.80 listed against four glasses gives the value that Ms Green places on increasing consumption from three to four glasses. It is the difference between the total utilities she attaches to consumption levels of three and four glasses per week.

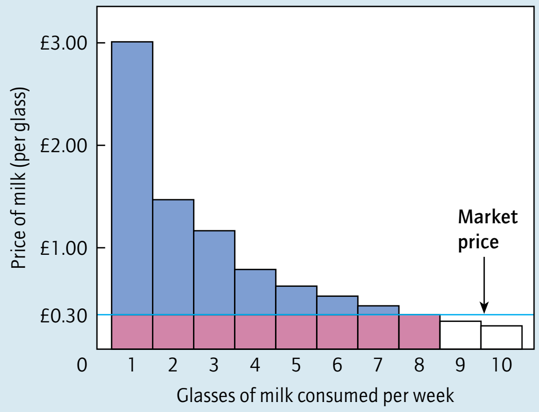

(i) Consumer's surplus for an individual

Consumer's surplus is the sum of the extra valuations placed on each unit above the market price paid for each. This figure is based on the data in Table 5.2. Ms Green pays the red area for the 8 glasses of milk she consumes per week when the market price is £0.30 a glass. The total value she places on these 8 glasses of milk is the entire shaded area (red and green). Hence her consumer's surplus is the green area.

(i) Consumer's surplus

Glasses of milk consumed per week (1)	Total utility (2)	Marginal utility (3)	Consumer's surplus on each glass if milk costs £0.30 per glass (4)
1	£3.00	£3.00	£2.70
2	4.50	1.50	1.20
3	5.50	1.00	0.70
4	6.30	0.80	0.50
5	6.90	0.60	0.30
6	7.40	0.50	0.20
7	7.80	0.40	0.10
8	8.10	0.30	0.00
9	8.35	0.25	—
10	8.55	0.20	—

Consumer's surplus on each unit consumed is the difference between the market price and the maximum price the consumer would pay to obtain that unit. The table shows the value that Ms Green puts on successive glasses of milk consumed each week. As long as she is willing to pay more than the market price for any glass, she obtains a consumer's surplus when she buys it. The marginal glass of milk is the eighth. This is the one she values at just the market price and on which she earns no consumer's surplus.

If Ms Green is faced with a market price of £0.30, she will maximize total utility by consuming eight glasses per week, because she values the eighth glass just at the market price, while valuing all earlier glasses at higher amounts. Because she values the first glass at £3.00 but gets it for £0.30, she makes a 'profit' of £2.70 on that glass; that is, she gets £3.00 worth of satisfaction for £0.30. Between her £1.50 valuation of the second glass and what she has to pay for it, she clears a 'profit' of £1.20. She clears £0.70 on the third glass. And so on. These 'profits', which are called her consumer's surpluses on each unit, are shown in the final column of the table. The total surplus is £5.70 per week. In the table, we calculate Ms Green's surplus by summing the surpluses on each glass. We arrive at the same total, however, by first summing the maximum that Ms Green would pay for all the glasses bought (which is £8.10 in this case) and then subtracting the £2.40 that she does pay.

The value placed by each consumer on his or her total consumption of some product can be estimated in at least two ways. The valuation that the consumer places on each successive unit may be summed, or the consumer may be asked the maximum that he or she would pay to consume the amount in question if the alternative were to have none. While other consumers would put different numerical values into the table, diminishing marginal utility implies that the figures in the final column would

→

Box 5.2 *continued*

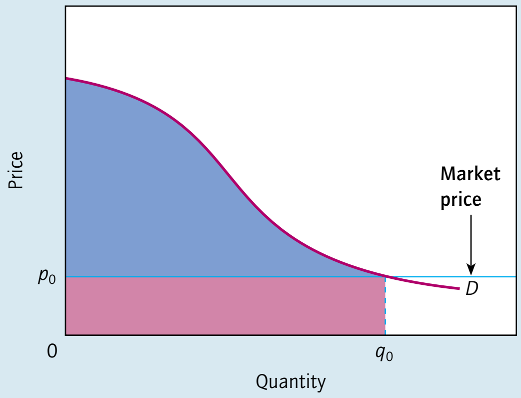

(ii) Consumer's surplus for the market

Total consumers' surplus is the area under the demand curve and above the price line. The area under the demand curve shows the total valuation that consumers place on all units consumed. For example, the total value that consumers place on q_0 units is the entire area shaded red and green under the demand curve up to q_0. At a market price of p_0 the amount paid for q_0 units is the red area. Hence consumers' surplus is the green area.

be declining for each person. Since a consumer will go on buying further units until the value placed on the last unit equals the market price, it follows that there will be a consumers' surplus on every unit consumed except the last one.

The data in columns (1) and (3) of the table give Ms Green's demand curve for milk. This is her demand curve, because she will go on buying glasses of milk as long as she values each glass at least as much as the market price she must pay for it. When the market price is £3.00 per glass she will buy only one glass; when it is £1.50 she will buy two glasses; and so on. The total consumption value is the area below her demand curve, and consumers' surplus is that part of the area that lies above the price line. This is shown in Figure (i).

Figure (ii) shows that the same relationship holds for the smooth market demand curve that indicates the total amount all consumers would buy at each price.*

* Figure 6.2 is a bar chart because we allowed the consumer to vary her consumption only in discrete units, one at a time. Had we allowed her to vary her consumption continuously, we could have traced out a continuous curve for Ms Green similar to the one shown in Figure 6.3.

Consumer optimization without measurable utility

The basic assumption here about consumer *motivation* is not changed from the last section. Consumers are assumed to maximize their satisfaction by allocating a given budget between the various goods and services that they wish to buy. Each consumer may be aware of exactly how much satisfaction is delivered by each of the goods consumed (though we do not need to assume this).

The key difference in this section is that, in explaining the consumer's behaviour, we do not need to know *how much* satisfaction he or she derives from consuming each product—nor indeed does the consumer need to know this. All that is needed is that each consumer can order any two bundles of goods by saying which gives more satisfaction and hence is the preferred bundle. Faced with a choice between many bundles, the maximizing consumer will then chose the one in the highest rank order of preference—and hence the one that is the most preferred of all available bundles.[6]

[6] This approach was originally due to the Italian economist Vilfredo Pareto (1848–1923). It was introduced to the English-speaking world (and greatly elaborated) by two British economists, John Hicks (1904–89) and R. G. D. Allen (1906–83).

First, we ask how we can find the consumer's equilibrium allocation of spending in this new framework. Once that is done, we will be able to study consumers' responses to changes in such things as prices and incomes.

The consumer's preferences

In the analysis that we are about to develop, the consumer's tastes or preferences, as they are variously called, are shown by indifference curves.

A single indifference curve

We start by deriving a single indifference curve. To do this we give an imaginary consumer, Kevin, some quantity of each of two products, say 18 units of clothing (C) and 10 units of food (F). This bundle is plotted as point b in Figure 5.1. Now think about the alternative combinations of these two products in the two shaded areas created by drawing vertical and horizontal lines through b. Would Kevin prefer the bundles of goods in these two shaded areas? To help answer this, we introduce our first assumption about tastes.

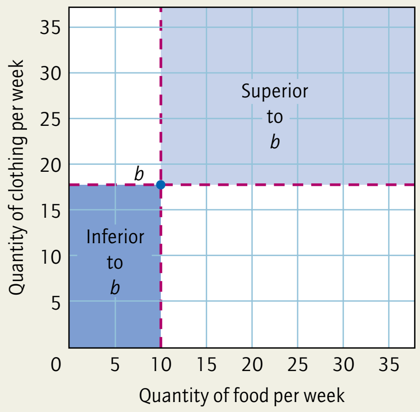

Figure 5.1 Some consumption bundles compared

According to assumption 1, bundle *b* is superior to bundles that have less of both goods and inferior to all bundles that have more of both. All points in the dark blue area are regarded as inferior to bundle *b* because they contain less of both commodities (except on the boundaries, where they have less of one and the same amount of the other).

Table 5.1 Bundles conferring equal satisfaction

Bundle	Clothing	Food
a	30	5
b	18	10
c	13	15
d	10	20
e	8	25
f	7	30

Since each of these bundles gives Kevin equal satisfaction, he is indifferent between them. None of the bundles contains more food and more clothing than any of the other bundles. Kevin's assumed indifference among these bundles is not, therefore, in conflict with the assumption that more is preferred to less of each product.

***Assumption 1*. Other things being equal, the consumer always prefers more of any one product to less of that same product.**

This allows us to rank the bundles of goods represented by the two shaded areas in Figure 5.1. Combinations on the edges of this space to the north-east of point *b* all have more of one good and no less of the other, while points inside this area represent bundles containing more of both goods. All points in this space, apart from *b* itself, will thus be preferred to *b*. By similar logic, all points to the south-west of *b* represent either fewer of both goods or fewer of at least one and no more of the other. These points will all be inferior for the consumer as they deliver a lower level of satisfaction.

But what about bundles that have more of some products and less of others? At point *b*, Kevin consumes 18 units of clothing and 10 of food. Let us ask how much extra clothing we would have to give him to make him equally satisfied if we took away one unit of food. The answer might be that 20 units of clothing and 9 units of food would leave Kevin just as satisfied as with the initial combination. If we do this again, taking away another unit of food, there will be some further increase in clothing that could just compensate. Table 5.1 shows that if we have taken away 5 units of food, Kevin will require 30 units of clothing to leave him feeling just as satisfied as at point *b*. This is also illustrated by point *a* in Figure 5.2. These combinations of fewer units of food and increased

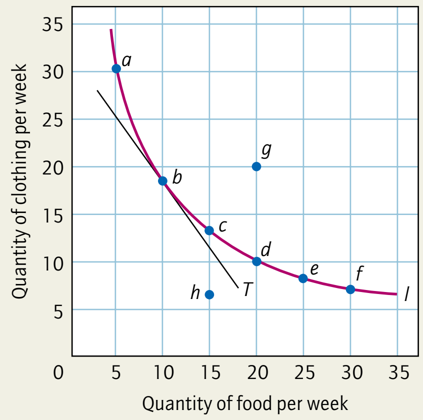

Figure 5.2 An indifference curve

The indifference curve shows combinations of food and clothing that yield equal satisfaction and among which the consumer is indifferent. Points *a* to *f* are plotted from Table 5.2 and an indifference curve is drawn through them. Compared with any point on the curve, point *g* is superior while point *h* is inferior. The slope of the tangent *T* gives the marginal rate of substitution at point *b*. Moving down the curve from *b* to *f*, the slope of the tangent flattens, showing that the more food and the less clothing Kevin has, the less willing he will be to sacrifice further clothing to get more food.

quantities of clothing that leave Kevin just as satisfied trace out the line segment from *b* to *a* in the figure.

Starting again at point *b*, we can now move in the opposite direction and ask how much extra food would Kevin need to leave him equally satisfied as we take successive units of clothing away from him. The answer to this question traces out the line through points *c, d, e,* and *f*.

By construction, the curved line drawn out in Figure 5.2 shows combinations of clothing and food all of which give Kevin the same level of satisfaction. He is indifferent between all of the different bundles of goods represented by that line (some specific combinations of which are listed in Table 5.1). For this reason this red line is called an **indifference curve**. The line joining points *a–f* in Figure 5.2 is one indifference curve.

An indifference curve shows combinations of products that yield the same satisfaction to the consumer. Thus, a consumer is indifferent between the combinations indicated by any two points on one indifference curve.

Points above and to the right of the indifference curve in Figure 5.2 show combinations of food and clothing that Kevin would prefer to combinations indicated by points on the curve. Consider, for example, the combination of 20*F* and 20*C*, which is represented by point *g* in the figure. Although it might not be obvious that this bundle is preferred to bundle *a* (which has more clothing but less food), Assumption 1 tells us that *g* is preferred to bundle *c*, because *g* has more clothing *and* more food than *c*. Inspection of the graph shows that *any* point above the curve will be obviously superior to *some* points on the curve in the sense that it will contain both more food and more clothing than those points on the curve. But since all points on the curve are equally valuable in Kevin's eyes, any point above the curve must, therefore, be superior to *all* points on the curve. By a similar argument, points such as *h*, which are below and to the left of the curve, represent bundles of goods that Kevin regards as inferior to all bundles on the curve. These comparisons are summarized in Figure 5.3.

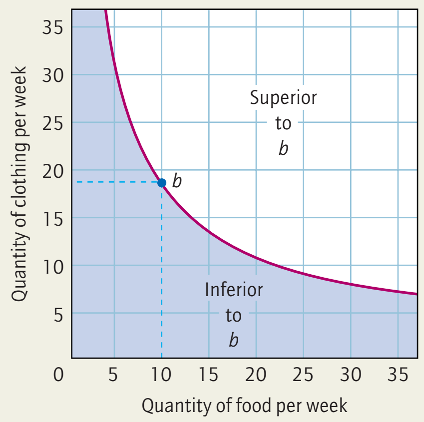

Figure 5.3 Consumption bundles compared

The indifference curve allows any bundle such as *b* to be compared with all others. Kevin regards all bundles in the dark blue area as inferior and all bundles in the light blue area as superior to *b*. The indifference curve is the boundary between these two areas. All points on the curve yield equal satisfaction and Kevin is, therefore, indifferent among them.

Diminishing marginal rate of substitution

What is the shape of a typical indifference curve? To answer this we need a second assumption.

***Assumption 2*. The less of one product that is presently being used by a consumer, the smaller the amount of it that the consumer will be willing to forgo in order to increase consumption of a second product.**

This is called the assumption of a **diminishing marginal rate of substitution**. The *rate of substitution* tells how much more of one product we need to compensate for successive lost units of the other. The *diminishing* of this rate of substitution may seem intuitively akin to diminishing marginal utility; however, for the latter we hold consumption of all but one good constant, while here we have more of one good compensating for fewer of the other.

Diminishing marginal rate of substitution is illustrated in Table 5.2, which is based on the example of food and clothing shown in Table 5.1. As we move down the table through points *a* to *f*, Kevin has bundles with fewer and fewer units of clothing and more and more food. In accordance with the hypothesis of diminishing marginal rate of substitution, he is willing to give up smaller and smaller amounts of clothing to further increase his consumption of food by one unit. When Kevin moves from *c* to *d*, for

Table 5.2 Diminishing marginal rate of substitution

Movement	Change in clothing (1)	Change in food (2)	Marginal rate of substitution (3)
From *a* to *b*	−12	5	2.4
From *b* to *c*	−5	5	1.0
From *c* to *d*	−3	5	0.6
From *d* to *e*	−2	5	0.4
From *e* to f	−1	5	0.2

The marginal rate of substitution measures the amount of one product a consumer must be given to compensate for giving up one unit of the other. This table is based on the data in Table 5.1. When Kevin moves from *a* to *b*, he gives up 12 units of clothing and gains 5 units of food, a rate of substitution of 12/5 or 2.4 units of clothing sacrificed per unit of food gained. When he moves from *b* to *c*, he sacrifices 5 units of clothing and gains 5 of food (a rate of substitution of 1 unit of clothing for each unit of food). Note that the marginal rate of substitution (MRS) is the absolute value of the ratio of ΔC to ΔF. Since these two changes always have opposite signs, the MRS is obtained by multiplying this ratio by −1.

example, the table tells us that he is prepared to give up 0.6 unit of clothing to get a further unit of food. When he moves from *e* to *f*, he will give up only 0.2 unit.

The geometrical expression of this hypothesis is found in the shape of the indifference curve. Look closely, for example, at the slope of the curve in Figure 5.2. Its negative slope indicates that, if Kevin is to have fewer units of one product, he must have more of the other to compensate. Diminishing marginal rate of substitution is shown by the curve being convex viewed from the origin: moving down the curve to the right, its slope gets flatter and flatter. The absolute value of the slope of the curve is the marginal rate of substitution, the rate at which the consumer is willing to reduce his consumption of the product plotted on the vertical axis in order to increase his consumption of the product plotted on the horizontal axis.

The slope of the indifference curve at any point is measured by the slope of the tangent to the curve at that point. The slope of tangent *T* drawn to the curve at point *b* shows the marginal rate of substitution at that point. It can be seen that, moving down the curve to the right, the slope of the tangent gets flatter and flatter, and hence the marginal rate of substitution is diminishing.[7]

The indifference map

So far we have constructed only a single indifference curve. There must, however, be a similar curve passing through any of the other points in Figures 5.2, in addition to those points on the single curve drawn. Starting at another point, such as *g*, and going through the same exercise, there will be other combinations that will yield Kevin equal satisfaction. If the line joining all of *these* combinations is drawn, another indifference curve will be constructed. This exercise can be repeated many times, generating a new indifference curve each time.

It follows from the comparisons given in Figure 5.3 that the further away any indifference curve is from the origin, the higher is the level of satisfaction given by the consumption bundles that it indicates. We refer to a curve that confers a higher level of satisfaction as a *higher curve*.

A set of indifference curves is called an **indifference map**. An example is shown in Figure 5.4. It specifies Kevin's tastes by showing his complete ordering of preferences between different bundles of these two products, and it shows his rate of substitution between them at each specific point. When economists say that a consumer's tastes are *given*, they do not mean merely that the consumer's current consumption pattern is given: rather, they mean that the consumer's entire indifference map is given.

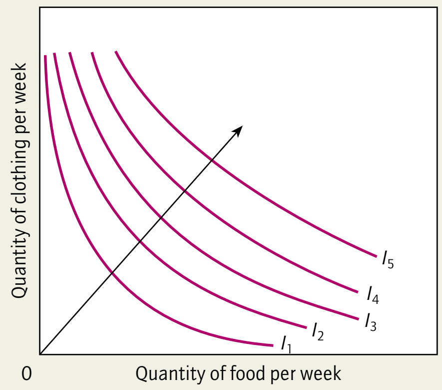

Figure 5.4 An indifference map

A set of indifference curves is called an indifference map. The further the curve from the origin, the higher the level of satisfaction it represents. If Kevin moves along the arrow, he is climbing a 'utility mountain', moving to ever-higher utility levels and crossing ever-higher equal-utility contours, which we call *indifference curves*.

Of course, there must be an indifference curve through *every* point in Figure 5.4. To graph them, we only show a few, but all are there. Thus, as Kevin moves upwards to the right starting from the origin, his utility is rising continuously. As he follows a route such as the one shown by the arrow, consuming ever more of both products, he can be thought of as climbing a continuous utility mountain. We show this 'mountain' by selecting a few equal-utility contours, labelled I_1 to I_5. But every point between each of the contours shown must also have a curve of equal utility passing through it. Thus, an indifference map is really like the continuous surface of one half of a cone, rather than a set of discrete lines.

In indifference theory we do not need to make any assumptions about how big the difference is between the level of satisfaction on one indifference curve and the next; i.e., we do not need to assume that utility can be quantified. Instead, all we assume is that the utility attached to I_5 exceeds that attached to I_4, which in turn exceeds the utility attached to I_3, and so on. We can say that the consumer is climbing a utility mountain as he moves along the arrow starting from the origin, but we do not need to know if the mountain is gentle or steep.

Box 5.3 shows some specific shapes of indifference curves that correspond to some specific taste patterns.

[7] Table 5.2 calculates the rate of substitution between distinct points on the indifference curve. Strictly speaking, these are the incremental rates of substitution between the two points. Geometrically, the slope of the chord joining the two points gives this incremental rate. The marginal rate refers to the slope of the curve at a single point and is given by the slope of the tangent to the curve at the point.

Box 5.3 Shapes of indifference curves

Any taste pattern can be illustrated with indifference curves. This box gives a few examples that will help you understand how indifference curves work. In each case the curve labelled I_2 indicates a higher utility than the curve labelled I_1.

Perfect substitutes: part (i) Drawing pins that came in red packages of 100 would be perfect substitutes for identical pins that came in green packages of 100 for a colour-blind consumer: he would be willing to substitute one type of package for the other at a rate of one for one. The indifference curves would thus be a set of parallel lines with a slope of −1, as shown in part (i) of the figure. *Indifference curves for perfect substitutes are straight lines whose slopes indicate the rate at which one good can be substituted for the other.*

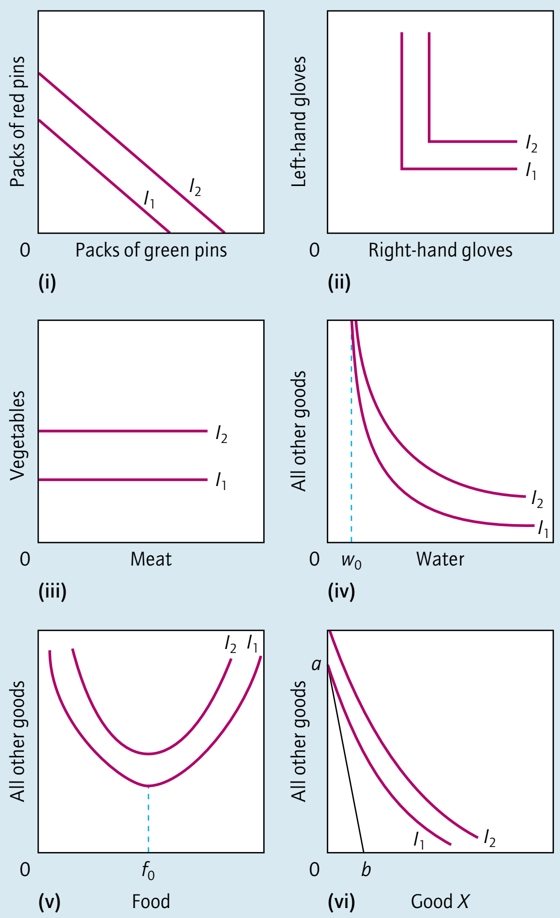

Perfect complements: part (ii) Left- and right-hand gloves are perfect complements, since one of them is of no use without the other. This gives rise to the indifference curves shown in part (ii) of the figure. There is no rate at which any consumer will substitute one kind of glove for the other when she starts with equal numbers of each. *Indifference curves for perfect complements are L-shaped.*

A good that gives zero utility: part (iii) When a good gives no satisfaction at all, a person will be unwilling to sacrifice even the smallest amount of other goods to obtain any quantity of the good in question. Such would be the case regarding meat for a vegetarian consumer, whose indifference curves are horizontal straight lines. *Indifference curves for a product yielding zero satisfaction are parallel to that product's axis.*

An absolute necessity: part (iv) There is some minimum quantity of water, w_0, that is necessary to sustain life. As consumption of water falls towards w_0, increasingly large amounts of other goods are necessary to persuade the consumer to cut down on his water consumption. Thus, each indifference curve becomes steeper and steeper as it approaches w_0, and the marginal rate of substitution increases. *The marginal rate of substitution for an absolute necessity approaches infinity as consumption falls towards the amount that is absolutely necessary.*

A good that confers a negative utility after some level of consumption: part (v) Beyond some point, further consumption of many foods and beverages, films, plays, or cricket matches would reduce satisfaction. Figure (v) shows a consumer who is *forced* to eat more and more food. At the amount f_0 she has all the food she could possibly want. Beyond f_0 her indifference curves have positive slopes, indicating that she gets *negative* value from consuming the extra food, and so will be willing to sacrifice some amount of other products to avoid consuming it. *When, beyond some level of consumption, the consumer's utility is reduced by further consumption, the indifference curves have positive slopes.*

This case does not arise if the consumer can dispose of the extra unwanted units at no cost. The indifference curves then become horizontal.

A good that is not consumed: part (vi) Typically, a consumer will consume only one or two of all of the available types of cars, TV sets, dishwashers, or tennis rackets. If a consumer is in equilibrium when consuming a zero amount of say, green peas, she is in what is called a *corner solution* (as shown in part (vi) of the figure by the budget line *ab* and the curve I_1). *When a good is not consumed, the indifference curve cuts the axis of the non-consumed good with a slope flatter than the budget line.*

Table 5.3 Data for Jane's budget line

Quantity of food	Value of food	Quantity of clothing	Value of clothing	Total expenditure
60	£120	0	£0	£120
50	100	5	20	120
40	80	10	40	120
30	60	15	60	120
20	40	20	80	120
10	20	25	100	120
0	0	30	120	120

The table shows combinations of food and clothing available to Jane when her income is £120 and she faces prices of £4 for clothing and £2 for food. Any row indicates a bundle of food and clothing that exactly exhausts Jane's income.

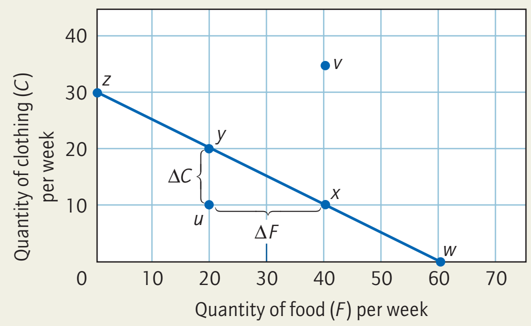

Figure 5.5 Jane's budget line

The budget line shows the quantities of goods available to Jane, given her money income and the price of the goods she buys. With an income of £120 a week and prices of £2 per unit for food and £4 per unit of clothing, the coloured line is Jane's budget line, showing all combinations of *F* and *C* that are obtainable. Bundle *u* (10*C* and 20*F*) does not use all of her income. Bundle *v* (35*C* and 40*F*) requires more than her present income.

If Jane moves from point *y* (20*F* and 20*C*) to point *x* (40*F* and 10*C*), she consumes 20 more *F* and 10 fewer *C*. These amounts are indicated by ΔF and ΔC in the figure. Thus, the opportunity cost of each unit of *F* added to consumption is 10/20 = 0.5 unit of clothing forgone. This is the absolute value of $\Delta C/\Delta F$, which is the slope of the budget line *zw* in the figure.

The choices available to the consumer

An indifference map tells us what any consumer *would like* to do: reach the highest possible indifference curve, that is, be as high up the utility mountain as possible. To see what that consumer *can* do, we need another construction, called the budget line.

We start by considering a single consumer, Jane, who is allocating the whole of her money income between two goods, food and clothing.[8]

The budget line

The **budget line** shows all those combinations of the goods that are just obtainable, given Jane's income and the prices of the products that she buys.[9]

Assume initially that Jane's income is £120 per week, the price of food is £2 per unit, and the price of clothing is £4 per unit. As in the earlier discussion, we denote food by *F* and clothing by *C*. Thus, for example, a bundle containing 20 units of food and 10 units of clothing is written as 20*F* and 10*C*. Table 5.3 lists a few of the bundles of food and clothing available to Jane, while the blue line running from *z* to *w* in Figure 5.5 shows all the possible bundles that she could buy with her income. At point *w*, for example, Jane is spending all her income to buy 60*F* and no clothing, while at point *z* she is spending all her income to buy 30*C* and no food. Points on the line between *z* and *w* indicate how much Jane could buy of both products.

The slope of the budget line

Marked on Figure 5.5 as points *x* and *y* are two of the specific spending combinations from Table 5.3. It is clear from the figure that the absolute value of the slope of the budget line measures the ratio of the change in *C* to the change in *F* as we move along the line. This ratio, $\Delta C/\Delta F$, is 0.5 in our present example (10/20).

How does the slope of the budget line relate to the prices of the two goods? This question is easily answered if we remember that all points on the budget line represent bundles of goods that just exhaust Jane's whole income. It follows that, when she moves from one point on the budget line to another, the change in expenditure on *C* must be of equal value, but opposite in sign, to the change in expenditure on *F*. Letting ΔC and ΔF stand for the changes in the quantities of clothing and food respectively, and p_c and p_f stand for the money prices of

[8] These assumptions are not as restrictive as they at first seem. Although just two goods are used so that the analysis can be handled graphically, the argument can easily be generalized to any number of goods with the use of mathematics. Savings are ignored because we are interested in the allocation of expenditure among commodities for current consumption. Saving and borrowing can be allowed for, but doing so affects none of the results in which we are interested here.

[9] A budget line is analogous to the production possibility boundary shown in Figure [1.1 on page ?]. The budget line shows the combinations of commodities available to one consumer given her income and prices, while the production possibility curve shows the combination of commodities available to the whole society given its supplies of resources and techniques of production.

clothing and food respectively, we can write this relation as follows:

$$\Delta C p_c = -\Delta F p_f.$$

There is nothing difficult in this. All it says is that, if any amount more is spent on one product, the same amount less must be spent on the other. A given income imposes this discipline on any consumer.

If we divide the above equation through, first by ΔF, and then by p_c, we get the following:

$$\frac{\Delta C}{\Delta F} = \frac{p_f}{p_c}.$$

So the slope of the budget line is the negative of the ratio of the two prices (with the price of the good that is plotted on the horizontal axis appearing in the numerator).

Notice that the slope of the budget line depends only on the ratio of the two prices, not on their absolute values. To check this, consider an example. If clothing costs £4 and food costs £2, then Jane must forgo 0.5 unit of clothing in order to be able to purchase one more unit of food. If clothing costs £8 and food costs £4, Jane must still forgo 0.5 unit of clothing to be able to purchase one more unit of food. As long as the price of clothing is twice the price of food, Jane must forgo half a unit of clothing in order to be able to purchase one more unit of food.

More generally, the amount of clothing that must be given up to obtain another unit of food depends only on the *ratio* of their two prices. If we take the money price of food and divide it by the money price of clothing, we have the opportunity cost of food in terms of clothing (the quantity of clothing that must be forgone in order to be able to purchase one more unit of food). This may be written:

$$\frac{p_f}{p_c} = \text{opportunity cost of food in terms of clothing.}$$

It is apparent that changing income and/or changing both prices in the same proportion leaves the ratio p_f/p_c unchanged.

This discussion helps to clarify the distinction between money prices and relative prices. Both p_f and p_c are money prices, while the ratio p_f/p_c is a relative price.

The consumer's equilibrium

The budget line tells us what consumers *can* do: they can select any consumption bundle on, or below, the line—but not above it. This means that they can spend only within the limits of a given income. To see what consumers *want* to do, we introduce our third assumption:

***Assumption 3*. Consumers seek to maximize total satisfaction, which means reaching the highest possible indifference curve.**

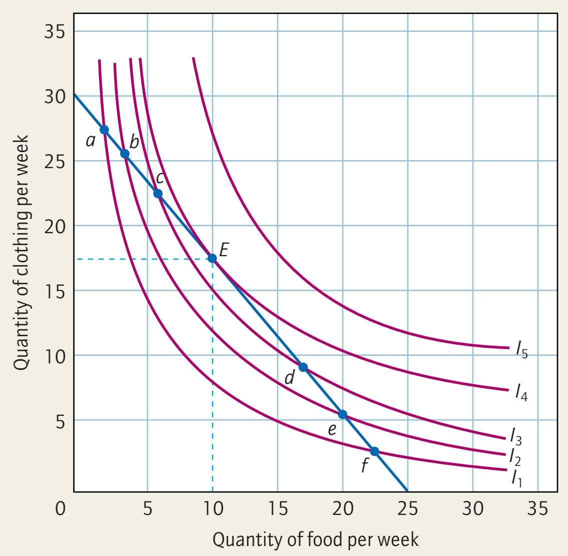

Figure 5.6 The equilibrium of a consumer

Equilibrium occurs at *E*, where an indifference curve is tangent to the budget line. Paul has an income of £150 a week and faces prices of £5 a unit for clothing and £6 a unit for food. A bundle of clothing and food indicated by point *a* is attainable, but by moving along the budget line to points such as *b* and *c*, higher indifference curves can be reached. At *E*, where the indifference curve I_4 is tangent to the budget line, Paul cannot reach a higher curve by moving along the budget line. If he did alter his consumption bundle by moving from *E* to *d*, for example, he would move to the lower indifference curve I_3 and thus to a lower level of satisfaction.

We have now developed representations of a consumer's tastes and available choices. Figure 5.6 brings together the budget line and the indifference curves for another consumer, Paul. Any point on the budget line can be attained. Which one will Paul actually choose?

Will Paul choose to consume 25 units of food and no clothing, as he could do with his income? Will he instead choose to consume 30 units of clothing and no food? The answer is no in both cases. By moving away from either of these combinations, he can move to a higher indifference curve. Indeed, he can get to higher and higher indifference curves by moving from each of the corners into the middle until he reaches point *E*, which is just touching—i.e. is tangent to—the highest possible indifference curve. When Paul is at this point of tangency between the indifference curve and the budget line, he cannot reach a higher indifference curve by varying the bundle consumed. Any move from this point that remains within the budget constraint will lead him to a lower indifference curve and thus to lower satisfaction.

Satisfaction is maximized at the point where an indifference curve is tangent to a budget line. At that point, the slope of

the indifference curve—which measures the consumer's marginal rate of substitution—is equal to the slope of the budget line—which measures the opportunity cost of one good in terms of the other as determined by market prices.

Notice that Paul is presented with market prices that he cannot change. He adjusts to these prices by choosing a bundle of goods such that, at the margin, his own relative valuation of the two goods conforms to the relative valuations given by the market. Paul's relative valuation is given by the slope of his indifference curve, while the market's relative valuation is given by the slope of his budget line.

When Paul has chosen the consumption bundle that maximizes his satisfaction, he will go on consuming that bundle unless something changes. The consumer is thus in equilibrium.

It is also worth noting that the equilibrium position we have just derived has the same characteristics as the one the utilitarians discovered and is expressed in equations (1) and (2) above. The price ratio p_x/p_y is the slope of the budget line in Figure 5.6. The slope of each indifference curve, which we have called the *marginal rate of substitution*, is the ratio of the marginal utilities of the two products, MU_x/MU_y; and so where the budget line is tangent to the highest possible indifference curve (i.e. where the consumer is maximizing utility) it will also be true that $MU_x/MU_y = p_x/p_y$.

The consumer's response to price and income changes

How do consumers change their spending patterns when there is a change in goods prices or in available income? To answer this, we take another hypothetical consumer called Karen. Karen's tastes are given, and are represented by an indifference map that does not change. We first show that changes in her income and the prices she faces can be represented as a shift in the budget line. We then investigate the change in spending induced by price and income changes.

Parallel shifts in the budget line

A change in money income

A change in Karen's money income will, other things being equal, shift her budget line. For example, if income rises Karen will be able to buy more of both goods. Her budget line will, therefore, shift out parallel to itself to indicate this expansion in her consumption possibilities. (The fact that it will be a parallel shift is established by our demonstration on page 114 that the slope of the budget line depends only on the relative price of the two products.)

A change in the consumer's income shifts the budget line parallel to itself—outwards when income rises and inwards when income falls.

The effect of income changes is shown in Figure 5.7. For each level of income, there is an equilibrium position at which an indifference curve is tangent to the relevant budget line. Each such equilibrium position means that Karen is doing as well as she possibly can for that level of income. If we join up all the points of equilibrium, we trace out what is called her **income–consumption line**. This line shows how the consumption bundle changes as income changes, with prices held constant.[10]

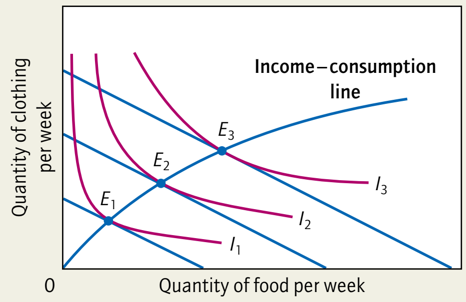

Figure 5.7 An income–consumption line

This line shows how Karen's purchases react to changes in income with relative prices held constant. Increases in income shift the budget line out parallel to itself, moving the equilibrium from E_1 to E_2 to E_3. The blue income–consumption line joins all these points of equilibrium.

[10] This income–consumption line can be used to derive the curve relating quantity demanded to income that was introduced on page [XX]. This is done by plotting the quantity of one of the goods consumed at the equilibrium position against the level of money income that determined the position of the budget line. Repeating this for each level of income produces the required curve.

A proportionate change in all prices

If all prices are cut in half, Karen can buy twice as much of both products. This causes the same shift in the budget line as when Karen's income doubles with prices held constant. On the other hand, a doubling of all prices will cause her budget line to shift inwards in exactly the same way as if her money income had halved with prices held constant.

This illustrates a general result:

An equal proportionate change in all money prices, with money income held constant, shifts the budget line parallel to itself—towards the origin when prices rise and away from the origin when prices fall.

From this point on the analysis is the same as in the previous section, since changing money prices proportionately has the identical effect to changing money income.

Offsetting changes in money prices and money incomes

The results in the last two sections suggest that we can have offsetting changes in money prices and money incomes. Consider a doubling of Karen's money income, shifting her budget line outwards. Let this be accompanied by a doubling of all money prices, which shifts her budget line inwards. The net effect is to leave her budget line where it was before the changes in her income and in the market prices. This illustrates a general result:

Multiplying money income by some constant λ, and simultaneously multiplying all money prices by λ, leaves the budget line unaffected and hence leaves consumer purchases unaffected.

The symbol λ is the lower-case Greek letter lambda, which is often used for some constant multiple. This result is sometimes referred to as the *homogeneity condition*.

Changes in the slope of the budget line

A change in relative prices

We already know that a change in the relative prices of the two goods changes the slope of the budget line. At a given price of clothing, Karen has an equilibrium consumption position for each possible price of food. Connecting these positions traces out a **price–consumption line**, as is shown in Figure 5.8. Notice that, as the relative prices of food and clothing change, the relative quantities of food and clothing purchased also change. In particular, as the price of food falls, Karen buys more food.

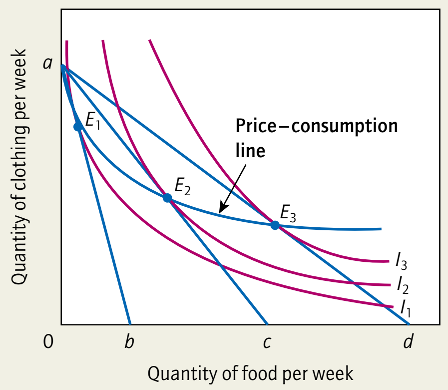

Figure 5.8 The price–consumption line

This line shows how a consumer's purchases react to a change in one price with money income and other prices held constant. Decreases in the price of food (with her money income and the price of clothing constant) pivot Karen's budget line from *ab* to *ac* to *ad*. Her equilibrium position moves from E_1 to E_2 to E_3. The blue price–consumption line joins all such equilibrium points.

Real and money income

The preceding analysis allows us to look deeper into the important distinction between two concepts of income. **Money income** measures a consumer's income in terms of some monetary unit, for example so many pounds sterling or so many dollars. **Real income** measures the *purchasing power* of the consumer's money income. A rise in money income of *x* per cent combined with an *x* per cent rise in all money prices leaves a consumer's purchasing power, and hence their real income, unchanged. When we speak of the *real value* of a certain amount of money, we are referring to the goods and services that can be bought with the money, that is, the purchasing power of the money.

Box 5.4 discusses the importance of relative prices and the problems created by inflation.

Box 5.4 Relative prices and inflation

Allocation of resources: the importance of relative prices

Price theory shows why the allocation of resources depends on the structure of relative prices. If the money value of all prices, incomes, debts, and credits were doubled, there would, according to our theory, be no noticeable effects. We have already seen that doubling money income and all money prices leaves each consumer's budget line unchanged. So, according to the theory of consumer behaviour, the combination of these changes gives the consumer no incentive to vary any purchases. As far as producers are concerned, if the prices of all outputs and inputs double, the relative profitabilities of alternative lines of production will be unchanged. Thus, producers will have no incentive to alter production rates so as to produce more of some things and fewer of others. The same set of relative prices and real incomes would exist, and there would be no incentive for any reallocation of resources. The economy would function as before.

In contrast, a change in *relative* prices will cause resources to be reallocated. Consumers will buy more of the relatively cheaper products and less of the relatively more expensive ones, and producers will increase production of those products whose prices have risen relatively, and reduce production of those whose prices have fallen relatively (since the latter will be relatively less profitable lines of production).

The theory of price and resource allocation is a theory of relative, not absolute, prices.

Inflation and deflation: the importance of absolute prices

The average level of all money prices is called the general price level, or more usually just the **price level**. If all money prices double, we say that the price level has doubled. An increase in the price level is called an **inflation**; a decrease is called a **deflation**. If a rise in all money prices and incomes has little or no effect on the allocation of resources, it may seem surprising that so much concern is expressed over inflation. Clearly, people who spend all their incomes, and whose money incomes go up at the same rate as money prices, lose nothing from inflation. Their real income is unaffected.

Inflation, while having no effect on consumers whose incomes rise at the same rate as prices, none the less does have many serious consequences. These arise mainly because all prices do not rise at the same rate and some assets are denominated in money terms, so that their value falls as the price level rises. These consequences are studied in detail later in this book. In the meantime, *we assume that the price level is constant.*

Under these circumstances, a change in one money price necessarily changes that price *relative to* the average of all other prices. The theory extends to situations in which the price level is changing. Under inflationary conditions, whenever shifts in demand or supply require a change in a product's relative price, its price rises *faster* (its relative price rising) or *slower* (its relative price falling) than the general price level is rising. Explaining this each time can be cumbersome. It is, therefore, simpler to deal with relative prices in a theoretical setting in which the price level is constant. It is important however to realize that, even though we develop the theory in this way, it is not limited to such situations. The propositions we develop can be applied to changing price levels merely by making explicit what is always implicit: in the theory of relative prices, 'rise' or 'fall' *always* means rise or fall *relative to the average of all other prices*.

The consumer's demand curve

We now establish the link between the above analysis of indifference curves and budget constraints, and the consumer's demand curve. To derive the consumer's demand curve for any product, we need to depart from the world of two products. We are now interested in what happens to the consumer's demand for some product, say petrol, as the price of that product changes, *all other prices being held constant*. We can do this with the tools developed above simply by making the bundle of 'all other goods' take the place of the second product.

Derivation of the demand curve

In part (i) of Figure 5.9 a new type of indifference map is plotted in which the horizontal axis measures litres of petrol and the vertical axis measures the value of all other goods consumed. We have in effect used *everything but petrol* as the second product. The indifference curves now give the rate at which another hypothetical consumer, Philip, is prepared to substitute petrol for money (which allows him to buy all other goods).

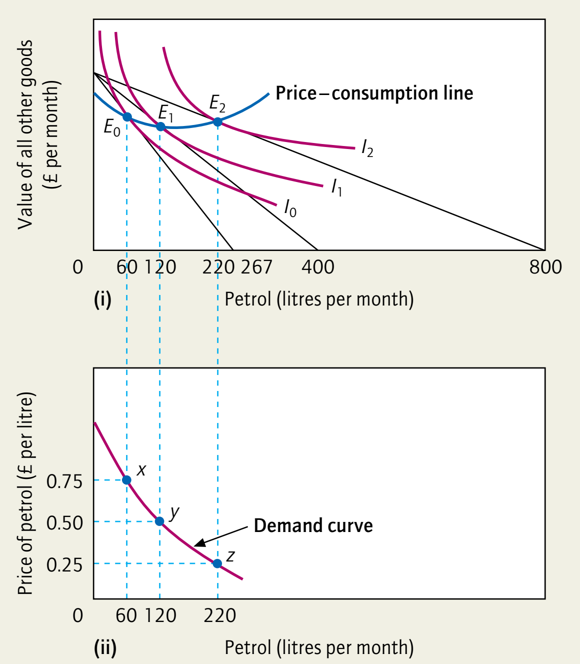

Figure 5.9 Derivation of an individual's demand curve

The points on a price–consumption line provide the information needed to draw a demand curve. In part (i) Philip has an income of £200 per month and alternatively faces prices of £0.75, £0.50, and £0.25 per litre of petrol, choosing positions E_0, E_1, and E_2. The information concerning the number of litres he demands at each price is then plotted in part (ii) to yield his demand curve. The three points *x*, *y*, and *z* in (ii) correspond to the three equilibrium positions E_0, E_1, and E_2 in (i).

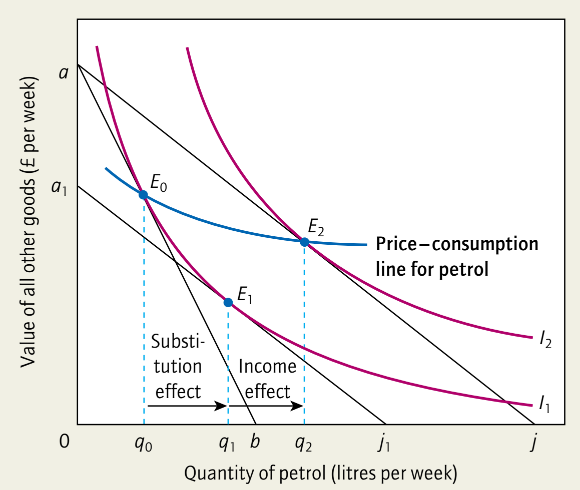

Figure 5.10 The income and substitution effects

The substitution effect is defined by sliding the budget line around a fixed indifference curve; the income effect is defined by a parallel shift of the budget line. The original budget line is *ab* and a fall in the price of petrol takes it to *aj*. The original equilibrium is at E_0 with q_0 of petrol consumed, and the final equilibrium is at E_2 with q_2 of petrol consumed. To remove the income effect, imagine reducing Philip's income until he is just able to attain his original indifference curve at the new price. We do this by shifting the line *aj* to a parallel line nearer the origin until it just touches the indifference curve that passes through E_0. The intermediate point E_1 divides the quantity change into a substitution effect, $q_1 - q_0$, and an income effect, $q_2 - q_1$. The point E_1 can also be obtained by sliding the original budget line *ab* around the indifference curve until its slope reflects the new relative prices.

The derivation of a demand curve is illustrated in part (ii) of Figure 5.9. For a given income, each price of petrol gives rise to a particular budget line and a particular spending choice. Plotting the quantity of petrol that Philip consumes for the specific budget line at any given price yields one point on his demand curve. Every other possible price yields a different point. The resulting price–quantity combinations trace out Philip's whole demand curve.

The slope of the demand curve

The price–consumption line in part (i) of Figure 5.9 indicates that, as price decreases, the quantity of petrol demanded increases. But it is possible to draw Philip's indifference curves in such a way that his response to a decrease in price is for less to be consumed rather than more. Such a positively sloped demand curve for a good is referred to as a **Giffen good** after the Victorian economist Sir Robert Giffen (1837–1910), who is reputed to have documented a case of such a curve. We now show how this case can be analysed using indifference curves.

Income and substitution effects

The key is to distinguish between *the income effect* and *the substitution effect* of a change in price. The separation of the two effects according to indifference theory is shown in Figure 5.10. We can think of it as occurring in the following way. After the price of the good has fallen, we reduce money income *until the original indifference curve can just be obtained*. Philip is now on his original indifference curve but facing the new set of relative prices. His response is defined as the **substitution effect:** the response of quantity demanded to a change in relative price, real income being held constant (i.e. staying on the original indifference curve). Then, to measure the income effect, we restore money income. Philip's response to this is defined as the **income effect:** the response of quantity

Box 5.5 The Slutsky decomposition of income and substitution effects

The discussion of income and substitution effects in the text is based upon the analysis of English Nobel Laureate Sir John Hicks (1904–89). An alternative approach was developed by the Russian mathematician Evgeny Slutsky (1880–1948).

Hicks's decomposition was derived in the context of developing the concept of indifference curves, so it was natural for him to ask the question: following a price change, how much income must be taken away in order that the consumer can return to the original indifference curve and thus have the same level of utility or satisfaction as prior to the price change?

Slutsky, when thinking about the same issue, did not have at his disposal the tool of indifference curves. Instead he asked the question: following a price change, how much income must be taken away so that the consumer is just able to buy the initial bundle of goods (and therefore could not be any worse off than in the initial position)?

The figure illustrates the difference between these two approaches. There is a fall in the price good 1 holding the price of good 2 constant. The initial consumption points is at A and after the price fall the consumption point is at B.

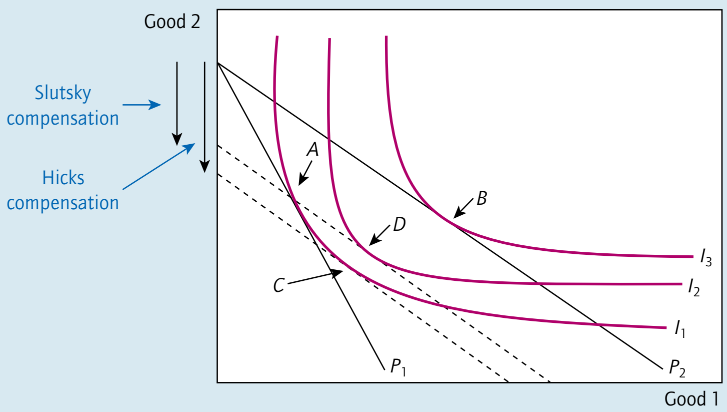

As we saw in the discussion of Figure 5.10, Hicks's decom position generates an income compensation that returns the consumer to the original indifference curve I_1 following the fall of price of good 1 and the associated shift of the budget constraint from P_1 to P_2. This is achieved by shifting the new budget line P_2 towards the origin until it is just tangent to the original indifference curve. Thus, the Hicks substitution effect takes the consumer from point A to point C, and the income effect takes her from C to B.

To find the Slutsky decomposition, we shift the new budget constraint inwards parallel to its new position until it just passes through the original consumption bundle at point A. If the consumer had faced this budget constraint with the original level of disposable income but at the new relative prices, she would have chosen to be at point D, which is on indifference curve I_2 and is thus at a higher utility level than the initial position.

There is no general reason why one of these methods is to be preferred to the others. They are answering slightly different questions. The Slutsky compensation is easier to calculate as it relies on observable income and prices, but the Hicks compensation is useful for welfare comparisons because it tells us the income change that leaves the consumer *feeling* just as well off as before. The choice of methods should therefore depend on the purpose to which it is put.

demanded to a change in real income, with relative prices held constant.

Box 5.5 explains an alternative method of isolating the income and substitution effects.

In Figure 5.10 the income and substitution effects work in the same direction, both tending to increase quantity demanded when price falls. Is this necessarily the case? The answer is no. It follows from the convex shape of indifference curves that the substitution effect is always in the same direction: more is consumed of a product whose relative price has fallen. The income effect, however, can be in either direction: it can lead to more or less being consumed of a product whose price has fallen. The direction of the income effect depends on the distinction between normal and inferior goods.

The slope of the demand curve for a normal good

For a *normal good*, an increase in any consumer's real income, arising from a decrease in the price of the product, leads to increased consumption, reinforcing the

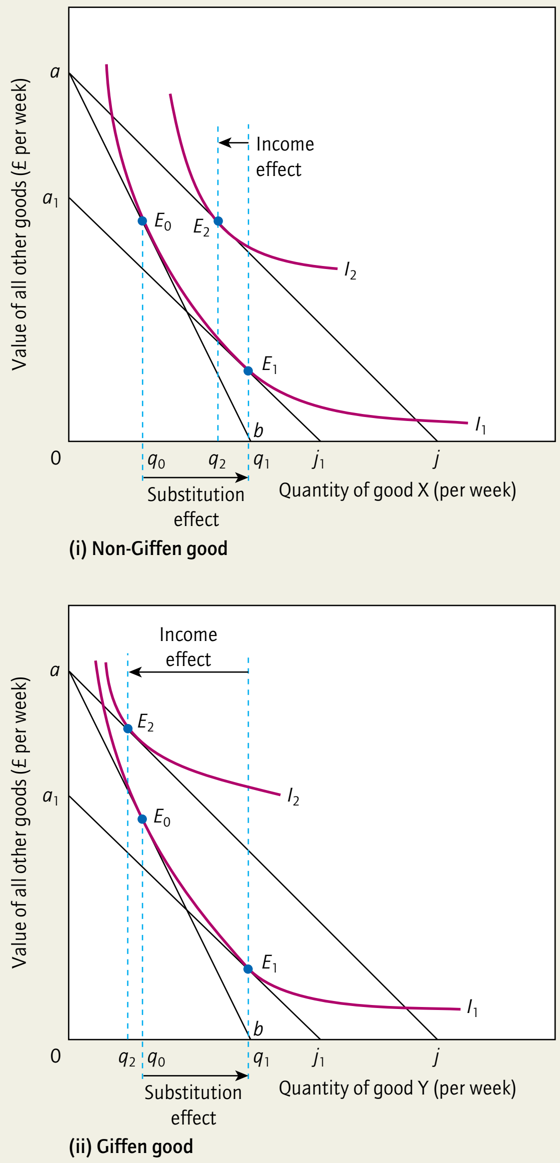

Figure 5.11 Income and substitution effects for inferior goods

A large enough negative income effect can outweigh the substitution effect and lead to a decrease in consumption in response to a fall in price. In each part of the diagram Philip is in equilibrium at E_0, consuming a quantity q_0 of the good in question. The price then decreases and the budget line shifts to *aj*, with a new equilibrium at E_2 and quantity consumed q_2. In each case the substitution effect increases consumption from q_0 to q_1. In (i) there is a negative income effect of $q_1 - q_2$. Because this is less than the substitution effect, the latter dominates, so good X has a normal, negatively sloped demand curve. In (ii) the negative income effect $q_1 - q_2$ is larger than the substitution effect, and quantity consumed actually decreases. Good Y is a Giffen good.

substitution effect. Because quantity demanded increases, the demand curve has a negative slope.[11] This is the case illustrated in Figure 5.9.

The slope of the demand curve for an inferior good

Figure 5.11 shows indifference curves for inferior goods. The income effect is negative in each part of the diagram. This follows from the nature of an *inferior good*: as income rises, less of the good is consumed. In each case the substitution effect serves to increase the quantity demanded as price decreases and is offset to some degree by the negative income effect. The final result depends on the relative strengths of the two effects. In part (i) the negative income effect only partially offsets the substitution effect, and thus quantity demanded increases as a result of the price decrease, though not as much as for a normal good. This is the typical pattern for inferior goods, and it too leads to negatively sloped demand curves, usually relatively inelastic ones.

In part (ii) the negative income effect outweighs the substitution effect and thus leads to a positively sloped demand curve. This is the Giffen case. For this to happen the good must be inferior. But that is not enough: the change in price must have a negative income effect *strong enough* to more than offset the substitution effect. These circumstances are unusual ones, because strong inferiority is rarely found. Such goods, if they ever existed, would tend to disappear from use as consumers get richer. Most goods are normal goods. A positively sloped market demand curve is thus a rare exception to the general rule that demand curves have negative slopes.

Equivalent and compensating variations

There are further concepts associated with the income effect of a price change that are commonly used in economics. These are known as **equivalent variation** and **compensating variation**.

[11] A possible exception to this arises from the *endowment income effect*. This arises in some models where the consumer is assumed to have an initial endowment of goods and may choose to be a net seller of some goods. If the price of these goods rises, the consumer has a higher income and therefore can buy more of all normal goods, including the goods for which he or she is a net seller. A practical example would be as follows. Suppose the price of haircuts rises (all other price remaining constant). For most consumers we would predict that the quantity demanded of haircuts would fall. However, hairdressers are now richer, so for them the price rise has generated a positive rather negative income effect. So for the hairdressers the income effect goes the other way, inasmuch as incomes rise as the price of haircuts rises.

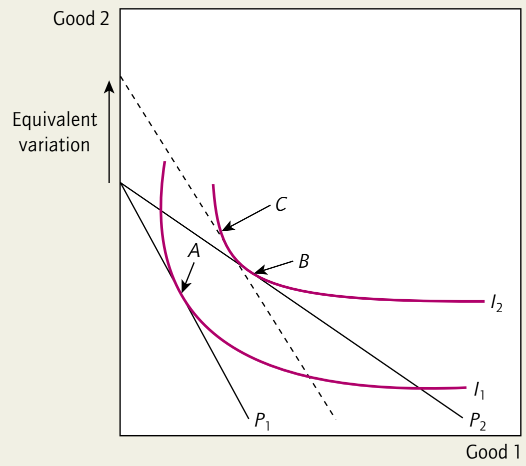

Figure 5.12 Equivalent variation of income

The equivalent variation is the change in income that leaves the consumer just as well as off as some specific change in the price of a good. The consumer is initially at point *A* on budget line P_1. The price of good 1 falls, the budget line shifts to P_2 and the consumer shifts her spending pattern to *B*, which is on a higher indifference curve. The equivalent variation in income is given by the size in the parallel shift in the original budget line that would have taken the consumer to the level of utility indicated by the higher indifference curve (which is achieved after the fall in the price of good 1.) The equivalent income variation would have generated consumption point *C* as the optimal choice, but the consumer is indifferent between points *B* and *C*.

Equivalent variation

The equivalent variation is the answer to the question: if we had given the consumer a sum of money instead of a lower price of one product, how much extra income would have made her feel just as well off? This is illustrated in Figure 5.12. It is calculated by shifting outwards the original budget line parallel to itself until it just touches the new indifference curve achieved after a price fall of one good.

Compensating variation

This works backwards rather than forwards. It is the same as the income effect shown in Figure 5.10, and is measured by the distance a–a_1 in that figure. It is the amount of income that has to be taken away from the consumer following a price fall of one good in order to return the consumer to the initial indifference curve, and thus leave her feeling just as well off as before.

The difference between these two effects is whether the income adjustment is made relative to the consumption bundle chosen at the original prices or that chosen at the new prices.

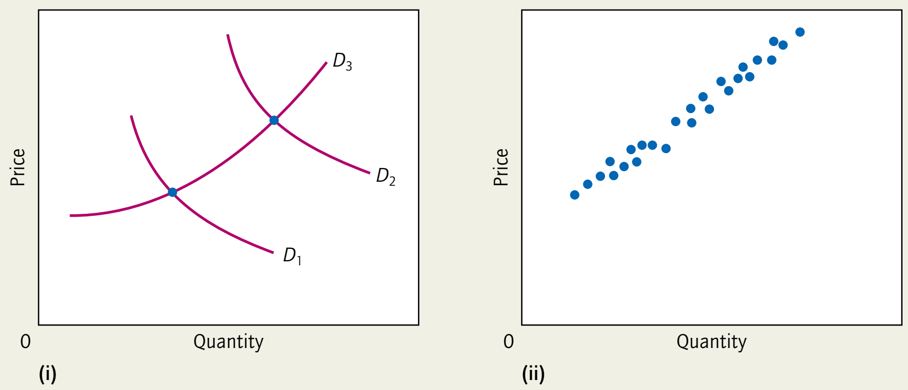

Figure 5.13 Which way does the demand curve slope?

Taste changes may explain isolated contradictions but not repeated ones. These observations may have been generated by changes in tastes that shifted a normal demand curve from D_1 to D_2 along an upward-sloping supply curve or by a supply curve that shifted along a positively sloped demand curve D_3. Part (ii) shows 26 weekly observations over a period when the product's price rose and fell (with incomes and other prices constant). The explanation that supply-curve shifts are operating on a positively sloped demand curve is more likely than the explanation that tastes changed each week to shift a normally sloped demand curve leftwards and rightwards along a upward-sloping supply curve.

CASE STUDIES

1. Income and substitution effects in practice

Although they sound highly abstract and 'theoretical' when first encountered, the income and substitution effects turn out to be useful tools. They help us to deal with many problems such as: Do high rates of income tax act as disincentives to work? Would cutting the rate of income tax increase the amount of work people will do? Would raising the wage rate of workers in some industry lead to a reduction in absenteeism?

Such questions frequently face decision-takers and they are often surprised at the results that the market produces. For example, many years ago the National Coal Board, which used to run the UK coal industry, raised miners' wages in an attempt to boost coal production and was surprised to find miners working fewer rather than more hours. In several countries increases in rates of income tax (within a moderate, not a confiscatory, range) have been found to be associated with people working more hours rather than fewer even though they earn less after-tax income for each hour worked; at other times reductions in tax rates seem to have caused people to work less even though they earn more after-tax income for each hour worked. The surprise in all these cases was the same. Intuition suggests that if you pay people more they will work more; experience shows that the result is sometimes the opposite: more pay, less work; less pay, more work.

The explanation of this surprising behaviour lies in distinguishing the income effect from the substitution effect of a change in the reward for work.

Think of Luke, starting with an endowment of 24 hours per day and deciding to consume some of it as 'leisure' (including sleeping time) and to trade the rest for income by working. If Luke works 9 hours a day at an after-tax rate of £10 per hour, he is consuming 15 hours a day of leisure and trading the other 9 for £90 worth of income which can be used to buy goods and services.

Now let the after-tax wage rate rise to £12 an hour, either because the wage rate rises or because the rate of personal income tax falls to produce that increase in after-tax earnings. Luke's response to this change will have an income and a substitution component.

The substitution effect works the way intuition suggested: more wages, more work. Gaining income is now cheaper in terms of the leisure Luke must sacrifice per £1 worth of income gained. At the new wage rate, 1/12 of an hour (i.e. 5 minutes) of work earns Luke £1 worth of income, whereas before it took 1/10 of an hour (6 minutes). Looked at the other way around, consuming leisure is now more expensive per amount of income that Luke must give up. An extra hour of leisure consumed requires sacrificing £12 of income instead of £10. The substitution effect leads to an increased consumption of the thing whose relative price has fallen—everything that income can buy, in this case—and to a reduced consumption of the thing whose relative price has risen—leisure.

So far so good. The surprise lies in the income effect. The rise in the after-tax wage rate has an income effect in the sense that Luke can have more goods *and* more leisure. He could, for example, consume an extra hour of leisure by cutting his hours worked from 9 to 8 while at the same time raising his income from £90 a day (9 hours @ £10) to £96 a day (8 hours @ £12). The income effect leads him to consume more goods and more leisure, that is to work fewer hours.

Only if the substitution effect is strong enough to overcome the income effect will the rise in the wage rate induce Luke to work more. If the substitution effect is strong enough, Luke might for example work 9.5 hours instead of 9 and increase his income from £90 to £114 a day. However, this means choosing this combination of income and leisure in preference to all combinations that give more income and more leisure, such as 8.5 hours of work (down from 9) and £105 of income (up from £90).

So we should not be surprised if increases in the after-tax hourly wage lead to less work; this merely means that the income effect is stronger than the substitution effect.

The above analysis helps to explain why employers separate higher overtime rates from normal rates of pay. If the normal rate of pay is increased the income effect is quite large, whereas if only the overtime rate is raised the income effect is much smaller but the substitution effect is unchanged. In the above example, raising the normal wage rate from £10 to £12 increases Luke's income by £18 if he continues to work an unchanged 9 hours a day. But introducing an overtime rate has an income effect only in so far as overtime hours are already being worked. If, in the previous example, the employer introduced a £15 hourly rate for work of over 9 hours a day, the income effect would be zero; Luke must work more in order to gain any benefit from the higher overtime rate.

2. Experimental economics and the concern for fairness

In the past decade or so economists and other scientists have co-operated in designing experiments to determine how people actually respond to various choices. This line of inquiry has provided evidence that individuals do not invariably maximize their own utility with no regard for what others around them are doing. The following extract summarizes the results of one such experiment.

Imagine that somebody offers you $100. All you have to do is agree with some other anonymous person on how to share the sum. The rules are strict. The two of you are in separate rooms and cannot exchange information. A coin toss decides which of you will propose how to share the money. Suppose that you are the proposer. You can make a single offer of how to split the sum, and the other person—the responder—can say yes or no. The responder also knows the rules and the total amount of money at stake. If her answer is yes, the deal goes ahead. If her answer is no, neither of you gets anything. In both cases, the game is over and will not be repeated. What will you do?

Instinctively, many people feel they should offer 50 per cent, because such a division is 'fair' and therefore likely to be accepted. More daring people, however, think they might get away with offering somewhat less that half of the sum.

Before making a decision, you should ask yourself what you would do if you were the responder. The only thing you can do as the responder is say yes or no to a given amount of money. If the offer were 10 per cent, would you take $10 and let someone walk away with $90, or would you rather have nothing at all? What if the offer were only 1 per cent? Isn't $1 better than no dollars? And remember, haggling is strictly forbidden. Just one offer by the proposer: the responder can take it or leave it.

So what will you offer?

According to utility maximizing theory you should keep the majority of the money for yourself and offer a very small amount to the responder. After all, her alternative is to get nothing. So if she is a maximizer, she will accept any offer greater than zero. But this is not what happens in such experiments!

Instead, between two thirds of the offers are between 40 and 50 per cent of the total sum. Only four in 100 people offer less than 20 per cent. Also more than half of all responders reject offers that are less than 20 per cent.

The motive of the person making the offer may be mixed. On the one hand, he may be concerned with what he thinks is fair. On the other hand, he may know that very small offers are likely to be refused because the responder will react strongly to what she perceives as an unfair offer.

But the motivation of the responder is not so complicated. She either accepts or rejects whatever offer she receives. According to maximization theory, there is a puzzle here: why should anyone reject an offer as 'too small'? The responder has just two choices: take what is offered or receive nothing at all. The only rational option for a maximizing individual is to accept any offer. Even $1 is better than nothing. A maximizing proposer who is also sure that the responder is also a maximizer will therefore make the smallest possible offer. In response, the responder will accept any offer greater than zero. The predictions of maximizing theory are clear on this one: offer as little as possible and accept anything positive, no matter how meagre.

The resolution of the puzzle in the case of both players is that they care about farness almost as much as they care about doing as well as they can for themselves.

The scenario just described, called the Ultimatum Game, belongs to a small but rapidly expanding field called experimental economics. . . . For a long time, theoretical economists postulated a being called Homo economics*—a rational individual relentlessly bent on maximizing a purely selfish reward. But the lesson from the Ultimatum Game and similar experiments is that real people are a cross-breed of* Homo economics *and* Homo emoticons*, a complicated hybrid species that can be ruled as much by emotion as by cold logic and selfishness. . . .*

Centuries ago philosophers such as David Hume and Jean-Jacques Rousseau emphasized the crucial role of 'human nature' in social interactions. Theoretical economists, in contrast, long preferred to study the selfish Homo economics*. They theorized about how an isolated individual—a Robinson Crusoe on some desert island—would choose among different bundles of commodities. We are, however, not Robinson Crusoes. Our ancestors' line has been social for hundreds of millions of years. And in social interactions, our preferences turn out to be far from selfish. (Sigmund, Fear, and Nowak 2002).*

The above extract does not imply that the assumption that individuals maximize their own utility or self-interest is useless. In everyday decisions, such as how many potatoes or holidays in Switzerland to buy, self-interest explains behaviour quite well. But the self-interest assumption is not applicable to many forms of group behaviour. We care about others as well as ourselves, and we also care about what others think of us. This often affects our behaviour, altering it from what a purely selfish individual would do.[12]

Conclusion

The demand curves for most products have negative slopes. Knowledge of the precise nature of the demand curve for a product is obviously important for firms that want to be able to predict the likely quantity demanded at various prices. An understanding of demand is also important for policy-makers, who might wish to impose taxes, intervene in markets in other ways, or predict the effects of sudden shortages of such things as food or energy. For economists, an understanding of demand is one important step along the road to understanding the detailed workings of a market economy.

12 For a fascinating discussion of the wider implications of altruistic behaviour, see Barber (2004).

SUMMARY

Early insights

- Consumers will maximize utility where the ratio of marginal utility to price is equal for all products.
- The paradox of value can be resolved when it is realized that marginal utilities and not total utilities determine market price.

Consumer optimization without measurable utility

- Indifference theory assumes only that individuals can order alternative consumption bundles, saying which bundles are preferred to which but not by how much.
- A single indifference curve shows combinations of products that give the consumer equal satisfaction, and among which he is therefore indifferent. An indifference map is a set of indifference curves.
- The basic assumption about tastes in indifference curve theory is that of a diminishing marginal rate of substitution: the less of one good and the more of another good the consumer has, the less willing she will be to give up some of the first good to get more of the second. This implies that indifference curves are negatively sloped and convex to the origin.
- While indifference curves describe the consumer's tastes, and therefore refer to what he or she *would like* to purchase, the budget line describes what the consumer *can* purchase.
- Each consumer achieves an equilibrium that maximizes his satisfaction at the point at which an indifference curve is tangent to his budget line.

The consumer's response to price and income changes

- The income–consumption line shows how quantity consumed changes as income changes with relative prices constant.
- The price–consumption line shows how quantity consumed changes as relative prices change. The consumer will normally consume more of the product whose relative price falls.
- The price–consumption line, relating the purchases of one particular product to all other products, contains the same information as an ordinary demand curve. The horizontal axis measures quantity, and the slope of the budget line measures price. Transferring this information to a diagram whose axes represent price and quantity leads to a conventional demand curve.

The consumer's demand curve

- A change in price of one product, all other prices and money income constant, changes both relative price and the real incomes of those who consume it. The effect of changes on consumption is measured by the substitution effect and the income effect.
- Demand curves for normal goods have negative slopes because both income and substitution effects work in the same direction, a decrease in price leading to increased consumption.
- A decrease in price of an inferior good leads to more consumption via the substitution effect and less consumption via the income effect. In the exceptional case of a Giffen good, the income effect more than offsets the substitution effect, causing the product's demand curve to have a positive slope.

TOPICS FOR REVIEW

- Marginal and total utility
- The paradox of value
- An indifference curve and an indifference map
- Slope of an indifference curve and diminishing marginal rate of substitution
- Budget line
- Absolute and relative prices, and the slope of the budget line
- Response of a consumer to changes in income and prices
- Derivation of the demand curve from indifference curves
- Income and substitution effects
- Hicks and Slutsky decomposition
- Normal goods, inferior goods, and Giffen goods
- Equivalent and compensating variations of income

QUESTIONS

1 Suppose a consumer's disposable income is £200 per week and she has a choice between spending this on meals and concerts. Concerts are £10 each and meals are £20 each. List the possible combinations of meals and concerts that could be bought with the income.

2 Using the same information as in question 1, the price of meals now falls to £10. What combinations of meals and concerts can now be purchased with the same income?

3 Assuming that (facing the prices in question 1) the consumer chose to consume 10 concerts and 5 meals per week, what change in income would leave the consumer still just able to consume this same combination of meals and concerts while facing the prices set in question 2? Would you expect this consumer to purchase the same combination of meals and concerts as before if faced by the new prices but with this amount less income?

4 Which of the following statements is true (there may be more than one or none)?

If the price of good X rises holding all other prices and income constant:

(*a*) the substitution effect alone will make a consumer buy more of X if X is inferior.

(*b*) the income effect alone will make a consumer buy more of X if it is a normal good.

(*c*) the income effect alone will make a consumer buy less of X if it an inferior good.

(*d*) the substitution effect will make a consumer buy less of X and it is irrelevant whether X is a normal or inferior good.

(*e*) the consumer will buy less of good X unless it is an inferior good, in which case she will always buy more.

5 Explain the difference between the income effect and substitution effect of a price change.

6 What is a Giffen good? Explain, using indifference curves, how it could arise.

7 Indifference curve analysis is not much use because it only tells us that demand curves slope down except when they don't. Discuss.

8 A company that normally pays its workers £400 per week in money decides to pay them instead with £400 worth of a specific good. Assume that there is no second-hand market in these goods, so that they cannot be sold for cash, but also assume that the workers would choose to consume some of these goods anyway. Using budget constraints and indifference curves, analyse whether the workers are likely to be just as happy with this arrangement as they were when they received their wages in money.