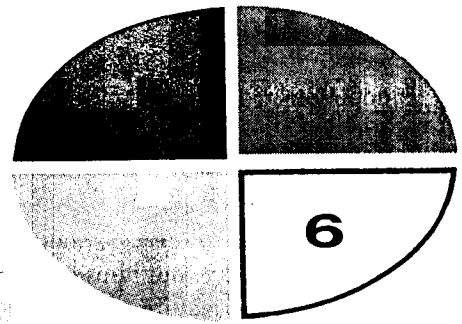

Household Economics and Demand

> Economy is going without something you do want in case you should, someday, want something you probably won't want.
>
> ANTHONY HOPE

THE HOUSEHOLD is a basic unit in the economy. Its spending decisions determine the demand for consumer goods, and indirectly, for raw materials and factors of production. Its decisions about saving influence the rate of capital formation. Its decisions about who in the family shall work, and how much work each shall do, determine the supply of labor.

These decisions are interrelated. Decisions about how much work to do, which means how much income to earn, are related to the satisfaction derived from spending this income on consumption. Saving decisions involve balancing the satisfaction from consuming right now against that of having higher consumption later on.

From an economic standpoint, the household is a satisfaction-producing organization, just as a business unit is a goods-producing organization. Its problem is to get greatest satisfaction with least effort or discomfort. More precisely, the problem is to maximize *net* satisfaction

that the two can be balanced against each other. This is what we assumed about Crusoe in Chapter 3. It is what we mean by treating the household as an economic *unit* rather than as a bundle of unrelated activities.

This chapter will explain briefly how a rational householder would go about solving the problem of maximizing net satisfaction. The discussion also has an important bearing on the issue of overall economic efficiency. A market economy is usually said to be "consumer-guided," that is, responsive to household preferences and decisions. For this to be an advantage, we must suppose that household decisions make sense, that people are actually able to do the best for themselves. What does it mean to do the best for yourself? We must pin down this notion more precisely.

DECISIONS ABOUT SPENDING

How would a rational consumer go about getting the most for his money? What *is* a rational consumer? In analyzing consumer behavior, economists usually assume that:

i. Consumers "know what they want." They have definite goals—freedom from hunger, shelter against the elements, security for the future, protection against domestic thieves and foreign foes, prestige in the community, rest and recreation. They also know how much they value one goal as against another.

ii. They know the effectiveness of various goods and services in attaining these goals—how much a new dress will contribute to prestige or how much a sirloin steak will reduce hunger.

iii. They know the price at which each good can be obtained in the market.

iv. Each consumer uses this information in such a way as to maximize his total satisfaction. This "does not imply (as almost all of its critics state) that the individual seeks to maximize money or wealth, that the human soul is a complex cash register. It does not affect the formal theory of demand in the least whether the individual maximizes wealth, religious piety, the annihilation of crooners or his waistline."[1]

These assumptions have been subjected to heavy criticism. It has been argued that people do not have definite goals, that they are interested mainly in developing new wants rather than in gratifying old ones; that consumers are poorly informed, led around by the nose by advertisers, and don't know what they are buying; and that we are creatures of custom and routine, buying usually on the basis of habit rather than rational choice.

The truth doubtless lies somewhere between the economist's picture

of rational man and the critics' picture of ignorant, habit-ridden man. Much purchasing is done on the basis of habit, and up to a point this is efficient. Conscious decisions take effort. It would be a nuisance to recalculate our whole life scheme every time we go to the store, to display what Professor J. M. Clark of Columbia has termed an "irrational passion for dispassionate rationality." It is also true, however, that habits respond to *marked* changes in the prices and availability of goods, provided one allows time for the word to get round. Consumers are better buyers of some things than of others. On items which are bought very frequently, people are bound in time to acquire some expertise. Ignorance is greater as regards unique purchases, where one cannot learn from repeated experience.

There is a further practical justification for our three assumptions: they enable us to say something definite about consumer behavior. Without them we could only say "people may do anything," which would not be much help. Moreover, hypotheses based on these assumptions bear a distinct resemblance to what we see people doing in the real world. They provide insight into, though not a full explanation of, actual consumer behavior.

Diminishing Marginal Utility

What shall we call the satisfaction which people get from consumption? Nineteenth-century economists labeled it *utility* and laid down an important principle about it. They reasoned that the more of a good I am already consuming, the less will an additional unit of it add to my satisfaction. They considered this a self-evident principle, which anyone can confirm by looking inside himself, and called it the law of *diminishing marginal* (= additional) *utility*.

Suppose Mr. A, a confirmed coffee drinker, could have only one pound of coffee a month. This would yield a large amount of utility. It would give him that first morning cup which makes it possible to face the day's work.

But perhaps instead of one pound a month he can have two. Now he can have a second cup of coffee at breakfast or perhaps a cup in midmorning. This will yield additional utility, though not as much as that yielded by the first pound. A third pound per month will mean that he can drink coffee at lunch or in the evening. This will add to his satisfaction, but by less than the second pound did. And so on. Eventually, when he has enough coffee to drink all he can hold, an additional pound will yield no utility at all.

This can be made more concrete by setting up a *utility schedule* for Mr. A (Table 1). Column 3 of the table shows how much additional satisfaction (marginal utility) Mr. A gets from each additional pound of coffee per month. The qualification *per month* is important. Consumption must always be shown as a *quantity per unit of time*. Otherwise we get

TABLE 1

Mr. A's Utility Schedule for Coffee

POUNDS PER MONTH (1)	TOTAL UTILITY (2)	MARGINAL UTILITY (3)
0	0	
		7
1	7	
		6
2	13	
		5
3	18	
		4
4	22	
		3
5	25	
		2
6	27	
		1
7	28	
		0
8	28	

into the absurdity of saying that 52 pounds of coffee (per year) can have as high a marginal utility as 1 pound of coffee (per week). Note that the entries in column 3 decline as we go down the column, in accordance with the principle of diminishing marginal utility. In this example, Mr. A would never wish to drink more than 7 pounds a month, since the eighth pound has a marginal utility of zero.

The second column shows Mr. A's *total utility* from consuming various amounts of coffee. Note that each entry in this column is obtained by adding up the entries in column 3 down to the level in question. Thus if Mr. A is consuming 4 pounds per month, his total utility is the sum of the marginal utilities of the first, second, third, and fourth pounds, i.e., $7 + 6 + 5 + 4 = 22$. As consumption increases, total utility rises as long as marginal utility remains positive. But when marginal utility falls to zero with the eighth pound, total utility ceases to rise.

Could marginal utility ever become negative? Yes, it could if one had so much of a product that part had to be thrown away and if this took trouble or expense. Tomatoes normally have a positive marginal utility; but a surplus of overripe tomatoes in my garden could be a nuisance.

The information in the table is shown graphically in Figure 1. The height of each rectangle shows the marginal utility yielded by an additional pound of coffee. The rectangles decline steadily in height,

Diminishing Marginal Utility from Consumption

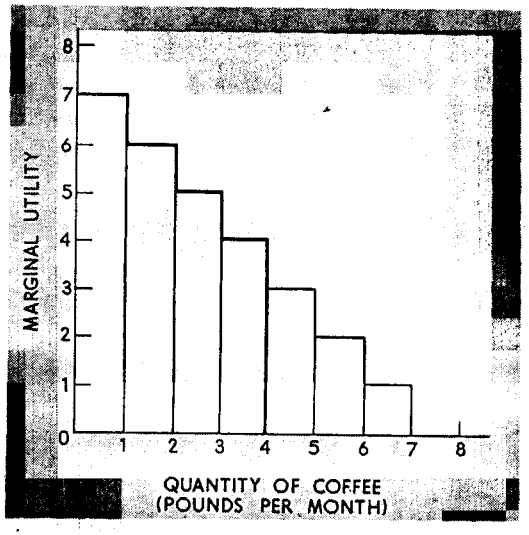

FIGURE 1. Each rectangle shows the amount of marginal (= additional) utility yielded by an additional pound of coffee per month. Note that each additional pound has lower marginal utility than the one before.

eventually reaching zero. The sum of the rectangular areas out to any point on the horizontal axis shows the total utility derived from consumption at that level. For example, if 4 pounds a month are being consumed, total utility is the area within the heavy black lines in the figure.

The staircase appearance of Figure 1 comes from the fact that we are dealing with whole pounds of coffee. Suppose we could buy coffee in very small amounts—ounces or half ounces—so that consumption could be varied by tiny steps. Then the rectangles in Figure 1 would become very thin. If they became thin enough, we should end up with a continuous line such as MU_c (marginal utility of coffee) in Figure 2. We shall normally use such lines from here on. They need not be straight lines, of course, but may show varying degrees of curvature. The only requirement is that they must slope downward from left to right.

The *height* of MU_c at any point shows the *marginal utility* of the last small unit of the product which has been added to consumption. The *area under* MU_c out to any point on the horizontal axis shows the *total utility* yielded by that rate of consumption. For example, at a consumption level of 4 pounds per month, total utility is the shaded area *OABC*.

Even with these simple ideas about utility, we can begin to draw some interesting conclusions. It is often said that the "economy of abundance" is either here or just around the corner. Professor Kenneth Galbraith of Harvard has argued that Americans now live in an "affluent society." What do such statements mean? Taken literally, "abundance" would mean that for every consumer and for all products, marginal utility has fallen to zero. This is clearly untrue. So being a bit more

Marginal and Total Utility

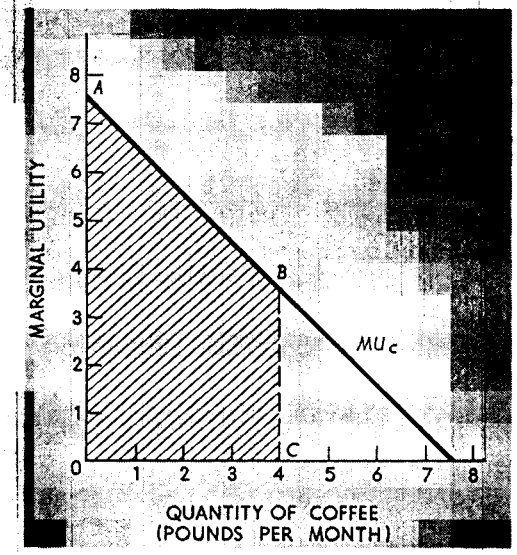

FIGURE 2. If consumption can be varied by very small amounts, the rectangles of Figure 1 become very thin and can conveniently be smoothed to a continuous line such as MU_c (marginal utility of coffee). The height of this line at any point shows the marginal utility of the last small unit of the product which has been added to consumption. The area under MU_c out to any point shows the total utility yielded by that rate of consumption. For example, at a consumption level of 4 pounds per month, marginal utility is *BC* and total utility is equal to the shaded area *OABC*.

charitable, what these writers must mean is that, for most people and products, marginal utility has reached a *low* level—low enough that total utility is high and people's wants are well satisfied. Is this actually true? Let each reader give his own answer.

Consider next the famous diamonds-water paradox, which puzzled the classical economists. Water is obviously more useful and necessary to life than diamonds. Yet diamonds command a much higher price. How can this be? Ingenious and even tortured reasoning was developed to explain the paradox. Today, we can give a simple explanation based on the distinction between marginal utility and total utility.

The total utility yielded by water is very great—indeed, the utility of the water needed to sustain life might be regarded as infinite. Since water is relatively abundant, however, the marginal utility of the last gallons used is low. Diamonds, on the other hand, are rare relative to the desire for them, so their marginal utility is high. The satisfaction yielded by the last unit of a good, and hence the price people can be induced to pay for it, is measured by marginal rather than total utility. This is why diamonds cost more than water.

Allocating the Consumer's Budget

It is all very well to set up a marginal utility schedule for one good. But the typical household has to distribute its spending over hundreds of different goods and services. How can it do this in a way which will yield greatest satisfaction? This is not an academic question. Every household faces it every day and answers it more or less effectively.

Once we begin talking in terms of money, we need a different version of the marginal utility curve. From now on, the horizontal axis of our diagrams will show, not *physical units* of a good, but *amounts of money* spent on the good. For clarity, it will be best to call this a *marginal utility of expenditure* (*MUE*) curve. In order to draw it we need to know both my liking for coffee, as expressed in the *MU* schedule, and the price of coffee. For example, suppose that my third pound of coffee per month yields 10 "utils" of satisfaction. The price of coffee is 50 cents per pound. Then the *MUE* at this point is $10/50 = 0.2$ utils per penny. By a similar calculation we might find that the *MUE* of money spent on a sixth pound is only 0.1 utils per penny. If we designate the price of coffee as *P*, then

$$MUE = MU \cdot \frac{1}{P}.$$

If *P* is constant, then *MUE* must fall continuously from left to right, because *MU* falls steadily as larger quantities are purchased.

Look now at a simple case in which a family buys only two goods, product A and product B, distributing its total budget of $5,000 a year between them. The "money" *MUE* curves for these goods are shown in Figure 3. Lay a string horizontally across the diagram, which we may call *II* (for income). Move the string up or down until you get a combination of expenditures, *OA* + *OB*, which just uses up the family's income. Note that *OA* + *OB* = $3,000 + $2,000 = $5,000. (This can be done by experiment. If expenditure is below income on the first try, so that there is

How to Get the Most for Your Money

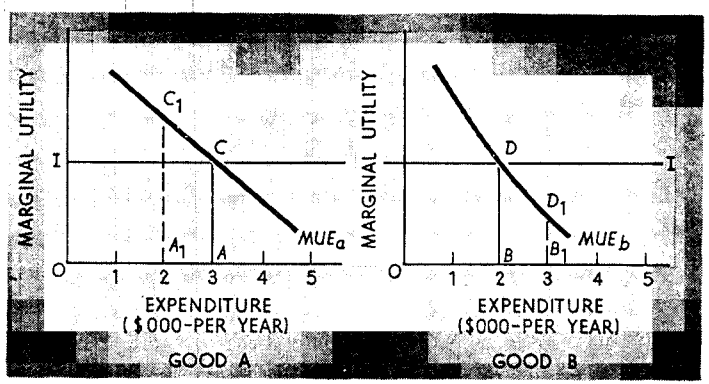

FIGURE 3. The problem is how to distribute my annual budget of $5,000 a year between goods A and B. Rule: the marginal utility of expenditure on A must equal the marginal utility of expenditure on B. So I should spend AO = $3,000 on A, and OB = $2,000 on B, because AC = BD. Why would OA_1 and OB_1 not be a correct combination of purchases?

money left over, lower *II* so that the family buys more of both goods; conversely if expenditure is too high. After all, this is what an actual family would do!)

Why is the combination *OA* + *OB* the best that the family can do? First, it is using its total income, so it cannot get more of one good without giving up some of the other. Second, *the last dollar spent on product A yields precisely the same utility as the last dollar spent on product B.* When the family is spending $3,000 on good A, the utility of the last dollar is *AC*. Similarly, the utility of the last dollar spent on B is *BD*. But *AC* and *BD* are equal, because *II* was drawn horizontally across both diagrams.

Suppose, instead, that the family had chosen to spend only OA_1 = $2,000 on A, and OB_1 = $3,000 on B. Then the last dollar spent on A would have a utility of A_1C_1, while the last dollar spent on B yields utility of B_1D_1. A_1C_1 is obviously larger. It follows that if the family shifted a little money from B to A, it would be gaining more than it would lose. Total satisfaction would be increased. If it shifted, say, $100 in the first instance, it would move a short distance down MUE_a and a short distance up MUE_b. But the marginal utility of expenditure on A would still be higher, so it would pay to shift another $100; and so on, until one reached the points *C* and *D*, where marginal utility is equal and no further gain is possible.

Now imagine hundreds of additional *MUE* curves for other products, extending out to the right of Figure 3, with *II* running across all of them. Then if the family spends on each good the amount indicated by the intersection of *II* with the *MUE* curve, it will be getting the same utility from the last dollar spent on each product. Since it cannot gain by shifting money from one good to another, this spending pattern yields maximum satisfaction.

The rule for rational budget allocation, then, is that *the marginal utility of expenditure on each good purchased must be equal.* If there are *n* commodities in the consumer's budget, the condition for maximum satisfation is that:

$$MUE_1 = MUE_2 = MUE_3 = \cdots = MUE_n.$$

We saw earlier that, for any good,

$$MUE = MU \cdot \frac{1}{P} = \frac{MU}{P}.$$

So the rule for maximum satisfaction can be written also as:

$$\frac{MU_1}{P_1} = \frac{MU_2}{P_2} = \frac{MU_3}{P_3} = \cdots = \frac{MU_n}{P_n}.$$

The marginal utility yielded by a unit of each good must be proportionate to its price.

We noted in Chapter 5 that a family's purchases of a good will be influenced by its tastes, by the price of the good, and by the prices of related goods. It is now clear that purchases also depend heavily on income. The effect of income is illustrated in Figure 4. For simplicity

My Spending Pattern Will Vary with My Income

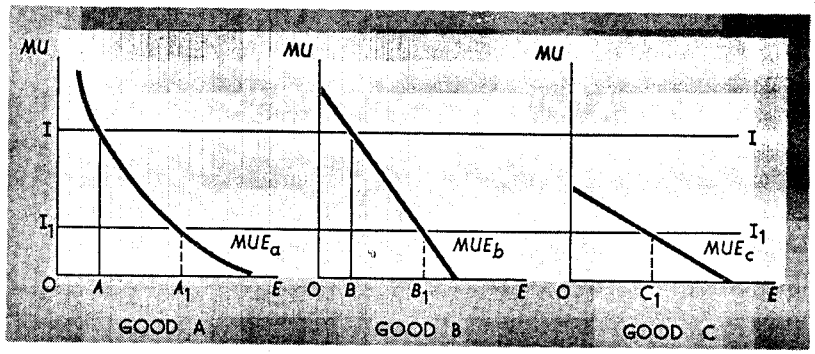

FIGURE 4. If I have a small income, corresponding to the income line *II*, I can buy only OA + OB and cannot afford good C at all. With a larger income corresponding to I_1I_1, I can buy $OA_1 + OB_1 + OC_1$.

show only 3 goods instead of 300. We ignore saving for the time being, and suppose that all income is spent on consumption. If the family's income is small, its consumption pattern might be defined by *II*. It buys small amounts of A and B, but "cannot afford" any of C, i.e., the MUE_c schedule falls entirely below the income line. (When you say you cannot afford something, you mean really that its marginal utility per dollar to you is not as high as that of the things you do buy. You can always afford the things that matter *enough*.) As usual, $OA + OB$ equals the family's income.

Suppose now the family's income were much larger. Then its consumption pattern might be defined by I_1I_1. It buys a good deal more of A and B, and also buys some C which it did not buy before.

While the effect of an increase in income is to add new goods to the family's budget and to increase its purchases of most things it was already consuming, this may not be true of *every* good. Family budget studies show that a higher income usually means *smaller* purchases of such things as margarine, potatoes, beans, and cheap lines of mass-produced clothing. These things are called *inferior goods*—not a term of reproach but simply an indication that use of these goods is negatively related to income.

DECISIONS ABOUT SAVING

So far we have assumed that this year's income is all spent on this

income for use in later years. What considerations influence this kind of decision?

Only a few misers put their savings under the mattress. Income saved is usually put into forms which are expected to return a larger income in the future. If you put money in a savings account, you receive interest each year and can get back the amount you put in whenever you choose. If you buy an 8-year government savings bond for $18.75, you will get back, say, $25 when the bond matures. If you buy a common stock, you will usually get cash dividends and can turn the stock back into cash whenever you wish.

In all these cases the saver is really *buying future income*. He is trading $1 of spending today for the possibility of spending something more than $1 sometime in the future. This future income can be regarded as a good with its own utility schedule. We can chart it and add it to all the other *MUE* curves which the consumer uses in reaching his budget decisions.

Look at the household illustrated in Figure 5. The left-hand chart

Balancing Consumption and Saving

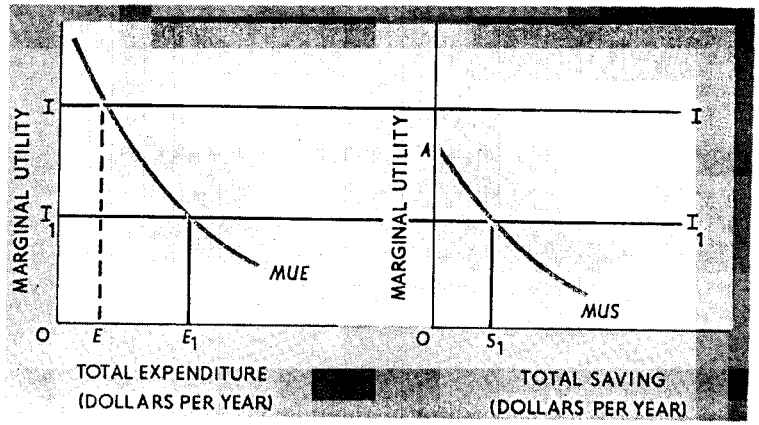

FIGURE 5. On the left-hand side, *MUE* shows the marginal utility of dollars spent on current consumption. On the right-hand side, *MUS* shows the marginal utility of dollars saved for use in the future. The usual rule for rational choice applies: the consumer should divide his budget so that *MUE* and *MUS* are equal. If his income is very low, corresponding to *II*, he "cannot afford" to save anything. But with a higher income, indicated by I_1I_1, he can spend more, OE_1, and save OS_1 as well.

shows its marginal utility curve for spending on current consumption. Where does this curve come from? It is based on the utility schedules for individual goods, which were illustrated in Figure 4. We saw there that if the consumer is allocating his budget correctly, the marginal utility of the

relation is shown by *MUE* in Figure 5. The curve must fall from left to right in the usual way: the greater the amount spent, the lower the marginal utility yielded by the last dollar.

The *MUS* (marginal utility of saving) schedule is shown on the right side of the chart. The dollars along the horizontal axis are *current* dollars—dollars saved today out of today's income. On the vertical axis, the distance *OA* shows how much utility the individual gets *today* from the first $1 of income saved *today;* and so on for lower points on *MUS*. Why does *MUS* slope down to the right? Because it represents expected satisfaction from *future consumption.*[1] Since consumption in any time period yields diminishing marginal utility, this should be true of future consumption as well as present consumption. Saving is also done partly for security reasons; but here the same rule applies. The more security I already have, the less an additional bit of security adds to my satisfaction.

Given the *MUE* and *MUS* schedules, we can apply the standard rule for rational choice: the consumer should save an amount such that his utility from the last dollar spent equals that from the last dollar saved. At the low income indicated by *II* in Figure 5, his income line lies entirely above the *MUS* curve. He "cannot afford" any saving, and will spend an amount *OE* equal to his total income. With the larger income corresponding to I_1I_1, he will maximize satisfaction by spending OE_1 and saving OS_1. As usual, OE_1 and OS_1 must add up to his total income. If his income were still larger, he would spend more *and* save more than before. A rise in family income normally raises both spending and saving.

The *MUS* schedules of different families must differ considerably, because at any income level—say, $10,000 a year—one finds some families that save a good deal and others that save nothing at all. Why is this so? What determines the height and shape of a family's *MUS* schedule?

One important consideration is how much property the family has already accumulated. Other things equal, a man who has already accumulated $100,000 will not be as eager to save as a man who is starting from zero.

A second consideration is whether this year's income is expected to continue in future. If a family has been used to getting $10,000 a year, its scale of living will be adjusted to that income level. This year, through an unexpected windfall, it has an income of $12,500, but this is not expected to continue. It would not be sensible, then, for the family to raise its consumption level substantially. Much of the extra $2,500 will probably go into saving.

[1] True, a wealthy man often builds up an estate which is expected to outlast his own lifetime. So *he* will not be spending the future income. Presumably, the prospect of his children or grandchildren being able to spend their legacies at some

There are important differences in families' future needs relative to expected future income. A young man who expects that his salary will be rising for many years will react differently from an older man looking forward to a low retirement income. A family faced with college expenses for children will behave differently from a family without such expenses. Research studies have shown that there are systematic differences in savings behavior at different stages of a family's life cycle.

It is sometimes argued that the rate of interest, which determines how much today's dollar will grow by next year, has an important effect on saving. If a dollar is expected to grow to $1.10 by next year, saving it will yield more utility than if it is expected to grow only to $1.05. So a rise in the interest rate will raise the *MUS* schedule and cause a higher proportion of household income to be saved; and conversely for a drop in interest. Many economists, however, doubt that this is very important compared with other influences on saving decisions. We shall go further into this in Chapter 11.

Apart from objective circumstances, people differ in their foresight or anticipation of the future. It is reasonable that a dollar of income next year should not be valued as highly as a dollar right now—next year I may not be here to enjoy it. Moreover, expected future satisfactions do not register as sharply on the mind as current satisfactions. So a future dollar shrinks in importance, and the farther away it is the more it shrinks, like railroad tracks disappearing in the distance. But some people discount the future at a much higher rate than others. People who *have a high time preference* for present over future income may appear shortsighted or spendthrift. Similarly, a man who values next year almost as much as this year *has low time preference,* and may appear thrifty, or farsighted. No praise or reproach is intended. These are preferences like any others, among which the economist must be neutral.

Given a family's *MUS* and *MUE* schedules, saving depends upon income. Taking the community as a whole, income is of dominant importance. Statistical studies show that, as we go up the income scale, the average amount saved per family rises steadily. Saving also rises as a percentage of family income.

DECISIONS ABOUT WORK

So far we have not asked where income comes from. Most of it comes from working; so decisions about work are a basic part of household economics. There is no law which says how many hours a man must work. He can get by with 30 hours a week in some occupations, or he may put in 40 hours in another, or he may get 50 hours by working overtime or holding a second job—"moonlighting." Several million people in the United States do hold two jobs at present.

How many hours a week does it pay to work? Each extra hour

utility, though at a diminishing rate. But time off for rest and recreation also yields utility. It takes time to spend and enjoy one's income. We must assume that these different kinds of satisfaction are commensurable and can be weighed against each other. People obviously do weigh them in reaching decisions about how much to work.

On this basis we can construct Figure 6. Along the horizontal axis

How Much Work to Do: Balancing Income and Leisure

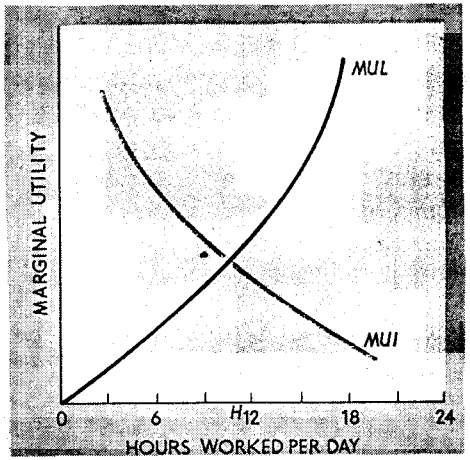

FIGURE 6. *MUI* shows the declining marginal utility of income as this man works more hours per day. *MUL* shows the rising marginal utility of leisure as work hours are lengthened, *i.e.*, leisure hours are decreased. The best workday for this man would be *OH*, where *MUI* and *MUL* are equal. Why would he get less satisfaction by working fewer or more hours than this?

we show the 24 hours a day which are available to be divided between work and leisure. Each point on the axis shows a possible combination of the two. As we move from left to right, work increases and leisure decreases. *MUI* shows the marginal utility of the income derived from additional work. It slopes downward to the right in the usual way. Since an hour of leisure is an hour of *not working*, we can show the marginal utility of leisure (*MUL*) on the same diagram. As we move to the right on the diagram, the amount of leisure is decreasing, and so its marginal utility is rising. Hence *MUL* slopes upward from left to right.

The reasoning behind this shape of *MUL* is simple. If work hours were zero, the worker would have so much leisure that the last hour would yield him little satisfaction. Leisure might even have negative utility, if he is ashamed of sitting at home idle. But as working hours are lengthened, each hour chopped off his leisure is more precious than the last. So as we move to the right on the diagram in the direction of more work and less leisure, the marginal utility of leisure rises steadily.

The man in Figure 6 would do best for himself by working *OH* hours per day. Why? Because if he worked less than this, the marginal utility of income would be higher than that of leisure. He could gain by trading a little more leisure for income, thus moving to the right toward *OH*. Any hours beyond *OH*, however, would involve a loss of satisfaction, since in this range *MUI* is below *MUL*.

People differ widely in their attitude toward leisure; and so some choose to work a great deal and others very little. If we want to be polite, we can say that Joe Doaks has a high preference for leisure, a high *MUL* curve. If we want to be rude, we can say he is lazy. Some people seem to be naturally work-shy, while others work with compulsive zeal. A person with a wide range of athletic, cultural, or other interests, a person who is skillful in getting satisfaction from leisure, will have a higher *MUL* schedule than a person of limited interests. (One of the objects of a liberal education is supposedly to raise your *MUL* curve!) The kind of work a person is doing also has a bearing. If the work is exhausting, monotonous, or otherwise unpleasant, leisure will seem more attractive than if the work is intrinsically pleasant and interesting.

This last point amounts to saying that work usually involves *disutility* or negative satisfaction. The extent of this disutility, however, varies widely from job to job. The more unpleasant the job, the more the balance of choice will be tilted in favor of leisure and the shorter will be the preferred work week.

A man may choose to work long hours, not because his valuation of leisure is low, but because his desire for income is great—he has a high *MUI* curve. This depends partly on the size of his family responsibilities. In boom periods, many factories work overtime, and their workers have a chance to bid for overtime hours at premium pay. Research studies show that married men with families put in substantially more overtime than young, single workers whose income needs are lower. Even with the same family size, however, some men will aspire to a high standard of living and work hard to attain it, while others will choose to work less and live less well.

People sometimes have income which is independent of work. Suppose the man in Figure 6 has a great-aunt, who leaves him some securities which bring in $2,000 a year. How will this affect the amount of work he does? If he earns $3,000 a year from work, his total income is now $5,000. The last dollar he gets from work is now the 5,000th dollar in his income instead of the 3,000th dollar; so it must have lower marginal utility than before. The marginal utility of income *from working* is reduced, and, everything else remaining the same, he will prefer to work less.

In all this we have taken the worker's *hourly wage rate* as given. To determine any point on *MUI*, say the point directly above 8 hours of work, we perform two operations. First, we multiply this number of hours by the wage rate to get the number of dollars earned per day. Then we must determine the marginal utility of the dollars earned in the last hour. If the wage rate changes, this changes our calculations. And we cannot be sure whether a wage increase, say, will lead a man to work more hours or fewer hours. The fact that an extra hour of work now yields more income is by itself an inducement to work the extra hour. But

the fact that his labors earn him greater income for *all* hours, not just the extra hour, means that his income is higher, and so the marginal utility of extra income is less than before. These two effects pull in opposite directions, and on balance the man may work either more hours or less.

Some economists have speculated that, on the average for the entire labor force, the income effect is predominant; and that higher wage rates will cause a reduction in average hours per worker. Over the past century, there has been both a great increase in real wage rates and a marked reduction in weekly and yearly hours of work. If the reduction of hours represents a deliberate choice based on workers' preferences, this seems to support the hypothesis of an inverse relation between wage levels and labor supply.

This question is of practical importance also in the realm of tax policy. A 20 percent income tax imposed on a man whose income comes entirely from work amounts to cutting his wage rate by 20 percent. Will the fact that an hour's work now yields less income cause him to work less than before? Or, on the contrary, will he work longer hours to restore his take-home pay to its previous level and maintain his standard of living? Either result is quite possible. Research studies on actual reactions to income taxation have so far been inconclusive.

THE EQUILIBRIUM OF THE HOUSEHOLD

To sum up: a supposed virtue of a market economy is that it is responsive to individual choices—choices about spending, about saving, about working. Each household has an opportunity to work out a pattern of life which, given its particular goals and preferences, will yield maximum satisfaction.

This chapter has attempted to define "maximum satisfaction." It appears that maximizing satisfaction involves the household in a set of interlocking calculations. The amount of work to be done (and this may include work by the wife or other household members as well as the husband) must be chosen so that the marginal utility of the income gained just equals the marginal utility of the leisure sacrificed. Saving must be carried to the point where the utility of the last dollar saved just equals the utility of the last dollar spent. Spending must be spread over the goods and services available so that the last dollar spent on each yields the same utility. If all these choices are being made correctly, the household is *in equilibrium*, that is, it cannot increase its satisfaction by a move in any direction.

This definition of overall equilibrium would be interesting even if no one behaved in precisely this way. But the question how far households actually do behave in this way is, of course, important. This divides into two subquestions: first, how far do existing economic institutions allow people to make the kinds of marginal calculation

outlined above? Could households exercise rational choice even if they were fully capable of it? Second, how rationally do households use the opportunities for choice which are open to them? A word on these points will provide a fitting conclusion for the chapter.

As to opportunity, households are perhaps best off in the distribution of *current purchases.* Rarely does any one dictate what you shall buy. Almost every conceivable good and service is available in some nearby retail outlet. The main problem is limited information. The quality and performance of some complex products is hard for the consumer to appraise; and a good deal of misinformation is spread deliberately by producers. Ignorance and misinformation is concentrated in certain kinds of product, however, and these do not form the bulk of consumer purchases.

As to *saving,* there is now a large volume of compulsory saving over which individuals have little control. The federal Old-Age, Survivors, and Disability Insurance system (OASDI) obliges most workers and employers to set aside certain amounts of money while a man is employed to provide him with a retirement pension, or a disability allowance should he become disabled before retirement, or death benefits for his family in the event of premature death. There are also many private pension systems, set up by companies or by union-employer negotiation, which are binding on the individual employee.

Establishment of compulsory savings systems does not affect our earlier argument so long as *the amount of compulsory saving is less than the person would have saved voluntarily.* His voluntary saving will now be lower than it would have been, but total saving may be the same. But OASDI and other compulsory savings systems have doubtless forced many low-income families to save *more* than they would have done voluntarily. In defense of this interference with free choice, it can be argued that most families are quite shortsighted, as witness the high proportion of old people who ended up living on their children or on relief in pre-social security days. Compulsory saving for old age or disability is probably a desirable corrective for this shortsightedness; but it does involve interference with individual choice.

The choice of *how much work* to do is somewhat constrained by law and custom, but perhaps not as seriously as appears at first glance. The federal Fair Labor Standards Act and union contracts have established a basic 40-hour week in many industries, and that is that. But why does the standard week hover around 40 hours? This must be within a reasonable distance of the hours which the average worker would choose voluntarily. If not, there would probably be agitation in the unions and in Congress to change the standard. Moreover, there are still local or intrastate industries which work 48 hours or more a week; and some industries work as little as 30 hours. There are opportunities to volunteer for overtime work or to hold a second job. There are a growing number

of part-time jobs. Several million workers, mainly women, voluntarily work less than a full week. Businessmen, professional men, and farmers have wide scope to vary their hours. In farm areas located near cities, many farmers hold full-time jobs in town and run a farm on the side. Long hours, but high income.

It seems, then, that there is wide scope for economic choice for households. But how fully do households exploit the opportunities which are open to them? How closely do they conform to the canons of rational choice? We do not have nearly as much evidence as one would like. If full evidence were available, the picture would doubtless come out quite mixed. Certainly not *every* household makes *precise* economic calculations *every day* of the year. Household heads differ in intelligence and temperament, and the accuracy of their economic calculations will differ accordingly. Calculations are often rough, and leave some difference in marginal utilities instead of reaching precise equality. People may not respond to small changes in income, prices, or wages, but may wait until the changes have become large enough to be worth thinking about. In between these spasms of decision, they follow habitual routines, which may be a sensible way of conserving mental effort.

With all these qualifications, there seems to be a *strong element* of rational choice in household behavior. Within the limits of their imperfect information, people do try to do the best for themselves. They do find their way gradually toward cheaper and better goods, toward better jobs, toward preferable hours of work, and so on.

SUMMARY

1. The marginal utility of a good is the satisfaction derived from the last unit consumed (per period of time). Marginal utility decreases as the number of units consumed increases. This is called the principle of *diminishing marginal utility.*

2. The individual's demand schedule for a good is derived from the marginal utility schedule. It slopes downward to the right because of the principle of diminishing marginal utility.

3. The marginal utility of *expenditure* on a good depends on the marginal utility of the good *and* on its price. $MUE = \frac{MU}{P}$.

4. A household maximizes its satisfaction from consumption by allocating its budget so that the last dollar spent on each good yields the same marginal utility. The rule can be stated as:

$$MUE_1 = MUE_2 = MUE_3 = \cdots = MUE_n$$

or

$$\frac{MU_1}{P_1} = \frac{MU_2}{P_2} = \frac{MU_3}{P_3} = \cdots = \frac{MU_n}{P_n}$$

5. Decisions about saving can be analyzed in the same way. The rational householder will divide his annual budget so that his utility from the last dollar spent on current consumption equals the utility he expects to derive from the last dollar saved.

6. Families with the same income this year may save different amounts, depending on how much property they have already accumulated, their prospective expenditure needs in the future, whether this year's income is expected to continue, and their valuation of the future compared with the present. On the average, however, saving rises as family income rises, both in absolute terms and as a *percentage of income.*

7. A worker will select his preferred hours of work by balancing the marginal utility of income against the marginal utility of additional hours of leisure. Individual preferences on this point differ substantially, both because of differences in the valuation of leisure and differences in desire for income.

8. An increase in a man's hourly wage rate may lead him to work either more hours or fewer hours. In the United States, taking the labor force as a whole over the past hundred years, rising wage levels have been associated with a marked reduction of hours per worker.

9. If a household reaches correct decisions on allocation of consumption expenditures, rate of saving, and amount of work to be offered in the market, it will have attained a position of maximum satisfaction. If all households in the economy do this, total consumer (and worker) satisfaction will be at a maximum.

DISCUSSION QUESTIONS

1. "Economists assume that consumer wants emerge spontaneously, and that producers respond by providing what consumers desire. It would be more accurate to say that producers create wants by developing new products, which consumers are then persuaded to purchase. The producer-dominated consumer is a reality, while the consumer-guided producer is largely a fiction." Discuss.
2. Review the main kinds of good you purchase. Can you think of any which do *not* obey the principle of diminishing marginal utility?
3. Does the "equal *MUE*" rule provide a sensible basis for consumer budget decisions? How close do you come to observing this rule in your own purchases?
4. Is saving always a Good Thing? Under what circumstances might zero saving or negative saving (borrowing) be quite rational?
5. "While the interest rate plays a large role in economic theory, it actually has a minor effect on savings decisions compared with other circumstances." Do you agree? What are some of the "other circumstances"?
6. High-income households save much more, both absolutely and as a percentage of income, than do low-income households. So it would seem that

the rich should get steadily richer. Actually, however, the proportion of all property owned by the richest 5 percent of the population has not increased in recent decades. Can you think of any explanations for this apparent paradox?

7. "The idea that workers can make a free choice between income and leisure is far from reality. When you take a job, you take the hours that go with it, and that is that." Discuss.
8. It is sometimes argued that union contracts and labor legislation have reduced working hours so much that most workers are getting less income and more leisure than they really prefer. What is your own hypothesis on this point? What kinds of facts might be gathered to test your hypothesis?

APPENDIX: THE INDIFFERENCE-CURVE APPROACH TO HOUSEHOLD DECISIONS

Many economists view the utility approach used in this chapter with considerable suspicion. It is not hard to see why. After all, what *are* the units on the vertical axis of our utility diagrams? Imaginative textbook writers sometimes label them "utils." Mr. A's first pound of coffee yields him 7 utils of satisfaction, the next pound 6 units, and so on. Nineteenth-century economists were inclined to regard a unit of satisfaction as a measurable quantity, like a pound or a foot. You could add and subtract them, show them as areas on a diagram, and so on.

But many modern economists regard this as dubious psychology. You can *feel* satisfaction, but you cannot *measure* it. You can also *compare* satisfactions. A man can say that one unit of a good, or one bundle of goods, yields more satisfaction than another; but he cannot say precisely how much more. All he can do is *rank* satisfactions as greater or less.

Moreover, it is argued, all the conclusions we have reached about household behavior can be reached equally well by another route, a route which carries no implication that utility is measurable. The key concept in this alternative approach is the *indifference curve*. This is a curve linking positions of *equal satisfaction*, among which the individual will consequently be indifferent. This may look complicated at first glance. But it is a quite simple and useful tool. We shall apply it first to decisions about current spending on consumer goods, and then see how it can be applied also to decisions about saving and working.

The Choice of Consumption Goods

Consider the case of Mr. Bibber (Figure 7). He has a fixed income each month to spend on food, and he eats only two foods, meat and potatoes. He is entitled to choose various combinations of these; but since his income is limited, he can get more of one item only by giving up some

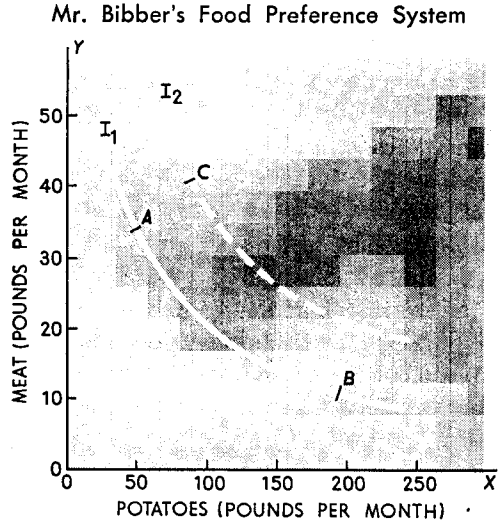

FIGURE 7. I_1 is an indifference curve. It shows combinations of meat and potatoes which would be equally acceptable to Mr. Bibber. So does I_2. The higher curve shows a higher level of satisfaction, since on it he can have more of both goods.

different combinations of the two articles which would be equally acceptable to him. He would be equally content with much meat and very little potatoes (point A), or far more potatoes and a small meat allowance (point B). As we move down the curve from A, its shape tells us how much additional potatoes it would take to compensate Mr. Bibber for the loss of a certain amount of meat. It shows the terms on which he is willing to substitute one product for the other. As we move closer to B and the amount of meat diminishes, it will take more and more potatoes to get Mr. Bibber to give up a pound of meat. This seems like a reasonable supposition. Finally, we may reach a rock-bottom level at which Mr. Bibber is unwilling to give up any more meat no matter how many potatoes he is offered. At this point I_1 becomes a horizontal line.

This construction, an indifference curve, is a very flexible and useful device. Different attitudes toward the two commodities can be shown simply by differences in the shape of the curve. If the consumer regards the two products as close substitutes, almost interchangeable with each other, I_1 will be a shallow curve, close to a straight line. If he does not regard them as very good substitutes for each other, the curve will bend down more toward the origin. In the extreme case, where Mr. Bibber insists on a specific "package" of so many pounds of meat and so many pounds of potatoes and is unwilling to substitute at all, the indifference curve would become a right angle with the point of the angle showing the one preferred combination.

"But," you may ask, "what about a point like C? Here he has *both* more meat and more potatoes than at A. So why won't he simply move up to C and stay there?" Of course, Mr. Bibber would prefer C to A if he

higher indifference curve representing a higher level of living. This second curve is shown by I_2. There are in fact a large number of indifference curves for Mr. Bibber, lying one above the other, and filling the whole space in the diagram. Each corresponds to a different satisfaction level, and Mr. Bibber naturally will climb to the highest one that his actual income level permits. The totality of his indifference curves for all levels of satisfaction is termed his *preference system.*

The notion of a preference system seems abstract and unreal on first acquaintance. One reason is that our illustration involves choice between two commodities only, whereas real-world choices are obviously more complex. But there is no logical difficulty in extending the idea of a preference system to embrace 10, 50, or any number of commodities which may be available for consumer choice.

So far we have not asked how much of any product—say meat—Mr. Bibber will actually buy. This depends, not just on Mr. Bibber's preferences, but also on his income and on product prices. We can demonstrate this by setting up the indifference diagram in a slightly different form (Figure 8). On the vertical axis we show pounds of meat bought per month. But on the horizontal axis we now show the dollar amount he spends per month. If prices are constant, this money income represents his command over *goods in general.* So the indifference curves, such as I_1 and I_2, show the rate at which he is willing to trade a reduction in meat for an increase in everything else he buys.

Suppose Mr. Bibber's monthly income is $OP = \$40$; and suppose the price of meat is \$1 per pound. Then if he spent all his money on meat, he could buy $OM = 40$ pounds of meat per month. Alternatively, if he bought no meat, he would have $OP = \$40$ per month to spend on other things. Or he can choose some intermediate combination. All the

Purchases Depend on Price and Income

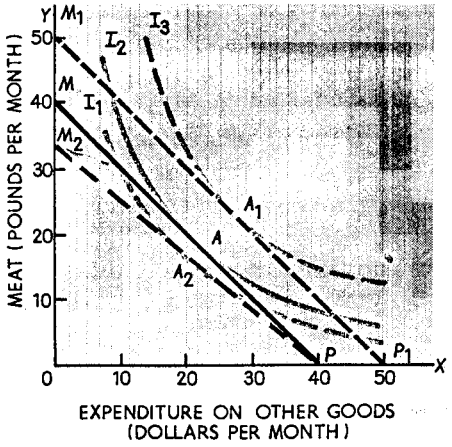

FIGURE 8. The budget line *MP* shows all combinations of meat and other things which Mr. Bibber can buy with his income. He will select the combination shown by *A*, where *MP* touches the highest indifference curve he can reach. If his income rises to P_1, his budget line becomes M_1P_1 and his purchases rise to A_1. If the price of meat rises, his budget line shifts to M_2P and his purchases fall to A_2.

combinations open to him lie along the straight line *MP*. (Why is *MP* a straight line? Because its slope depends on the price of meat, which we have taken as constant.)

We may call *MP* his *consumption possibilities line* or his *budget line,* since it shows all the combinations of meat and other goods which he can buy with his income. He cannot move to any position above *MP* because he hasn't enough income. He will not choose any position below *MP* because that would mean leaving some of his income unspent—"throwing money away."

How much meat will Mr. Bibber actually buy? He must settle for some point on *MP*. Given this limitation, he will naturally try to reach the highest possible indifference curve (or satisfaction level). The highest indifference curve which is touched by *MP*, therefore, will be the best he can do. This is I_2 in Figure 8, and his preferred combination is shown by *A*, the point at which *MP* is tangent to I_2. Here he is buying about 22 pounds of meat and has about \$18 per month left to spend on other things.

A little thought reveals why Mr. Bibber cannot profitably move in any direction from *A*. If he tries to move right or left along I_2, he cannot do so because all points on I_2 except *A* lie *above MP*, i.e., they represent combinations which cost more than his income. If he tries to move up or down along *MP*, he can do so but will find himself on a *lower indifference curve* representing a lower level of satisfaction.

The tangency of *MP* and I_2 has an important economic meaning. At the point of tangency, *A*, the slope of *MP* is just equal to the slope of I_2. But what is the slope of I_2? It shows the rate at which Mr. Bibber *is willing* to substitute meat for other goods in his consumption pattern. The slope of *MP*, on the other hand, shows the rate at which he *can* substitute in view of existing prices. He is in equilibrium when his *subjective* rate of substitution just equals the *objective* rate permitted by the market.[2]

How will Mr. Bibber's meat purchases be affected by changes in his income or in the price of meat? Suppose first that his income rises to $OP = \$50$ per month, while the price of meat remains unchanged. The maximum amount of meat he can buy is now $OM_1 = 50$ pounds a month. His new budget line is M_1P_1, parallel to *MP* at a higher level. (Why is it parallel? Because the slope of the budget line depends on the price of meat. Since this has remained unchanged, M_1P_1 will have the same slope as *MP*.) The highest indifference curve he can reach is now I_3, and he

[2] This is the same thing as the $\frac{MU_a}{P_a} = \frac{MU_b}{P_b}$ rule reached in the body of the chapter. This can be converted to the form $\frac{MU_a}{MU_b} = \frac{P_a}{P_b}$ The first term is the subjective rate of substitution, while the second is the objective or market rate.

will settle at A_1, where M_1P_1 is tangent to I_3. He buys more meat than before and also more of other things.[3]

Now let's keep Mr. Bibber's income constant at $40 per month and raise the price of meat to $1.25 a pound. The most meat he can buy is now $OM_2 = 32$ pounds a month. The budget line becomes M_2P. (Note that the slope of the budget line has changed because of the change in price.) The best position which Mr. Bibber can now reach is A_2, the point at which M_2P is tangent to the lower indifference curve I_1. The price increase without any increase in income has lowered his satisfaction level, which is what common sense would suggest.

Mr. Bibber now buys less meat than he did before. His meat purchases have fallen for two reasons: (1) the fact that meat prices have risen while his money income remains unchanged means that his purchasing power is reduced. He is less "well off" than before. Note that he also buys less of other goods than he did before, even though their prices are unchanged. This is termed the *income effect*. (2) Since meat is now *relatively* more expensive than before, he buys *relatively* less meat and more of other things. He substitutes other goods for meat in his consumption pattern. This is termed the *substitution* effect.

The income and substitution effects normally work in the same direction. The income effect can run in the opposite direction, however, and in rare cases might be big enough to outweigh the substitution effect. Then one encounters the "Giffen paradox" mentioned in Chapter 5.

To sum up: Mr. Bibber's meat purchases depend on (1) his income; (2) the price of meat; (3) the prices of all other goods he buys; and of course (4) his basic tastes and preferences, reflected in the shape of the indifference curves. This is fully consistent with our earlier results using the utility approach.

The Choice of Spending or Saving

We saw earlier in the chapter that a saver is really *buying future income*. Consider a man wondering whether to save part of this year's income in order to earn interest which will add to his income next year.

On the horizontal axis of Figure 9 we show this year's consumption. On the vertical axis we show interest earnings which would form additions to next year's income. The indifference curves show combinations of present expenditure and future interest income which would be equally satisfactory to Mr. Goldbilt. As he cuts down his present spending, it will take larger and larger interest payments to compensate him; so the indifference curves, as usual, become steeper as we move up to the left. If there is a minimum level of living below which he will not go, his indifference curves will become vertical at that level.

[3] While this is the general rule, it is not an invariable rule. The indifference

Saving Varies with Income and the Rate of Interest

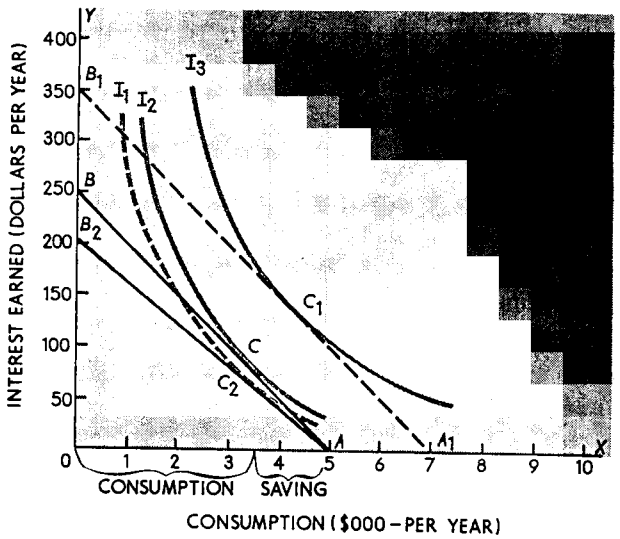

FIGURE 9. If the market rate of interest is 5 percent, Mr. Goldbilt can earn $50 interest for each $1,000 saved. If his income this year is OA = $5,000, his possibilities curve is *AB*. He will choose position C, where *AB* is tangent to the highest indifference curve he can reach. If his income rises to $7,000, his possibilities curve becomes A_1B_1, and he will move to C_1, where he saves more and spends more than before. If with the $5,000 income the rate of interest falls to 4 percent, his possibilities curve becomes AB_2 and he will be forced down to C_2.

What possibilities are actually open to him? This depends on his income and on the rate of interest. Suppose his income is $OA = \$5{,}000$ and the market rate of interest is 5 percent. Then if he saved all his income, he could earn $OB = \$250$ of interest next year. If he spends all his income for consumption, he will, of course, earn zero interest; and he can do anything in between. His budget line is *AB*, showing all combinations of current spending and future interest earnings which are possible within the limits of his income. Its slope depends on the rate of interest, and will be steeper as the interest rate rises.

Mr. Goldbilt's best position in these circumstances will be at *C*, where *AB* is tangent to I_2, the highest indifference curve he can reach. So he cannot gain by a change in any direction. Here the rate at which he *is willing* to substitute present for future income (shown by the slope of I_2) just equals the rate at which he *can* substitute them through the money market (shown by the slope of *AB*).

Suppose now that his income rises to $OA_1 = \$7{,}000$. By saving all of

becomes A_1B_1, parallel to AB (since the rate of interest is unchanged) but at a higher level. He can now reach the higher indifference curve I_3, and his preferred position will be at C_1. Note that he is now saving more *and* spending more than before, which is the general rule as income rises.

The other thing that may happen is that the rate of interest may change. Go back to the first position in which his income is OA. Suppose now that the interest rate falls to 4 percent. The maximum he could earn even by saving all his income is now only \$5,000 × 4 percent = \$200. His budget line is AB_2 instead of AB.

This forces him down to a lower indifference curve, I_1, and he will settle at O_2, where AB_2 is tangent to I_1. Here he is spending more of his current income and saving less than he was before. It seems reasonable that if the reward for saving falls, the amount of saving done will decline.

It is possible, however, that some savers may not react in this way. Suppose you are determined to have a fixed amount of interest income in the future—say, \$1,000 per year—for retirement or some other purpose. You are what has been termed a "target saver." Then if the interest rate is 5 percent, you can attain your objective by saving \$20,000. But if the interest rate falls to 4 percent, you will have to save \$25,000. Thus your saving will *increase* as the interest rate *decreases*. We could illustrate this kind of reaction by giving a different shape to the indifference curves in Figure 9.

The Choice of Work versus Leisure

Let's look finally at how this same approach can be applied to decisions about hours of work. Mr. Restwell's work preference system is shown in Figure 10. Mr. Restwell is, say, a commission salesman who can work as few or as many hours per week as he wishes. He can get more income by working longer hours, but this means a sacrifice of leisure. On the vertical axis we show weekly income in dollars. On the horizontal axis we show various ways in which Mr. Restwell might divide, say, 100 waking hours per week between work and leisure. Hours worked appear above the line, leisure hours below the line, and the two, of course, move in opposite directions. As we move to the right, leisure increases and hours worked decline.

The indifference curves I_1 and I_2 show different combinations of income and leisure which Mr. Restwell would consider equally attractive. At a point far to the right on I_1, indicating short hours and low income, he will be willing to work longer for only a little more income. As we move up I_1 to the left, however, and his work week becomes steadily longer, it takes more and more money to persuade him to work an extra

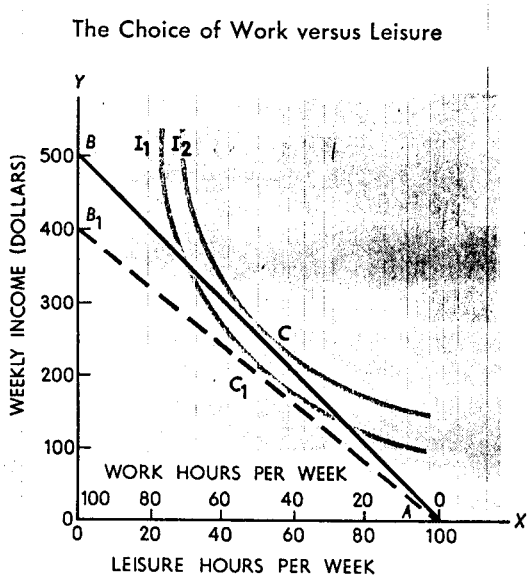

FIGURE 10. The horizontal axis shows different ways in which Mr. Restwell can divide his 100 waking hours between work and leisure. As we move to the right, leisure increases and work decreases. The indifference curves I_1 and I_2 show combinations of income and leisure which would be equally acceptable to Mr. Restwell. If the wage on his job is \$5 an hour, his possibilities curve is AB, and his preferred position is C, where AB is tangent to the highest possible indifference curve. A drop in the wage rate to \$4 would shift the possibilities curve to AB_1, and force him down to a less satisfactory position at C_1.

Where Mr. Restwell will come out depends on the wage rate for his job. Suppose this is \$5 an hour. Then if he worked the absolute maximum of 100 hours per week, he could earn OB = \$500. If he doesn't work at all, shown by OA on the horizontal axis, he will, of course, earn zero; and he can do anything in between. His budget line is AB, the slope of which depends on the hourly wage rate. His preferred work week is shown by point C, where AB is tangent to the highest possible indifference curve. At this point he is working about 50 hours and earning \$250 per week.

As always, this result depends on the shape which we give to the indifference curves. With a stroke of the pencil we can make Mr. Restwell greedy or abstemious, idle or diligent, and bring him out with a long work week or a very short one. And while for illustrative purposes we can do anything we choose, it must never be forgotten that each individual's work preference system is a real and unique thing.

The outcome depends also on the slope of AB, i.e., on the hourly wage rate. Suppose that, instead of \$5, the wage rate were only \$4. Mr. Restwell's possibilities curve is now AB_1, and the best adjustment he can make is at C_1, where AB_1 just touches the indifference curve I_1. Since I_1 is below I_2, this is a less satisfactory position than before—it is impossible to be as well off at a lower wage rate as at a higher one.

Will the lower wage cause Mr. Restwell to work longer hours or shorter hours than before? In Figure 10, he ends up working longer hours. But this will not necessarily be true in all cases. The best way to see this is to regard the hourly wage rate as *the price of leisure*. How can

I get an additional hour of leisure? By working one hour less than before. But this means giving up the wage that I might have earned in that hour. So the higher the wage, the higher is the price of leisure; and conversely.

A lower wage, then, means a cut in the price of leisure. Just like the meat price change noted earlier, this sets up two conflicting reactions: (1) Since leisure is now cheaper than before, there is an incentive to buy more of it, that is, to work shorter hours. This is the *substitution effect.* (2) But at the lower wage level Mr. Restwell is also poorer than before. So he will feel less able to afford all the good things of life, *including leisure.* This *income effect* will stimulate him to work longer hours, in order to bring his weekly income back closer to its previous level.

The outcome depends on which effect is stronger in a particular case. If the income effect predominates, as we have shown it doing in Figure 10, a lower wage rate will mean longer hours, and a higher wage rate shorter hours. The opposite result is quite possible, however, and could be illustrated by altering the shape of the indifference curves.