

Price Theory: An Intermediate Text

by

David D. Friedman

Published by South-Western Publishing Co.

©David D. Friedman 1986, 1990

Table of Contents

Introduction

[Preface](#)

Section I ECONOMICS FOR PLEASURE AND PROFIT

Chapter [1](#) What is Economics?

[2](#) How Economists Think.

Section II PRICE=VALUE=COST: COMPETITIVE EQUILIBRIUM IN A SIMPLE ECONOMY

Chapter [3](#) The Consumer: Choice and Indifference Curves

[4](#) The Consumer: Marginal Value, Marginal Utility, and Consumer Surplus

[5](#) Production

[6](#) Simple Trade

[7](#) Markets&endash;Putting it All Together

[8](#) The Big Picture

[Halftime](#)

Section III COMPLICATIONS, OR ONWARD TO REALITY

Chapter [9](#) The Firm

[10](#) Small-Numbers Problems: Monopoly and All That

[11](#) Hard Problems: Game Theory, Strategic Behavior, and Oligopoly

[12](#) Time...

[13](#) ...and Chance

[14](#) The Distribution of Income and the Factors of Production

Section IV JUDGING OUTCOMES

Chapter [15](#) Economic Efficiency

[16](#) What is Efficient?

[17](#) Market Interference

[18](#) Market Failures

Section V APPLICATIONS & CONVENTIONAL AND UN

Chapter [19](#) The Political Marketplace

[20](#) The Economics of Law and Law Breaking

[21](#) The Economics of Love and Marriage

Section VI WHY YOU SHOULD BUY THIS BOOK

Chapter [22](#) Final Words

Additional Chapters from the First Edition not included in the Second

Chapter [21](#) The Economics of Heating

[22](#) Inflation and Unemployment

The author retains all rights in this material, save that users of the World Wide Web are permitted to reproduce it to the extent, and only to the extent, that doing so is a necessary part of reading it on the web.

The printed version of the book, along with supplementary materials, is available from South-Western Publishing, Cincinnati, OH. My new book *Hidden Order: The Economics of Everyday Life* offers a similar approach to explaining economics in a shorter form, aimed at the intelligent layman rather than at students taking intermediate micro. Click here for the [table of contents](#), and here for a link to [My Publisher's Page](#).

Section 1

Economics for Pleasure and Profit

Chapter 1

What Is Economics?

Economics is often thought of either as the answers to a particular set of questions (How do you prevent unemployment? Why are prices rising? How does the banking system work? Will the stock market go up?) or as the method by which such answers are found. Neither description adequately defines economics, both because there are other ways to answer such questions (astrology, for example, might give answers to some of the questions given above, although not necessarily the right answers) and because economists use economics to answer many questions that are not usually considered "economic" (What determines how many children people have? How can crime be controlled? How will governments act?).

I prefer to define economics as a particular way of understanding behavior; what are commonly thought of as economic questions are simply questions for which this way of understanding behavior has proved particularly useful in the past:

Economics is that way of understanding behavior that starts from the assumption that people have objectives and tend to choose the correct way to achieve them.

The second half of the assumption, that people tend to find the correct way to achieve their objectives, is called *rationality*. This term is somewhat deceptive, since it suggests that the way in which people find the correct way to achieve their objectives is by rational analysis--analyzing evidence, using formal logic to deduce conclusions from assumptions, and so forth. No such assumption about how people find the correct means to achieve their ends is necessary.

One can imagine a variety of other explanations for rational behavior. To take a trivial example, most of our objectives require that we eat occasionally, so as not to die of hunger (exception--if my objective is to be fertilizer). Whether or not people have deduced this

fact by logical analysis, those who do not choose to eat are not around to have their behavior analyzed by economists. More generally, evolution may produce people (and other animals) who behave rationally without knowing why. The same result may be produced by a process of trial and error; if you walk to work every day, you may by experiment find the shortest route even if you do not know enough geometry to calculate it. Rationality in this sense does not necessarily require thought. In the final section of this chapter, I give two examples of things that have no minds and yet exhibit rationality.

Half of the assumption in my definition of economics was rationality; the other half was that people have objectives. In order to do much with economics, one must strengthen this part of the assumption somewhat by assuming that people have *reasonably simple objectives*; with no idea at all about what people's objectives are, it is impossible to make any prediction about what people will do. Any behavior, however peculiar, can be explained by assuming that the behavior itself was the person's objective. (Why did I stand on my head on the table while holding a burning \$1,000 bill between my toes? I *wanted* to stand on my head on the table while holding a burning \$1,000 bill between my toes.)

To take a more plausible example of how a somewhat complicated objective can lead to apparently irrational behavior, consider someone who has a choice between two identical products at different prices. It seems that for almost any objective we can think of, he would prefer to buy the less expensive item. If his objective is to help the poor, he can give the money he saves to the poor. If his objective is to help his children, he can spend the money he saves on them. If his objective is to live a life of pleasure and luxury, he can spend the money on Caribbean cruises and caviar.

But suppose you are taking a date to a movie. You know you are going to want a candy bar, which costs \$1.00 in the theater and \$0.50 in the Seven-Eleven grocery you pass on your way there. Do you stop at the store and buy a candy bar? Do you want your date to think you are a tightwad? You buy the candy bar at the theater, impressing your date (you hope) with the fact that you are the sort of person who does not have to worry about money.

One could get out of this problem by claiming that the two candy bars are not really identical; the candy bar at the theater includes the additional characteristic of impressing your date. But if you follow this line of argument, no two items are identical and the statement that you prefer the lower priced of two identical items has no content. I would prefer to say that the two items are identical enough for our purposes but that in this

particular case your objective is sufficiently odd so that our prediction (based on the assumption of reasonably simple objectives) turns out to be wrong.

WHY ECONOMICS MIGHT WORK

Economics is based on the assumption that people have reasonably simple objectives and choose the correct means to achieve them. Both halves of the assumption are false; people sometimes have very complicated objectives and they sometimes make mistakes. Why then is the assumption useful?

Suppose we know someone's objective and also know that half the time that person correctly figures out how to achieve it and half the time acts at random. Since there is generally only one right way of doing things (or perhaps a few) but very many wrong ways, the "rational" behavior can be predicted but the "irrational" behavior cannot. If we predict this person's behavior on the assumption that he is rational, we will be right half the time. If we assume he is irrational, we will almost never be right, since we still have to guess *which* irrational thing he will do. We are better off assuming he is rational and recognizing that we will sometimes be wrong. To put the argument more generally, the tendency to be rational is the consistent (and hence predictable) element in human behavior. The only alternative to assuming rationality (other than giving up and concluding that human behavior cannot be understood and predicted) would be a *theory* of irrational behavior--a theory that told us not only that someone would not always do the rational thing but also *which particular irrational thing* he would do. So far as I know, no satisfactory theory of that sort exists.

There are a number of reasons why the assumption of rationality may work better than one would at first think. One is that we are often concerned not with the behavior of a single individual but with the aggregate effect of the behavior of many people. Insofar as the irrational part of their behavior is random, its effects are likely to average out in the aggregate.

Suppose, for example, that the rational thing to do is to buy more hamburger the lower its price. People actually decide how much to buy by first making the rational decision then flipping a coin. If the coin comes up heads, they buy a pound more than they were planning to; if it comes up tails, they buy a pound less. The behavior of each individual

will be rather unpredictable, but the total demand for hamburger will be almost exactly the same as without the coin flipping, since on average about half the coins will come up heads and half tails.

A second reason why the assumption works better than one might expect is that we are often dealing not with a random set of people but with people who have been selected for the particular role they are playing. Consider the heads of companies. If you selected people at random for the job, the assumption that they want to maximize the company's profits and know how to do so would not be a very plausible one. But people who do not want to maximize profits, or do not know how to, are unlikely to be chosen for the job; if they are, they are unlikely to keep it; if they do, their companies are likely to become increasingly unimportant in the economy, until eventually the companies go out of business. So the simple assumption of profit maximization plus rationality turns out to be a good way to predict how firms will behave.

A similar argument applies to the stock market. We may reasonably expect that the average investment is made by someone with an accurate idea of what companies are worth--even though the average American, and even the average investor, may be poorly informed about such things. Investors who consistently bet wrong on the stock market soon have very little to bet with. Investors who consistently bet right have an increasing amount of their own money to risk--and often other people's money as well. Hence the well-informed investors have an influence on the market out of proportion to their numbers as a fraction of the population. If we analyze the workings of the market on the assumption that all investors are well informed, we may come up with fairly accurate predictions in spite of the inaccuracy of the assumption. In this as in all other cases, the ultimate test of the method is whether its predictions turn out to describe reality correctly. Whether something is an economic question is not something we know in advance. It is something we discover by trying to use economics to answer it.

SOME SIMPLE EXAMPLES OF ECONOMIC THINKING

So far, I have talked of economics in the abstract; it is now time for some concrete examples. I have chosen examples involving issues not usually considered economic in order to show that economics is not a particular set of questions to be answered but a particular way of answering questions. I will begin with two very simple examples and

then go on to some slightly more complicated ones.

You are laying out a college campus as a rectangular pattern of concrete sidewalks with grass between them. You know that one of the objectives of many people, including many students, is to get where they are going with as little effort as possible; you suspect most of them realize that a straight line is the shortest distance between two points. You would be well advised to take precautions against students cutting across the lawn. Possible precautions would be constructing fences or diagonal walkways, adding tough ground cover, or replacing the grass with cement and painting it green.

One point to note. It may be that everyone will be better off if no one cuts across the lawn (assuming the students like to look at green lawns without brown paths across them). Rationality is an assumption about individual behavior, not group behavior. The question of under what circumstances individual rationality does or does not lead to the best results for the group is one of the most interesting questions economics investigates. Even if a student is in favor of green grass, he may correctly argue that his decision to cut across provides more benefit (time saved) than cost (slight damage to the grass) *to him*. The fact that his decision provides additional costs, but no additional benefits, to other people who also dislike having the grass damaged is irrelevant unless making those other people happy happens to be one of his objectives. The total costs of his action may be greater than the total benefits; but as long as the costs to him are less than the benefits to him, he takes the action. This point will be examined at much greater length in Chapter 18, when we discuss public goods and externalities.

A second simple example of economic thinking is Friedman's Law for Finding Men's Washrooms--"Men's rooms are adjacent, in one of the three dimensions, to ladies' rooms." One of the builder's objectives is to minimize construction costs; it costs more to build two small plumbing stacks (the set of pipes needed for a washroom) than one big one. So it is cheaper to put washrooms close to each other in order to get them on the same stack. That does not imply that two men's rooms on the same floor will be next to each other (although men's rooms on different floors are usually in the same position, making them adjacent vertically). Putting them next to each other reduces the cost, but separating them gets them close to more users. But there is no advantage to having men's and ladies' rooms far apart, since they are used by different people, so they are almost always put on the same stack. The law does not hold for buildings constructed on government contracts at cost plus 10 percent.

As a third example, consider someone making two decisions--what car to buy and what politician to vote for. In either case, the person can improve his decision (make it more likely that he acts in his own interest) by investing time and effort in studying the alternatives. In the case of the car, his decision determines with certainty which car he gets. In the case of the politician, his decision (whom to vote for) changes by one ten-millionth the probability that the candidate he votes for will win. If the candidate would be elected without his vote, he is wasting his time; if the candidate would lose even with his vote, he is also wasting his time. He will rationally choose to invest much more time in the decision of which car to buy--the payoff to him is enormously greater. We expect voting to be characterized by *rational ignorance*; it is rational to be ignorant when the information costs more than it is worth.

This is much less of a problem for a concentrated interest than for a dispersed one. If you, or your company, receives almost all of the benefit from some proposed law, you may well be willing to invest enough resources in supporting that law (and the politician who wrote it) to have a significant effect on the probability that the law will pass. If the cost of the law is spread among many people, no one of them will find it in his interest to discover what is being done to him and oppose it. Some of the implications of that will be seen in Chapter 19, where we explore the economics of politics.

In the course of this example, I have subtly changed my definition of rationality. Before, it meant making the right decision about *what to do*--voting for the right politician, for example. Now it means making the right decision about *how to decide what to do*--collecting information on whom to vote for only if the information is worth more than the cost of collecting it. For many purposes, the first definition is sufficient. The second is necessary where an essential part of the problem is the cost of getting and using information.

A final, and interesting, example is the problem of winning a battle. In modern warfare, many soldiers do not fire their guns in battle, and many of those who fire do not aim. This is not irrational behavior--on the contrary. In many situations, the soldier correctly believes that nothing he can do is very likely to determine who wins the battle; if he shoots, especially if he takes time to aim, he is more likely to get shot himself. The general and the soldier have two objectives in common. Both want their army to win. Both also want the soldier to survive the battle. But the relative importance of the second objective is much greater for the soldier than for the general. Hence the soldier rationally does not do what the general rationally wants him to do.

Interestingly enough, studies of U.S. soldiers in World War II revealed that the soldier most likely to shoot was the member of a squad who was carrying the Browning Automatic Rifle. He was in a situation analogous to that of the concentrated interest; since his weapon was much more powerful than an ordinary rifle (an automatic rifle, like a machine gun, keeps firing as long as you keep the trigger pulled), his actions were much more likely to determine who won--and hence whether he got killed--than the actions of an ordinary rifleman.

The problem is not limited to modern war. The old form of the problem (which still exists in modern armies) is the decision whether to stand and fight or to run away. If you all stand, you will probably win the battle. If everyone else stands and you run, your side may still win the battle and you are less likely to get killed (unless your own side notices what you did and shoots you) than if you fought. If everyone runs, you lose the battle and are quite likely to be killed--but less likely the sooner you start running.

One proverbial solution to this problem is to burn your bridges behind you. You march your army over a bridge, line up on the far side of the river, and burn the bridge. You then point out to your soldiers that if your side loses the battle you will all be killed, so there is no point in running away. Since your troops do not run and the enemy troops (hopefully) do, you win the battle. Of course, if you lose the battle, a lot more people get killed than if you had not burned the bridge.

We all learn in high school history how, during the Revolutionary War, the foolish British dressed their troops in bright scarlet uniforms and marched them around in neat geometric formations, providing easy targets for the heroic Americans. My own guess is that the British knew what they were doing. It was, after all, the same British Army that less than 40 years later defeated the greatest general of the age at Waterloo. I suspect the mistake in the high school history texts is not realizing that what the British were worried about was controlling their own troops. Neat geometric formations make it hard for a soldier to advance to the rear unobtrusively; bright uniforms make it hard for soldiers to hide after their army has been defeated, which lowers the benefit of running away.

The problem of the conflict of interest between the soldier as an individual and the soldiers as a group is nicely illustrated by the story of the battle of Clontarf, as given in *Njal Saga*. Clontarf was an eleventh century battle between an Irish army on one side and a mixed Irish-Viking army on the other side. The Vikings were led by Sigurd, the Jarl of the Orkney Islands. Sigurd had a battle flag, a raven banner, of which it was said that as

long as the flag flew, his army would always go forward, but whoever carried the flag would die.

Sigurd's army was advancing; two men had been killed carrying the banner. The Jarl told a third man to take the banner; the third man refused. After trying unsuccessfully to find someone else to do it, Sigurd remarked, "It is fitting the beggar should bear the bag," cut the banner off the staff, tied it around his own waist, and led the army forward. He was killed and his army defeated. The story illustrates nicely the essential conflict of interest in an army, and the way in which individually rational behavior can prevent victory. If one or two more men had been willing to carry the banner, Sigurd's army might have won the battle--but the banner carriers would not have survived to benefit from the victory.

And you thought economics was about stocks and bonds and the unemployment rate.

PUZZLE

You are a hero with a broken sword (Conan, Boromir, or your favorite Dungeons and Dragons character) being chased by a troop of bad guys (bandits, orcs, . . .). Fortunately you are on a horse and they are not. Unfortunately your horse is tired and they will eventually run you down. Fortunately you have a bow. Unfortunately you have only ten arrows. Fortunately, being a hero, you never miss. Unfortunately there are 40 bad guys. The bad guys are strung out behind you, as shown.

Problem: Use economics to get away.

Note: You cannot talk to the bad guys. They are willing to take a substantial chance of being killed in order to get you--after all, they know you are a hero and are still coming. They know approximately how many arrows you have.

OPTIONAL SECTION

SOME HARDER EXAMPLES--ECONOMIC EQUILIBRIA

So far, the examples of economic reasoning have not involved any real interaction among the rational acts of different people. We dealt either with a single rational individual--the architect deciding where in the building to put washrooms--or with a group of rational individuals all doing more or less the same thing. Very little in economics is this simple. Before we start developing the framework of price theory in the next chapter, you may find it of interest to think through some more difficult examples of economic reasoning, examples in which the outcome is an equilibrium produced by the interaction of a number of rational individuals.

I will use economics to analyze two familiar situations (supermarket lines and crowded expressways), showing how economics can produce useful and nonobvious results and how the argument can be expanded to deal with successively higher levels of complexity. The logical patterns that appear in these examples reappear again and again in economic analysis. Once you clearly understand when and why supermarket lines are all the same length and lanes in the expressway equally fast, and why and under what circumstances they are not, you will have added to your mental tool kit one of the most useful concepts in economics.

Supermarket Lines

You are standing in a supermarket at the far end of a row of checkout counters with your arms full of groceries. The line at your end blocks your view of the other lines; you know your line is long, but you do not know if the others are any shorter. Should you stagger from line to line looking for the shortest line, or should you get in the nearest one?

The first and simplest answer is that all the lines will be about the same length, so you should get into the one next to you; it is not worth the cost of searching for a shorter one. Why?

Consider any two adjacent lines in Figure 1-1, say Lines 4 and 5. Some shoppers will approach the checkout area not from one end, as you did, but from the aisle that lies between those two lines. Since those shoppers can easily see both lines, they will go to whichever one appears shorter. By doing so, they will lengthen that line and shorten the other; the process continues until both lines are the same length. The same argument holds for every other pair of adjacent lines, so all lines will be the same length. It is not worth it for you to make a costly search for the shortest line.

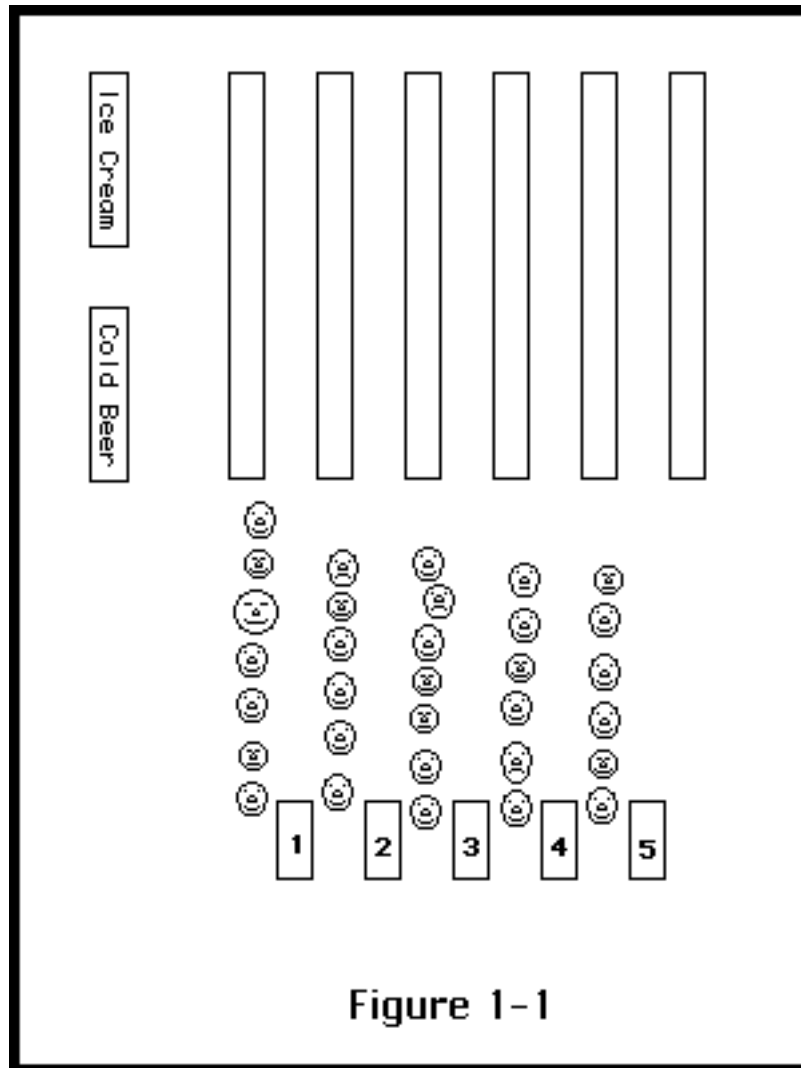
There are a number of implicit assumptions in this argument. When these assumptions are false the argument may break down. Suppose, for example, that you are at the far end of the row of checkout counters because that is where the ice cream freezer and the refrigerator with the cold beer are located. Many other customers also choose to get these things last and so enter the checkout area from that end. Even if everyone who comes in between Lines 1 and 2 goes to Line 2, there are not enough such people to make Line 2 as long as Line 1. If everyone understands the argument of the previous paragraph and acts accordingly, Line 1 will be longer than Line 2 (and probably much longer than the other lines), and the conclusion of the argument will be wrong.

Imagine that you program a computer to assign customers to lines in a way that equalizes the length of the two lines, as described above, and tell it that 10 people per minute are entering the checkout area at one end (where they can only see Line 1) and 6 per minute are entering between the two lines. The computer informs you that of the 6 customers coming in between the two lines, 8 must go to Line 2 and -2 to Line 1. Since 10 customers are going to Line 1 from the end, the total number going to Line 1 is 10 plus -2, which equals 8--the number going to Line 2. The computer, having solved the problem you gave it, sits there with a satisfied expression on its screen.

You then reprogram it, pointing out that fewer than zero customers cannot go anywhere. Mathematically speaking, you are asking the computer to solve the problem subject to the condition that a certain number (the number of customers coming in between the two lines and going to one of them) cannot be negative. The computer replies that in that case, the best it can do is to send all six customers to Line 2--leaving the lines still unequal.

This sort of result is called a *corner solution* because it corresponds to the mathematical situation where the maximum of a function is not at the top of its graph but instead at a corner where the graph ends, as shown in Figure 1-2a. In such a situation, the normal conclusion (in the supermarket case, that all the lines must be the same length) may no

longer hold. The corresponding result in Figure 1-2a is that the graph is not horizontal at its maximum--as it would be if the maximum were at an *interior solution*, as it is in Figure 1-2b. In economics--especially mathematical economics--the usual role of corner solutions is to provide annoying exceptions to general theorems.



Supermarket, viewed from above. Lines tend to be equal; Line 1 is a special case because many customers get ice cream and cold beer last.


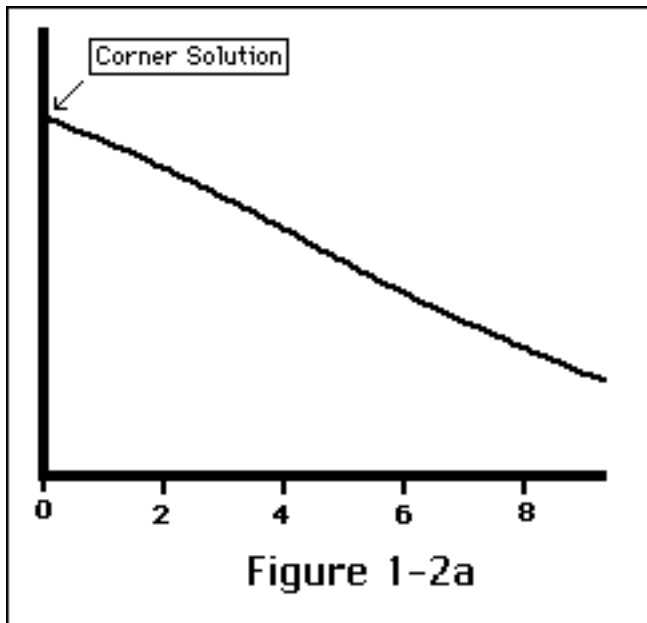
Are there other situations in which the conclusion--that all lines will be the same length--does not hold? Yes.

So far, I have assumed that for people coming in between two lines, it is costless to see

which line is shorter. This is not always true. The relevant length, after all, is not in space but in time; you would rather enter a line of ten customers with only a few items each than a line of eight customers with full carts. Estimating which line is shorter requires a certain amount of mental effort. If the system works so well that all lines are exactly the same length (in time), then it will never be worth that effort. Hence no one will make it; hence there will be nothing keeping the lines the same length. In equilibrium the length of lines must differ by just enough to repay (on average) the effort of figuring out which line is shorter. If it differed by more than that, everyone would look for the shortest line, making all lines the same length (assuming no corner solution). If it differed by less than that, nobody would.

It may have occurred to you that I am assuming all customers have the same ability to estimate how long a line will take. Suppose a few customers know that the checker on Line 3 is twice as fast as the others. The experts go to Line 3. Line 3 appears to be longer than the other lines (to nonexperts, that is; allowing for the fast checker, the line is actually shorter, in time although not in length). nonexperts avoid Line 3 until it shrinks back to the same length as the others. The experts (and some lucky nonexperts--the ones who are still in Line 3) get out twice as fast as everyone else.

Word spreads; the number of experts increases. As long as, with all the experts going through Line 3, Line 3 can still be as short (in appearance) as the other lines, the increasing number of experts does not reduce the payoff to being an expert. Every time one more expert enters the line (making it appear slightly longer than the others), one more nonexpert decides not to enter it.



Two maxima--a corner solution (a) and an interior solution (b). At the interior maximum, the slope of the curve is zero; at the corner maximum, it need not be.

Eventually the number of experts becomes large enough to crowd out all the nonexperts from that line. As the number of experts increases further, Line 3 begins to lengthen. It cannot be brought back to the same length as the other lines by the defection of nonexperts (who mistakenly believe that it is longer in waiting time as well as length) because there are none of them going to it and the experts know better. Eventually the number of experts becomes so great that Line 3 is twice as long as the other lines and takes the same length of time as they do; the gain from being an expert has now vanished.

To put the same argument in more conventional economic language, rational behavior (in the sense of "making the right decision") requires information. If that information is itself costly, rational behavior consists of acquiring information (paying information costs) only as long as the return from additional information is at least as great as the cost of getting it. If certain minimal information is required to equalize the time-length of lines, then the time-length of lines must be sufficiently unequal so that the saving from knowing which line is shorter just pays the cost of acquiring that information. That principle applies to both the cost of looking at lines to see which is shortest and the cost of studying checkers to learn which ones are faster. The initial argument was given in an approximation in which information was costless; such an approximation greatly simplifies many economic arguments but should be used with care.

There is at least one more hidden assumption in the argument as given. I have assumed that everyone in the grocery store wants to get out as quickly as possible. Suppose the grocery store (Westwood Singles Market) is actually the local social center; people come to stand in long lines gossiping with and about their friends and trying to make new ones. Since they do not want to get out as fast as possible, they do not try to go to the shortest line; so the whole argument breaks down.

Rush Hour Blues

A similar analysis can be applied to lanes on the freeway. When you are driving on a crowded highway, it always seems that some other lane is going faster than yours; the obvious strategy is to switch to the faster lane. If you actually try to follow such a strategy, however, you discover to your amazement that a few minutes after you switch lanes, the battered blue pickup that was behind you in the lane you left is now in front of you.

To understand why it is so difficult to follow a successful strategy of lane changing, consider that by moving into a lane you slow it down. If there is a faster lane then people will move into it, equalizing its speed with that of the other lanes, just as people moving into a short line lengthen it. So a lane remains fast only as long as drivers do not realize it is.

Here again, a more sophisticated analysis would allow for the costs (in frayed nerves and dented fenders) of continual lane changes. On average, if everyone is rational, there must be a small gain in speed from changing lanes--if there were not, nobody would do it and the mechanism described above would not work. The payoff must equal the cost for the *marginal* lane changer--the least well qualified of those following the lane-changing strategy. If the payoff were less than that, he would not be a lane changer; if it were more, someone else would. In principle, if you knew how much a strategy of lane changing cost each driver (in dents and nerves--less for those with strong nerves and old cars) and how many lane changers it took to reduce the benefit from lane changing by any given amount, you could figure out who would be the marginal lane changer and how much the gain from lane changing would be. By the end of the course, you should see how to do this. If you see it now, you are already an economist--whether or not you have studied economics.

Even More Important Applications to Think About

Doctors make a lot of money. Doctors also spend many years as medical students and interns. The two facts are not unrelated. Different wages in different professions are set by a process similar to that described above. If one profession is, on net, more attractive than another (taking account of wages, risks, costs of learning the profession, and so on), more people go into the more attractive profession and by so doing drive down the wages. All professions are in some sense equally attractive--to the marginal person. In deciding what profession you want to enter, it is not enough to ask what profession pays the highest wage. Not only are there other factors, there is also reason to expect that the other factors will be worst where the wage is best. What you should ask instead is what profession you are particularly suited for in comparison to other people making similar choices. This is like deciding whether to follow a lane-switching strategy by considering how old your car is compared to others, or deciding whether to look for a shorter line in the grocery store according to how much you are carrying.

A similar argument applies to the stock market. It is often said that if a company is doing very well, you should buy its stock. But if everyone else knows that the company is doing well, then the price of its stock already reflects that information. If buying it were really such a good deal, who would sell? The company you should buy stock in is one that you know is doing better than most other investors think it is--even if in some absolute sense it is not doing very well.

A friend of mine has been investing successfully for several years by following almost the opposite of the conventional wisdom. He looks for companies that are doing very badly and calculates how much their assets would be worth if they went out of business. Occasionally he finds one whose assets are worth more than its stock. He buys stock in such companies, figuring that if they do well their stock will go up and if they do badly they will go out of business, sell off their assets--and the stock will again go up.

If all of this is obvious to you the first time you read it (or even the second), then in your choice of careers you should give serious consideration to becoming an economist.

NEGATIVE FEEDBACK

Several of the situations described in this chapter involved a principle called negative feedback. A familiar example of negative feedback is driving a car. If the car is going to the right of where you want it, you turn the wheel a little to the left; if it is going to the left of where you want, you turn it a little to the right. This is called feedback because an error in the direction you are going "feeds back" into the mechanism that controls your direction (through you to the steering wheel). It is negative feedback because an error in one direction (right) causes a correction in the other direction (left). An example of positive feedback is the shriek when the amplifier attached to a microphone is turned up too high. A small noise comes into the mike, is amplified by the amplifier, comes out of the speaker, and feeds back into the mike. If the amplification is high enough, the noise becomes louder each time around, eventually overloading the system.

In the supermarket line example, the lines are kept at about the same length by negative feedback: If a line gets too long compared to other lines people stop going to it, which makes it get shorter. Similarly, when a lane on the expressway speeds up, cars move into it, slowing it down. In each case, what we are mostly interested in are not the details of the feedback process but rather the nature of the stable equilibrium--the situation such that deviations from it cause correcting feedback.

RATIONALITY WITHOUT MIND

In defending the assumption of rationality, I pointed out that it is not the same as the assumption that people reason logically. Logical reasoning is not the only, or even the most common, way of getting a correct answer. I will demonstrate this with two extreme examples--cases in which we observe rationality in something that cannot reason, since it has no mind to reason with. In the first case, I will show how a mindless object--a collection of matchboxes filled with marbles--can learn to play a game rationally. In the second, I will show how the rational pursuit of objectives by genes--mindless chains of atoms inside your cells--explains a striking fact about the real world, something so fundamental that it never occurs to most of us to find it surprising.

Computers that Learn

Suppose you want to build a computer to play some simple game, such as tic-tac-toe. One way is to build in the correct move for every situation. Another, and in some ways more interesting, approach is to let the computer teach itself how to play. Such a learning computer starts out moving randomly. Each time a game ends, the computer is told whether it won or lost and adjusts its strategy accordingly, lowering the probability of moves that led to losses and increasing the probability of moves that led to wins. After enough games, the computer may become a fairly good player.

The computer does not think. Its "mind" is simply a device that identifies the present situation of the game, chooses a move by some random mechanism, and later adjusts the probabilities according to whether it won or lost. A simple version consists of a bunch of matchboxes filled with black and white marbles, laid out on a diagram of the game. Moves are chosen by picking a marble at random, with the color of the marble determining the move. The mix of marbles in each matchbox is adjusted at the end of the game to make moves that led to a win more likely and moves that led to a loss less likely.

A matchbox computer, or its more sophisticated electronic descendants, does not think, yet it is rational. Its objective is to win the game and, after it has played long enough to "learn" how to win, it tends to choose the correct way of achieving that objective. We can understand and predict its behavior in the same way that we understand and predict the behavior of humans. "Rationality" is simply the ability to get the right answer; it may be the result of many things other than rational thinking.

Economics and Evolution

There is a close historical connection between economics and evolution. Both of the discoverers of the theory of evolution (Darwin and Wallace) said they got the idea from Thomas Malthus, an economist who was also one of the originators of the so-called Ricardian Theory of Rent (named after David Ricardo, who used it but did not invent it), one of the basic building blocks of modern economics.

There is also a close similarity in the logical structure of the two fields. The economist expects people to choose correctly how to achieve their objectives but is not very much concerned with the psychological question of how they do so. The evolutionary biologist expects genes--the fundamental units of heredity that control the construction of our

bodies--to construct animals whose structure and behavior are such as to maximize their reproductive success (roughly speaking, the number of their descendants), since the animals that presently exist are descended from those that were reproductively successful in the past and carry the genes that made them successful. The biologist need not be concerned very much with the detailed biochemical mechanisms by which the genes control the organism. Many of the same patterns appear in both economics and evolutionary biology; the conflict between individual interest and group interest that I mentioned earlier reappears in the conflict between the interest of the gene and the interest of the species.

A nice example is Sir R.A. Fisher's explanation of observed sex ratios. In many species, including ours, male and female offspring are produced in roughly equal numbers. There is no obvious reason why this is in the interest of the species; one male suffices to fertilize many females. Yet the sex ratio remains about 1:1, even in some species in which only a small fraction of the males succeed in reproducing. Why?

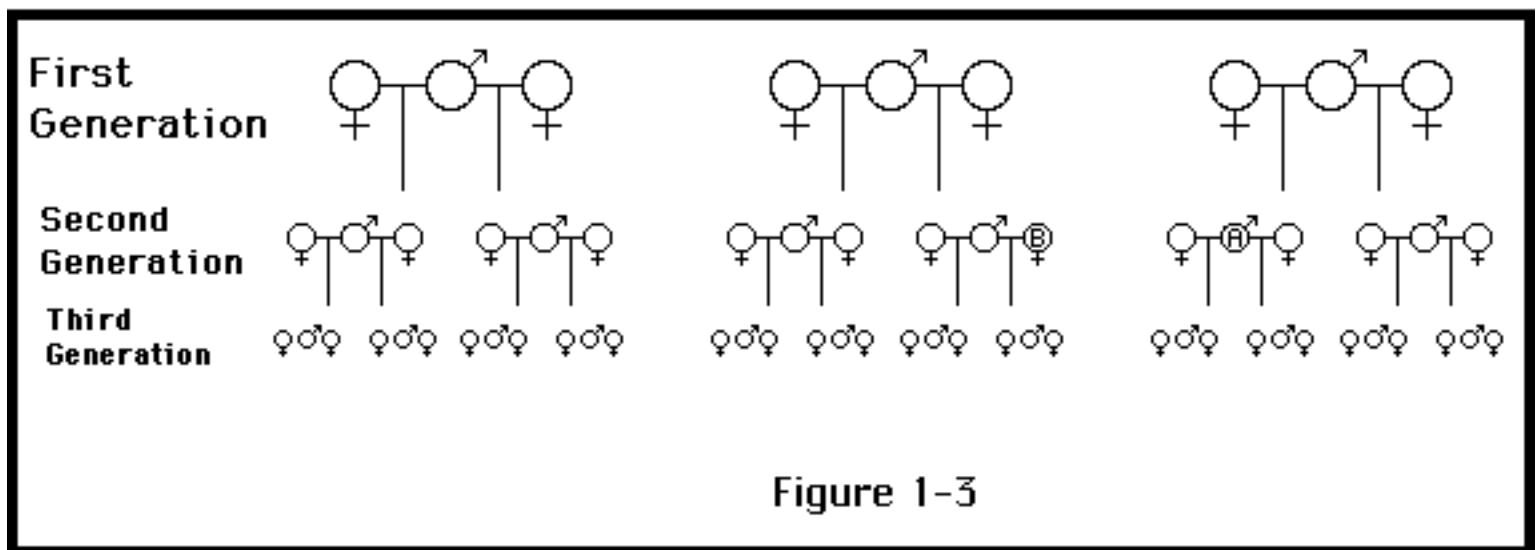
Fisher's answer is as follows. Imagine that two thirds of offspring are female, as shown in Figure 1-3. Consider three generations. Since each individual in the third generation has both a father and a mother, if there are twice as many females as males in the second generation, the average male must have twice as many children as the average female. This means that an individual in the first generation who produces a son will, on average, have twice as many grandchildren as one who produces a daughter. Individual A on Figure 1-3, for example, has six children, while Individual B only has three. A's parents got twice as great a return in grandchildren for producing A as B's parents did for producing B.

If there are more females than males in the population, couples who produce sons have more descendants, on average, than those who produce daughters. Since couples who produce sons have more descendants, more of the population is descended from them and has their genes--including the gene for having sons. Genes for producing male offspring increase in the population.

The initial situation, in which two thirds of the population in each generation was female, is unstable. As long as more than half of the children are female, genes for having male children spread faster than genes for having female children; so the percentage of female children falls. Similarly, if more than half the children were male, genes for having female children would have the advantage and spread. Either way, the situation must swing back

towards an even sex ratio.

In making this argument, I implicitly assumed equal cost for producing male and female offspring. In a species with substantial sexual dimorphism (male and female babies of different size), the argument implies that the total weight of female offspring (weight per offspring times number of offspring) will be about the same as that for male offspring. One could add further complications by considering differences in the costs of raising male and female offspring to maturity. Yet even the simple argument is strikingly successful in explaining one of the observed regularities of the world around us by the "rational" behavior of microscopic entities. Genes cannot think--yet in this case and many others, they behave as if they had carefully calculated how to maximize their own survival in future generations.



Three generations of a population with a male:female ratio of 1:2. Members of the first generation who have a son produce twice as many grandchildren as those who have a daughter, so genes for having sons increase in the population, swinging the sex ratio back toward 1:1.

PROBLEMS

1. In defending the rationality assumption, I argued that while people sometimes make mistakes, their correct decisions are predictable and their mistakes are not. Can you think

of any alternative approaches to understanding human behavior that claim to predict the mistakes? Discuss.

2. Give examples (other than buying candy for your date--the example discussed in the text) of apparently irrational behavior that consists of choosing the correct means to achieve an odd or complicated end.

3. In this chapter and throughout the book I treat individual preferences as givens--I neither judge whether people have the "right" preferences nor consider the possibility that something might change individual preferences.

a. Do you think some preferences are better than others? Give examples. Discuss.

b. Describe activities that you believe can only be understood as attempts to change people's preferences. How would you try to analyze such activities in economic terms?

4. Friedman's Law for Finding Men's Washrooms could be described as fossilized rationality--whether the architect lives or dies, his rationality remains set in concrete in the building he designed.

a. Can you think of other examples? Discuss.

b. Can you describe any cases where instead of deducing the shape of something from the rationality of its maker, we deduce the rationality of its maker from its shape? Discuss.

5. What devices (other than those discussed in the text) are used by generals, ancient and modern, to prevent soldiers from concluding that it is in their interest to run away, not aim, or in some other way act against the interest of the army of which they are a part?

6. The problem I have discussed exists not only in your army but in the enemy's army as well. Discuss ways in which a general might take advantage of that fact, giving real-world examples if possible.

7. In a recent conversation with one of our deans, I commented that I was rather absent-minded--I had missed two or three faculty meetings that year--and wished I could get him to make a point of reminding me when I was supposed to be somewhere. He replied that he had already solved that problem, so far as the (luncheon) meetings he was responsible for. He made sure I would not forget them by always arranging to have a scrumptious

chocolate dessert.

a. Is this an economic solution to the problem of getting me to remember things? Discuss.

b. In what sense does or does not the success of this method indicate that I "choose" to forget to go to meetings? Discuss.

8. This chapter discusses situations where rational behavior by each individual leads to results that are undesirable for all. Give an example of such a situation in your own experience; it should not be one discussed in the chapter.

9. Many voters are rationally ignorant of the names of their congressmen. List some things you are rationally ignorant of. Explain why your ignorance is rational. Extra credit if they are things that many people would say you ought to know.

The following problems refer to the optional section:

10. The analyses of supermarket lines, freeway lanes, and the stock market all had the same form. In each case, the argument could be summarized as "The outcome has a particular pattern because if it did not, it would be in the interest of people to change their behavior in a way that would push the outcome closer to fitting the pattern." Such a situation is called a stable equilibrium. Can you think of any examples not discussed in the text?

11. Analyze express lanes in supermarkets. Is the express lane always faster? If not, when is it and when is it not?

12. In the supermarket example, I started by assuming that you had your arms full of groceries. Why? How does that assumption simplify the argument?

13. The friend whose investment strategy I described is a very talented accountant. When I met him, he was in his early twenties and was making a good income teaching accounting to people who wanted to pass the CPA exam. Does this have anything to do with his investment strategy?

14. Is there any reason why my accountant friend should prefer that this book, or at least this chapter, not be published?
15. Give some examples of negative and positive feedback in your own experience.
16. Certain professions are very attractive to their members and very badly paid. Consider the stereotype of the starving artist--or a friend of mine who is working part-time as a store clerk while trying to make a career as a professional lutenist. Is the association between job attractiveness and low pay accidental, or is there a logical connection? Discuss.
17. You have been collecting data on the behavior of a particular stock over many years. You notice that every Friday the 13th, the stock drops substantially, only to come back up over the next few weeks; your conclusion is that superstitious stockholders sell their stock in anticipation of bad luck. What can you do to make use of this information? What effect does your action have? Suppose more people notice the behavior of the stock and react accordingly; what is the effect?
18. Generalize your answer to the previous question to cover other situations where a stock price changes in a predictable way. What does this suggest about schemes to make money by charting stock movements and using the result to predict when the market will go up?
19. Suppose that in Floritania the total cost of bringing up a son is three times the cost of bringing up a daughter, since Floritanians do not believe in educating women. Floritanians simply love grandchildren; every couple wants to have as many as possible. Due to a combination of modern science and ancient witchcraft, Floritanian parents can control the gender of their offspring. What is the male/female ratio in the Floritanian population? Explain.
20. The principal foods of the Floritanians are green eggs and ham. It costs exactly twice as much to produce a pound of green eggs as a pound of ham. The more green eggs that are produced, the lower the price they sell for, and similarly with ham.
 - a. You are producing both green eggs and ham. Green eggs sell for \$3/pound; so does ham. How could you increase your revenue without changing your production cost?
 - b. What will be the result on the prices of green eggs and ham?

c. If everyone acts rationally, what can you say about the eventual prices of green eggs and ham in Floritania?

FOR FURTHER READING

For a good introduction to the economics of genes I recommend Richard Dawkins's *The Selfish Gene* (New York: Oxford University Press, 1976).

A more extensive discussion of the economics of warfare can be found in my essay, "The Economics of War," in J.E. Pournelle (ed.), *Blood and Iron* (New York: Tom Doherty Associates, 1984).

For a very different application of economic analysis to warfare, I recommend Donald W. Engels's *Alexander the Great and the Logistics of the Macedonian Army* (Berkeley: University of California Press, 1978). The author analyzes Alexander's campaigns while omitting all of the battles. His central interest is in the problem of preventing a large army from dying of hunger or thirst and the way in which that problem determined much of Alexander's strategy. Consider, as a very simple example, the fact that you cannot draw water from a well, or 5 wells, or 20 wells, fast enough to keep an army of 100,000 people from dying of thirst.

The relationship between individual rationality and group behavior is analyzed in Thomas Schelling's *Micromotives and Macrobehavior* (New York: W.W. Norton and Co., 1978).

[Next Chapter](#)

[Table of Contents](#)

Chapter 2

How Economists Think

This chapter consists of three parts. The first describes and defends some of the fundamental assumptions and definitions used in economics. The second attempts to demonstrate the importance of price theory, in part by giving examples of economic problems where the obvious answer is wrong and the mistake comes from not having a consistent theory of how prices are determined. The third part briefly describes how, in the next few chapters, we are going to create such a theory.

PART I -- ASSUMPTIONS AND DEFINITIONS

There are a number of features of the economic way of analyzing human behavior that many people find odd or even disturbing. One such feature is the assumption that the different things a person values can all be measured on a single scale, so that even if one thing is much more valuable than another, a sufficiently small amount of the more valuable good is equivalent to some amount of the less valuable. A car, for example, is probably worth much more to you than a bicycle, but a sufficiently small "amount of car" (not a bumper or a headlight but rather the use of a car one day a month, or one chance in a hundred of getting a car) has the same value to you as a whole bicycle--given the choice, you would not care which of them you got.

This sounds plausible enough when we are talking about cars and bicycles, but what about really important things? Does it make sense to say that a human life--as embodied in access to a kidney dialysis machine or the chance to have an essential heart operation--is to be weighed in the same scale as the pleasure of eating a candy bar or watching a television program?

Strange as it may seem, the answer is yes. If we observe how people behave with regard to

their own lives, we find that they are willing to make trade-offs between life and quite minor values. One obvious example is someone who smokes even though he believes that smoking reduces life expectancy. Another is the overweight person who is willing to accept an increased chance of a heart attack in exchange for some number of chocolate sundaes.

Even if you neither smoke nor overeat, you still trade off life against other values. Whenever you cross the street, you are (slightly) increasing your chance of being run over. Every time you spend part of your limited income on something that has no effect on your life expectancy, instead of using it for a medical checkup or to add safety equipment to your car, and every time you choose what to eat on any basis other than what food comes closest to the ideal diet a nutritionist would prescribe, you are choosing to give up, in a probabilistic sense, a little life in exchange for something else.

Those who deny that this is how we do and should behave assume implicitly that there is such a thing as enough medical care, that people should (and wise people do) first buy enough medical care and then devote the rest of their resources to other and infinitely less valuable goals. The economist replies that since additional expenditures on medical care produce benefits well past the point at which one's entire income is spent on it, the concept of "enough" as some absolute amount determined by medical science is meaningless. The proper economic concept of enough medical care is that amount such that the improvement in your health from buying more would be worth less to you than the things you would have to give up to pay for it. You are buying too much medical care if you would be better off (as judged by your own preferences) buying less medical care and spending the money on something else.

I have defined *enough* in terms of money only because the choice you face with regard to the goods and services you buy is whether to give up a dollar's worth of one in exchange for getting another dollar's worth of something else. But market goods and services are only a special case of the general problem of choice. You are buying enough safety when the pleasure you get from running across the street to talk to a friend just balances the value to you of the resulting increase in the chance of getting run over.

So far, I have considered the trade-off between small amounts of life and ordinary amounts of other goods. Perhaps it has occurred to you that we would reach a different conclusion if we considered trading a large amount of life for a (very) large amount of some other good. My argument seems to imply that there should be some price for which

you would be willing to let someone kill you!

There is a good reason why most people would be unwilling to sell their entire life for any amount of money or other goods--they would have no way of collecting. Once they are dead, they cannot spend the money. This is evidence not that life is infinitely valuable but that money has no value to a corpse.

Suppose, however, we offer someone a large sum of money in exchange for his agreeing to be killed in a week. It still seems likely he would refuse. One reason (seen from the economist's standpoint) is that as we increase the amount we consume in a given length of time, the value to us of additional amounts decreases. I am very fond of Baskin-Robbins ice cream cones, but if I were consuming them at a rate of a hundred a week, an additional cone would be worth very little to me. I weigh life and the pleasure of eating ice cream on the same scale, yet no quantity of ice cream I can consume in a week is worth as much to me as the rest of my life. That is why, when I initially defined the idea that everything can be measured on a single scale, I put the definition in terms of a comparison between the value of a given amount of the less valuable good and a sufficiently small amount of the more valuable, instead of comparing a given amount of the more valuable to a sufficiently large amount of the less valuable.

Wants or Needs?

The economist's assumption that all (valued) goods are in this sense comparable shows itself in the use of the term *wants* rather than *needs*. The word *needs* suggests things that are infinitely valuable. You need a certain amount of food, clothing, medical care, or whatever. How much you need could presumably be determined by the appropriate expert and has nothing to do with what such things cost or what your particular values are. This is the typical attitude of the noneconomist, and it is why the economist's way of looking at things often seems unrealistic and even ugly. The economist replies that how much of each of these things you will, and should, choose to have depends on how much you value them, how much you value other things you must give up to get them, and how much of such other things you must give up to get a given amount of clothing, medical care, or whatever. Your choices depend, in other words, on your tastes and on the costs to you of the alternative things that you desire.

One reply the noneconomist (perhaps I mean the antieconomist) might make is that we ought to have enough of everything. If you have enough movies and enough ice cream cones and enough of everything else you desire, you no longer have any reason to choose less medical care or nutrition in order to get more of something else (although combining good nutrition with enough ice cream cones could be a problem for some of us). Perhaps our objective should be a society where everybody has enough. Perhaps, it is sometimes argued, the marvels of modern technology, combined with the right economic system, could bring such a society within our reach, making the problems of choosing among different values obsolete.

This particular argument was more popular 20 years ago than it is now. Currently the fashion has changed and we are being told that limitations in natural resources (and in the ability of the environment to absorb our wastes) impose stringent limitations on how much of everything we can have. Yet even if that is not true, even if (as I suspect) resource limits are no more binding now than in the past, "enough of everything" is still not a reasonable goal. Why?

It is often assumed that if we could only produce somewhat more than we do, we would have everything we want. In order to consume still more, we would each have to drive three cars and eat six meals a day. This argument confuses increasing the value of what you consume with increasing the amount you consume. A modern stereo is no bigger and consumes no more power than its predecessor of 30 years ago, yet moving from one to the other represents an increase in "consumption." I have no use for three cars, but I would like a car three times as good as the one I now have. There are many ways in which my life could be improved if I consumed things that are more costly to create but no larger than those I now have. My desire for pounds of food is already satiated and my desire for number of cars could be satiated with a moderate increase in my income, but my desire for quality of food or quality of car would remain even at a much higher income, and my desire for more of *something* would remain unsatiated as long as I remained alive and conscious under any circumstances I can imagine.

From both introspection and conversation, I have formulated a general law on this subject. Everyone feels that there is a level of income above which all consumption is frivolous. For everyone, that level is about twice his own. An Indian peasant living on \$500/year believes that if only he had \$1,000/year, he would have everything he could want with a little left over. An American physician living on \$50,000/year (after taxes) doubts that anyone has any real use for more than \$100,000/year.

Both the peasant and the physician are wrong, but both opinions are the result of rational behavior by those who hold them. Whether you are living on \$500/year or \$50,000/year, the consumption decisions you make, the goods you consider buying, are those appropriate to such an income. Heaven would be a place where you had all the things you have considered buying and decided not to. There is little point wasting your time learning or thinking about consumption goods that cost ten times your yearly income, so the possession of such goods is not part of your picture of the good life.

Value

So far I have discussed, and tried to defend, two of the assumptions that go into economics: *comparability*, the assumption that the different things we value are comparable, and *non-satiation*, the assumption that in any plausible society, present or future, we cannot all have everything we want and must give up some things we desire in order to have others. In talking about value, I have also implicitly introduced an important definition--that *value* (of things) means how much we value them and that how much we value them is properly estimated not by our words but by our actions. In discussing the trade-off between the value of life and the value of the pleasure of smoking, my evidence that the two are comparable was that people choose to smoke, even though they believe doing so lowers their life expectancy. This definition is called the *principle of revealed preference*--meaning that your preferences are revealed by your actions.

The first part of the definition of value embodied in the principle of revealed preference might be questioned by those who prefer to base value on some external criterion--what we should want or what is good for us. The second might be questioned by those who believe that their values are not fairly reflected in their actions, that they value health and life but just cannot resist one more cigarette. But economics is supposed to describe how people act, and we are therefore concerned with value as it relates to action. A smoker's statement that he puts infinite value on his own life may help to explain what he believes, but it is less useful for understanding what he will do than is the kind of value expressed when he takes a cigarette out and lights it.

Even if revealed preference is a useful concept for our purpose, should we call what it reveals value? Does not the word carry with it an implication of something beyond mere individual preference? That is a philosophical question that goes beyond the subject of this

book. If using the word *value* to refer equally to a crust of bread in the hands of a starving man and a syringe of heroin in the hands of an addict makes you uncomfortable, then substitute *economic value* instead. But remember that the addition of "economic" does not mean "having monetary value," "being material," "capable of producing profit for someone," or anything similar. Economic value is simply value to individuals as judged by them and revealed in their actions.

Economics Joke #1: *Two economists walked past a Porsche showroom. One of them pointed at a shiny car in the window and said, "I want that." "Obviously not," the other replied.*

Choice or Necessity?

The difference between the approaches to human behavior taken by economists and by noneconomists comes in part from the economist's assumptions of comparability and insatiability, in part from the definition of value in terms of revealed preference, and in part from the fundamental assumption of rationality that I made and defended in the previous chapter. One form in which the difference often appears is the economist's insistence that virtually all human behavior should be described in terms of choices. To many noneconomists, this seems deceptive. What, after all, is the point of saying that you choose not to buy something you cannot afford?

When you say that you cannot afford something, you usually mean only that there are other things you would rather spend the money on. Most of us would say that we could not afford a \$1,000 shirt. Yet most of us could save up \$1,000 in a year if it were sufficiently important--important enough that you were willing to spend only a dollar a day on food (roughly the cost of the least expensive full-nutrition diet--powdered milk, soy beans, and the like), share a one-room apartment with two roommates, and buy your clothing from Goodwill.

Consider an even more extreme case, in which you have assets of only a few hundred dollars and there is something enormously valuable to you that costs \$100,000 and will only be available for the next month. In a month, you surely cannot earn that much money. It seems reasonable, in this case at least, to say that you cannot afford it. Yet even here, there is a legitimate sense in which what you really mean is that you do not want it.

Suppose the object were so valuable that getting it made your life wonderful forever after and failing to get it meant instant death. If you could not earn, borrow, or steal \$100,000, the sensible thing to do would be to get as much money as possible, go to Reno or Las Vegas, work out a series of bets that would maximize your chance of converting what you had into exactly \$100,000, and make them. If you are not prepared to do that, then the reason you do not buy the object is not that you cannot afford its \$100,000 price. It is that you do not want it--enough.

In part, the claim that people do not really have any choice confuses the lack of alternatives with the lack of attractive or desirable ones. Having chosen the best alternative, you may say that you had little choice; in a sense you are correct. There may be only one best alternative.

One example of this confusion that I find particularly disturbing is the argument that the poor should be "given" essential services by government even if (as is often the case) they end up having to pay for the services themselves through increased taxes. Poor people, it is said, do not really choose not to go to doctors--they simply cannot afford to. Therefore a benevolent government should take money from the poor and use it to provide the medical services they need.

If this argument seems convincing, try translating it into the language of choice. Poor people choose not to go to doctors because to do so they would have to give up things still more important to them--food, perhaps, or heat. It may sound heartless to say that someone chooses not to go to a doctor when he can do so only at the cost of starving to death, but putting it that way at least reminds us that if you "help" him by forcing him to spend his money on doctors, you are compelling him to make a choice--starvation--that he rejected because it was even worse than the alternative--no medical care--that he chose.

The question of how much choice individuals really have reappears on a larger scale in discussions of how flexible the economy as a whole is--to what extent it can vary the amount of the different resources it uses. Our tendency is to look at the way things are now being done and assume that that way is the only possible one. But the way things are now done is the solution to a particular problem--producing goods as cheaply as possible given the present cost of various inputs. If some input--unskilled labor, say, or energy or some raw material--were much more or less expensive, the optimal way of producing would change.

A familiar example is the consumption of gasoline. If you suggest to someone that if gasoline were more expensive he would use less of it, his initial response is that using less gasoline would mean giving up the job he commutes to or walking two miles each way to do his shopping. Indeed, when oil prices shot up in the early 1970's, many people argued that Americans would continue to use as much gasoline as before at virtually any price, unless the government forced them to do otherwise.

There are many ways to save gasoline. Car pooling and driving more slowly are obvious ones. Buying lighter cars is less obvious. Workers choosing to live closer to their jobs or employers choosing to locate factories nearer to their workers are still less obvious. Petroleum is used to produce both gasoline and heating oil; the refiners can, to a considerable degree, control how much of each is produced. One way of "saving" gasoline is to use less heating oil and make a larger fraction of the petroleum into gasoline instead. Insulation, smaller houses, and moving south are all ways of saving gasoline.

PART 2 -- PRICE THEORY--WHY IT MATTERS

This book has two purposes--to teach you to think like an economist and to teach you the set of ideas that lie at the core of economic theory as it now exists. That set of ideas is *price theory*--the explanation of how relative prices are determined and how prices function to coordinate economic activity.

There are at least two reasons to want to understand price theory (aside from passing this course). The first is to make some sense out of the world you live in. You are in the middle of a very highly organized system with nobody organizing it. The items you use and see, even very simple objects such as a pen or pencil, were each produced by the coordinated activity of millions of people. Someone had to cut down the tree to make the pencil. Someone had to season the wood and cut it to shape. Someone had to make the tools to cut down the trees and the tools to make the tools and the fuel for the tools and the refineries to make the fuel. While small parts of this immense enterprise are under centralized control (one firm organizes the cutting and seasoning of the wood, another actually assembles the pencil), nobody coordinates the overall enterprise.

Someone who had visited China told me about a conversation with an official in the ministry of materials supply. The official was planning to visit the United States in order

to see how things were done there. He wanted, naturally enough, to meet and speak with his opposite number--whoever was in charge of seeing that U.S. producers got the materials they needed in order to produce. He had difficulty understanding the answer--that no such person exists.

A market economy is coordinated through the price system. Costs of production--ultimately, the cost to a worker of working instead of taking a vacation or of working at one job instead of at another, or the cost of using land or some other resource for one purpose and so being unable to use it for another--are reflected in the prices for which goods are sold. The value of goods to those who ultimately consume them is reflected in the prices purchasers are willing to pay. If a good is worth more to a consumer than it costs to produce, it gets produced; if not, it does not.

If new uses for copper increase demand, that bids up the price, so existing users find it in their interest to use less. If supply decreases--a mine runs out or a harvest fails--the same thing happens. Prices provide an intricate system of signals and incentives to coordinate the activities of several million firms and several billion individuals. How this is done you will learn in the next few months.

Four Wrong Answers

The first reason to understand price theory is to understand how the society around you works. The second reason is that an understanding of how prices are determined is essential to an understanding of most controversial economic issues while a misunderstanding of how prices are determined is at the root of many, if not most, economic errors. Consider the following four examples of cases where the obvious answer is wrong and where the error is an implicit (wrong) assumption about price theory. I shall not prove what the right answer is, although I shall give you some hints about where the counterintuitive result comes from.

Rental Contracts. Tenants rent apartments from landlords. Cities often have laws restricting what lease agreements are legal. For example, the law may require the landlord to give the tenant three months' notice before evicting him, even if the lease provides for a shorter term.

It seems obvious that the effect of such a law is to benefit tenants and hurt landlords. That may be true for those tenants who have already signed leases when the law goes into effect. For most other tenants, it is false. The law either has no effect or it injures both tenants and landlords (on average; there may be particular tenants, or particular landlords, who benefit).

The reason most people expect such a law to benefit tenants is that they have, without realizing it, assumed that the law does not affect how much rent the tenant must pay. If you are paying the same rent and have a more favorable lease, you are better off. But this assumption is implausible. Although the law says nothing about rents, it indirectly affects both the operating costs of landlords (they are higher, since it is harder to get rid of bad tenants) and the attractiveness of the lease to tenants (who are now guaranteed three months' notice). With both supply and demand conditions for rental housing changed, you can hardly expect the market rent to remain the same--any more than you would expect the market price of cars to be unaffected by a law that forced the manufacturers to produce cars that were more costly to build and more desirable to buy. It turns out that either the law has no effect at all (the landlords would have chosen to offer the guarantee anyway in order to attract tenants and so be able to get more rent) or it injures both parties (the advantage of greater security does not compensate the average tenant for the resulting increase in his rent). I am asserting this, not proving it; the argument will be worked out in detail in Chapter 7.

Popcorn Prices. The second counterintuitive result concerns popcorn. Movie theaters normally sell popcorn (and candy and sodas) for substantially higher prices than they are sold for elsewhere. There is an obvious explanation--the movie theater has a captive audience. While it is obvious, it is also wrong. Assuming that both customers and theater owners are rational, a straightforward economic argument can be constructed to show that selling food at above-cost prices lowers the net income of the theater owner. Explaining the observed prices requires a more complicated argument.

Here again, the error is in assuming that a price--this time the price the theater can get for a ticket--is fixed, when it will in fact depend on the characteristics of what is being sold, including, in this case, how much the theater charges for food. If that does not seem plausible to you, imagine that instead of exploiting its captive market with high food prices, the theater exploits it by charging an additional dollar per customer for seat rental. Just as the customers have nowhere else to buy their popcorn so they have nowhere else to rent seats in the movie theater. If the price the theater can sell tickets for is unaffected by

the price of popcorn, why should it be affected by the availability or price of other amenities--such as seats?

Obviously the conclusion is absurd. The theater charges the ticket price it does because any increase costs it more in lost customers than it gains from the higher price per ticket. Since an additional fee for seats is equivalent to raising the ticket price (unless customers are willing to watch the movie while standing), it will lower, not raise, the theater's profits.

The effect of raising popcorn prices is more complicated than the effect of renting seats, since it is easier to vary the amount of popcorn you eat according to its price than to vary the number of seats you sit in; we will return to the question of why popcorn in theaters is expensive in later chapters. But the error in the obvious explanation of expensive popcorn--assuming the price at which tickets can be sold is unaffected by changes in the quality of the product--is the same.

Why Price Control Makes Gasoline More Expensive. A third counterintuitive result is that although price control on gasoline lowers the price consumers pay for gasoline in dollars per gallon, it raises the cost to consumers of getting gasoline, where the cost includes both the price and nonmonetary costs such as time spent waiting in line.

To see why this is true, imagine that the uncontrolled price is \$1/ gallon. At that price, producers produce exactly as much gasoline as consumers want to consume (which is why it is the market price). The government imposes a maximum price of \$0.80/gallon. As a first step in the argument, assume producers continue producing the same quantity as before. At the lower price, consumers want to consume more. But you cannot consume gasoline that is not produced, so stations start running out. Consumers start coming to the stations earlier in the day, just after the stations have received their consignments of gasoline. But although this may enable one driver to get gasoline instead of another, it still does not allow drivers as a group to consume more than is produced, so the stations still run out. As everyone tries to be first, lines start to form. The cost of gasoline is now a cost in money plus a nonmonetary cost--waiting time (plus getting up early to go to the gas station); you can think of the latter as equivalent, from the consumer's standpoint, to an additional sum of money. As long as the money equivalent of the nonmonetary cost is less than \$0.20, the total cost per gallon (waiting time plus money) is less than \$1/ gallon. Consumers still want to consume more than is being produced (remember that \$1 /gallon was the market price at which quantity demanded and quantity supplied were equal), and the lines continue to grow. Only when the cost--time plus money--reaches the old price are

we back in a situation where the amount of gasoline that consumers want to buy is equal to the amount being produced.

So far, we have assumed that the producers produce the same amount of gasoline when they are receiving \$0.80/gallon as when they are receiving \$1/gallon. That is unlikely. At the lower price, producers produce less--marginal oil wells close down, older and more inefficient refineries go out of use, and so on. Since less is being produced than at a price of \$1/gallon, consumers are still trying to consume more than is being produced even when the cost to them (price plus time) reaches \$1/gallon; the lines have to grow still longer, making the cost even higher, before quantity demanded is reduced to quantity supplied. So price control raises the cost of gasoline. In Chapter 17, this analysis will be applied in more detail to price control under a variety of arrangements.

Improved Light Bulbs. The final example concerns light bulbs. It is sometimes argued that if a company with a monopoly of light bulbs invents a new bulb that lasts ten times as long as the old kind, the company will be better off suppressing the invention. After all, it is said, if the new bulb is introduced, the company can only sell one tenth as many bulbs as before, so its revenue and profit will be one tenth as great.

The mistake in this reasoning is the assumption that the company will sell the new bulb, if introduced, at the same price as the old. If consumers were willing to buy the old light bulbs for \$1 each, they should be willing to buy the new ones for about \$10 each. What they are really buying, after all, are light bulb hours, which are at the same price as before. If the company sells one tenth as many bulbs at ten times the price, its revenue is the same as before. Unless the new bulb costs at least ten times as much to produce as the old, costs are less than before and profits therefore are higher. It is worth introducing the new bulb.

In all of these cases, when I say something is true on average, what I mean is that it is strictly true if all consumers are identical to each other and all producers are identical to each other. This is often a useful approximation if you wish to distinguish distributional effects within a group from distributional effects between groups.

Naive Price Theory

All of these examples have one element in common. In each case, the mistake is in

assuming that one part of a system will stay the same when another part is changed. In three of the four cases, what is assumed to stay the same is a price. I like to describe this mistake as naive price theory--the theory that the only thing determining tomorrow's price is today's price. Naive price theory is a perfectly natural way of dealing with prices--if you do not understand what determines them. In each of the three cases--theater tickets, light bulbs, and apartments--we were considering a change in something other than price. In each case, a reader unfamiliar with economics might argue that since I said nothing about the price changing when the problem was stated, he assumed it stayed the same.

If that seems like a reasonable defense of naive price theory, consider the following analogy. I visit a friend whose month-old baby is sleeping in a small crib. I ask him whether he plans to buy a larger crib or a bed when the child gets older. He looks puzzled and asks me what is wrong with the crib the child is sleeping in now. I point out that when the child gets a little bigger, the crib will be too small for him. My friend replies that I had asked what he planned to do when the child got older--not bigger.

It makes very little sense to assume that as a baby grows older he remains the same size. It makes no more sense to assume that the market price of a good remains the same when you change its cost of production, its value to potential purchasers, or both. In each case, the assumption "If you did not say it was going to change, it probably stays the same" ceases to make sense once you understand the causal relations involved. That is what is wrong with naive price theory.

Why, you may ask, do I dignify this error by calling it a price theory? I do so in order to point out that in each of these cases, and many more, the alternative to correct economic theory is not doing without theory (sometimes referred to as just using common sense). The alternative to correct theory is incorrect theory. In order to analyze the effect of introducing longer lasting light bulbs (or the other cases I have just discussed), you must, explicitly or implicitly, assume something about the effect on the price; you do not avoid doing so by assuming that there is no effect.

PART 3 -- THE BIG PICTURE, OR HOW TO SOLVE A HARD PROBLEM

In order to understand how prices are determined, we must somehow untangle a complicated, intricately interrelated problem. How much of a good a consumer chooses to

consume depends both on the total resources available to him--his income--and, as the earlier discussion suggested, on how much of other things he must give up to get that good--in other words, on how much it costs. How much it costs depends, among other things, on how much he consumes, since his demand affects what producers can sell it for. How much producers sell and at what price will affect how much labor (and other productive resources) they choose to buy, and at what price. Since consumers get their income by selling their labor (and other productive resources they own), this will in turn affect the income of the consumers, bringing us full circle. It seems as though we cannot solve any one part of the problem until we have first solved the rest.

The solution is to break the problem into smaller pieces, solve each piece in a way sufficiently general that it can be combined with whatever the solutions of the other pieces turn out to be, then reassemble the whole in such a way that all of the solutions are consistent with each other. First, in Chapters 3 and 4, we consider a consumer with either a given income or a given endowment of goods, confronted with a market and a set of prices, and analyze his behavior. Next, in Chapter 5, we consider a producer producing either for his own consumption or for sale; the producer can transform his labor into goods and either consume them or sell them on the market. In Chapter 6, we consider trade among individuals, mostly in the context of a two-person (or two-country) world. In Chapter 7, we put together the material of Chapters 3, 4 and 5, showing how the interaction of (many) consumers who wish to buy goods and (many) producers who wish to sell them produces market prices. Finally, in Chapter 8, we close the circle, combining the results of the previous five chapters to recreate the whole interacting system.

What we will be analyzing, in this section of the book, is a very simple economy. Production and consumption are by and for individuals; there are no firms. The world is predictable and static; complications of change and uncertainty are assumed away. Once we understand the logic of that simple economy, we will be ready to put back into it, one after another, the complications initially left out.

PROBLEMS

1. Give examples of ways in which you yourself make trade-offs between your life and relatively minor values; they should not be examples given in the chapter.

2. Suppose we were talking not about what people *do* value but about what they *should* value. Do you think comparability would still hold? Discuss. If your answer is no, give examples of incomparable values.
3. State the principle of revealed preference in your own words. Give an example, in your own or your friends' behavior, where stated values are different from the values deduced from revealed preference.
4. Life is not the only thing that is said to be beyond price. Other examples are health, love, salvation, and the welfare of our contry. Give examples, for yourself and others, of ways in which (small amounts of) such "priceless" things are routinely given up in exchange for minor values.
5. Figure 2-1 shows how the total pleasure I get from eating ice cream cones varies with the number of ice cream cones I eat each week. Figure 2-2 shows how the total pleasure I get from all the goods I buy varies with the number of dollars worth of goods I buy each week. Discuss and explain the similarities and the differences in the two figures.
6. Describe some likely short-run and long-run adjustments that people would make to each of the following changes. Assume in each case that the change is permanent, reflecting some underlying change in technology, resource costs, or the like.
 - a. Large chunks of the country fall into the Atlantic and Pacific oceans; land prices go up tenfold.
 - b. Electricity prices go up tenfold.
 - c. All heating costs triple.
 - d. The government imposes a \$20,000 baby tax for every baby born in the United States.
 - e. Solar power satellites start beaming energy down to earth; electricity prices go down by a factor of 100.
 - f. Due to extensive immigration, hard working (but unskilled) workers are readily available for a dollar/hour.
7. You are an economist, you have a child, and you decide you should make him wash out

his mouth with soap whenever he uses a bad word or phrase. The first forbidden word on your list is "need." What would other words or phrases be? If possible, give examples that have not been discussed in the chapter.

8. The chapter describes how millions of people cooperate to produce a pencil. Describe how you or someone you know is involved in producing a pencil. A computer. An atomic bomb. The examples should be real ones.

9. Which of the principles discussed in the chapter did the Porsche joke illustrate? Explain.

10. It is part of American folklore that from time to time some genius invents a razor blade that lasts forever or a car engine that runs on water, only to have his invention bought up and suppressed by companies that want to continue making money selling razor blades or gasoline. Does this seem plausible? Discuss.

11. I recently received a letter from a credit card company (call it ACCCo.) urging me to support a law that would make it illegal for merchants to charge a higher price to customers who used credit cards. Such a law currently exists but is about to expire. The letter argues as follows:

To begin with, present law already permits merchants to offer discounts to customers who choose to pay with cash. Such discounts can benefit customers--and we have long been for them. They allow you to either pay the regular price and have the convenience of using your credit card, or pay cash and receive a discount.

We think you and all consumers should have this freedom of choice. It is a choice with no penalty and numerous benefits.

A credit card surcharge, however, is entirely different. It would penalize you whether you used cash or a credit card. If you paid cash, you would be charged the full price. If you wanted--or needed--to use your credit card, you would be charged a penalty over and above the regular price.

a. Is the distinction made by ACCCo. between permitting cash discounts and permitting a surcharge for use of a credit card a legitimate one? Discuss.

b. ACCCo. apparently believes that it is in its interest to have credit card surcharges prohibited (how do I know that?). Is it obvious that it is right? From the standpoint of credit card companies, what are the *advantages* of permitting such surcharges?


12. While negotiating with a firm that wished to publish this book, I got into a conversation on the subject of the secondhand market for textbooks, The editor I was talking with complained that sales of a textbook typically drop sharply in the second year because students buy secondhand copies from other students who bought the books new the year before. While she had no suggestions for eliminating the secondhand market, she clearly regarded it as a bad thing.

I put the following question to her and her colleagues. Suppose an inventor walks in your door with a new product--timed ink. Print your books in timed ink and activate it when the books leave the warehouse. At the end of the school year, the pages will go blank. Students can no longer buy second-hand textbooks. Do your profits go up--or down?


To make the problem more specific, assume that presently textbooks are sold for \$30 each, that students resell them to other students for \$15, and that each textbook costs the publisher \$20 to produce and lasts exactly two years. Discuss.

13. "We should require every barber to have a year of training and to pass an exam. The barbers would be a little worse off, since they would have to be trained, but the rest of us would obviously be better off, since our hair would be cut better."

Discuss. Is the last sentence of the quote true?



Total pleasure per week from eating ice cream cones, as a function of the rate at which they are eaten.



Total pleasure per week from all consumption, as a function of weekly expenditure.

[Next Chapter](#)

[Table of Contents](#)

[Previous Chapter](#)

Section II

Price = Value = Cost:

Competitive Equilibrium in a Simple Economy

Chapter 3

The Consumer: Choice and Indifference Curves

PRICE, COST, AND VALUE

A very old puzzle in economics is the relation between price, value to the consumer, and cost of production. It is tempting to say that the price of a good is determined by its value to the user. Why, after all, would anyone buy a good for more or sell it for less? But if this is so, why are diamonds, which are relatively unimportant (most of us could get along quite well if they did not exist), worth so much more per pound than water, which is essential for life? If the answer is that diamonds are rare and that it is rarity rather than usefulness that determines price, I reply that signatures of mine written in yellow ink are even rarer than original autographs of Abraham Lincoln but bring a (much) lower price.

Perhaps it is cost of production that determines price. When I was very young, I used to amuse myself by shooting stalks of grass with a BB gun. That is a costly way of mowing the lawn, even considering that the cost per hour of a nine-year-old's time is not very high. I think it unlikely that anyone would pay a correspondingly high price to have his lawn mowed in that fashion.

This puzzle--the relation between value to the consumer, cost of production, and price--was solved a little over 100 years ago. The answer is that price equals both cost of production and value to the user, both of which must therefore be equal to each other.

How market mechanisms arrange that triple equality will be discussed in the next few chapters. In this chapter and the next, we shall analyze the behavior of a consumer who must decide what to buy with his limited resources; among the things we shall learn in the process is why, as a consequence of rational behavior by the consumer, price equals (marginal) value.

LANGUAGES

There are several different languages in which the problem of consumer behavior--and many other problems in economics--can be analyzed. Each of these languages has advantages and disadvantages. One may use the language of calculus, making assumptions about the form of the "utility function" that describes the individual's preferences among different goods and deducing the characteristics of the bundle of goods that maximizes it. This has the advantage of allowing compact and rigorous mathematical arguments and of producing very general results, applicable to a wide range of possible situations. It has the disadvantage that even if you know calculus, you probably do not know it in the same sense in which you know English. Unless you are very good at intuiting mathematics, you can follow a proof step by step from assumptions to conclusions and still not know why the result is true. For these reasons, calculus and utility functions will be used only in some of the optional sections of this text. The ordinary sections will not assume any knowledge of either, although a few concepts borrowed from calculus will be explained in simple terms and used where necessary.

Another possible language is geometry. Most of us can understand abstract relations better as pictures than as equations; hence geometric arguments are easier to intuit. One disadvantage of geometry is that it limits us to situations that can be drawn in two dimensions--typically, for example, to choices involving only two different goods. A second disadvantage is that we may, in drawing the picture, inadvertently build into it assumptions about the problem--possibly false ones.

The third language is English. While not as good as mathematical languages for expressing precise quantitative relations, English has the advantage of being, for most of us, our native tongue. Insofar as we think in words at all, it is the language we are used to thinking in. Unless we have very good mathematical intuition, all mathematical arguments eventually get translated, in our heads, into words, and it is only then that we really

understand them. Alfred Marshall, possibly the most important economist of the past century, wrote that economic ideas should be worked out and proved in mathematical form and then translated into words; if you find that you cannot put your analysis into words, you should burn your mathematics. Since it is often hard to keep track of quantitative relations in a verbal argument, explanations given in English will frequently be supplemented by tables.

This chapter presents the logic of a consumer, first in verbal form, then in a simple geometrical form suitable for describing the choice between two different goods. The analysis is continued in the next chapter, which uses a somewhat more complicated geometric argument, designed to produce calculus results without actually using calculus. Among the results are the answers to three interesting questions: How does the amount you buy depend on price? How much do you benefit by being able to buy something at a particular price? What is the relation between price and value?

THE CONSUMER I: ENGLISH VERSION

Your problem as a consumer is to choose among the various bundles of goods and services you could purchase or produce with your limited resources of time and money. There are two elements to the problem--your preferences and your opportunity set. Your preferences could be represented by a gigantic table showing all possible *bundles*--collections of goods and services that you could conceivably consume--and showing for every pair of bundles which one you prefer. We assume that your preferences are consistent; if you prefer A to B and B to C, you also prefer A to C. Your *opportunity set* can be thought of as a list containing every bundle that you have enough money to buy. Your problem as a consumer is to decide which of the bundles in your opportunity set you prefer.

I will simplify the problem in most of this chapter by considering only two goods at a time--in this part of the chapter, apples and oranges. We may imagine either that these are the only goods that exist or else that you have already decided how much of everything else to consume. We assume that both apples and oranges are *goods*, meaning that you prefer more to less. Things that are not goods are either *neutral* (you do not care how much you have) or *bads* (you prefer less to more: garbage, strawberry ice cream, acid rock). As these examples suggest, whether something is a good or a bad for you depends

on your preferences; some people like strawberry ice cream.

Preferences: Patterns on a Table

Table 3-1

Bundle	Apples	Oranges	Utility	Bundle	Apples	Oranges	Utility
A	10	0	5	F	2	8	5
B	7	1	5	G	10	1	6
C	5	2	5	H	8	2	6
D	4	3	5	J	7	3	6
E	3	5	5	K	9	1	?
				L	7	5	?

Table 3-1 is a list of bundles of apples and oranges. For each bundle, the table shows its name (A-L), how many apples and oranges it contains, and its *utility*--an abstract measure of how much you value the bundle. The statement "Bundle A and bundle C have the same utility" is equivalent to the statement "Given a choice between A and C, you would not care which you got." The statement "Bundle G has greater utility than bundle B" is equivalent to "Given a choice between B and G, you would choose G." Listing a utility for each bundle is a simple way of describing your preferences; by comparing the utilities of two bundles, we can see which you prefer.

Utility is being used here as an *ordinal* measure--the order matters (bundle G has more utility than bundle F, so you prefer G to F) but the amount does not. In Chapter 13, we will expand the idea of utility in a way that converts it into a *cardinal* measure--one for which both order (bundle G has more utility than bundle F) and size (bundle G has 1 utile more utility than bundle F) matter. Since in this chapter, utility describes the choices of one individual, we need not worry about *interpersonal utility comparisons*--questions such as "Does an orange have more utility to me or to you?" We will say a little more about that question in Chapter 15, when we are trying to evaluate changes that make some people better off and some worse off.

Since bundles A-F have the same utility (5), you are indifferent among them. If you started with 4 apples and 3 oranges (D) and somehow gained an apple and lost an orange, moving from D to C, you would be neither better nor worse off. We will say that in such a situation an apple and an orange have the same value to you, or alternatively, that the value of an apple is 1 orange; the value of an orange is 1 apple.

It is important to note that the statement "The value of an apple is 1 orange" is true only between C and D. As we move up or down the table, values change. If you start with 5 apples and 2 oranges, you must receive not 1 but 2 apples to make up for losing 1 orange; in this situation (between B and C), the value of an orange is equal to that of 2 apples. An orange is worth 2 apples; an apple is worth half an orange.

The numbers in bundles A-F follow a pattern--as you move up the table, it takes more and more apples to equal 1 orange; as you move down, it takes more and more oranges to equal 1 apple. The numbers are set up this way because, as a rule, the more you are consuming of something, the less you value consuming one more. If you have very few oranges, you will be willing to give up a good deal to have one more (assuming you like oranges). If you are already consuming 12 oranges per day, you will be willing to give up very little to have 13 instead. As we move up the column, each successive bundle has fewer oranges and more apples, so in each successive case oranges are worth more to you and apples less, making each orange worth more apples. This general pattern is referred to as a declining *marginal rate of substitution* --the rate at which additional apples substitute for additional oranges declines with increases in the number of apples or decreases in the number of oranges.

Another way of seeing the pattern is to ask how many oranges it takes to raise your utility by 1. If you start at A, the answer is that it takes 1 orange; adding 1 orange to your bundle puts you at G, with a utility of 6, up 1 from 5. If you start at B, it takes 2 extra oranges to move you to J, increasing your utility by 1. At B you already have 1 orange, so the extra utility you get from an additional orange (the *marginal utility* of an orange) is less than at A, where you start with no oranges at all.

The name for this pattern is the *principle of declining marginal utility*--marginal utility because what is declining as you have more and more oranges is the additional utility to you of having one more orange. It is the same principle that was introduced in the previous chapter when I discussed why I would not trade my life for any quantity of Baskin-Robbins ice cream cones. Figure 2-1 showed that as the rate at which I consumed

ice cream cones increased, the additional utility from each additional cone became less and less. Eventually I reached a rate of consumption at which increased consumption resulted in decreased utility--the additional utility from additional ice cream cones was negative.

Trading toward an Optimal Bundle

Suppose you start with bundle A on Table 3-1, and someone offers to trade oranges for your apples at a rate of 1 for 1. You accept the offer, and trade 1 apple for 1 orange. That gives you bundle K. Since K has more apples than B and as many oranges, you prefer K to B; since B is equivalent to A, you prefer K to A. We do not know what K's utility is, but it must be more than 5 (and less than 6. Why?).

To figure out how many apples you would be willing to exchange for oranges at a rate of 1 for 1, we would need to add many more bundles to the table. That problem is more easily solved using the geometric approach, which will be introduced in the next part of the chapter. There are, however, a number of lessons that can be drawn from this rather simple analysis of consumer choice.

The first is that the value of something is whatever we are (just) willing to give up for it. Two things have the same value if gaining one and losing the other leaves us neither better nor worse off--meaning that we are indifferent between the situation before the exchange and the situation after the exchange. This is an application of the principle of revealed preference discussed in the previous chapter--our values are defined by the choices we make.

A second lesson is that the value of goods (to you) depends not only on the nature of the goods and your preferences but also on how much of those goods you have. If you have 1 apple and 12 oranges, an orange will be worth very little (in apples). If you have 10 apples and no oranges, an orange is worth quite a lot of apples--3, according to Table 3-1.

The third lesson is that the price (or cost) of a good is the amount of something else you must give up to get it. In our example, where someone is willing to trade oranges for apples at a rate of 1 for 1, the price of an apple is 1 orange and the price of an orange is 1 apple. This is called *opportunity cost*--the cost of getting one thing, whether by buying it

or producing it, is what you have to give up in order to get it. The cost of an A on a midterm, for example, may turn out to be three parties, two football games, and a night's sleep. The cost of living in a house that you already own is not, as you might think, limited to expenditures on taxes, maintenance, and the like; it also includes the interest you could collect on the money you would have if you sold the house to someone else instead of living in it yourself.

Opportunity cost is not a particular kind of cost but rather the correct way of looking at all costs. The money you spend to buy something is a cost only because there are other things you would like to spend the money on instead; by buying A, you give up the opportunity to buy B. Not getting the most valuable of the B's that you could have bought with the money--the one you would have bought if A had not been available--is then the cost to you of buying A. That is why, if you were certain that the world was going to end at midnight today, money would become almost worthless to you. Its only use would be to be spent today--so you would "spend as if there were no tomorrow."

The final lesson is that you buy something if and only if its cost is less than its value. In the example we gave, the cost of an orange was 1 apple. The value of an orange, between bundles A and B, was 3 apples. So you bought it. That put you at bundle K. If, starting from there, the value of an orange was still more than 1 apple, you would have bought another. As you trade apples for oranges, the number of apples you have decreases and the number of oranges increases. Because of the principle of declining marginal utility, additional oranges become less valuable and apples become more valuable, so the value of (one more) orange measured in apples falls. When, as a result of trading, you reach a bundle for which the value of yet another orange is no more than its price, you stop trading; you have reached the best possible bundle, given your initial situation (bundle A) and the price at which you can trade apples for oranges.

So far, I have only considered trading (and valuing) whole apples and oranges. As long as we limit ourselves in this way, concepts such as the value of an apple are somewhat ambiguous. If you have 4 apples and 3 oranges, is the value of an apple the number of oranges you would give up in order to get 1 more apple (1 orange) or the number of oranges you would accept in exchange for having 1 fewer apple (2 oranges)? This ambiguity disappears if we consider trading very small amounts of the two goods

If this sounds messy with apples and oranges, substitute apple juice and orange juice. If we move from four quarts of apple juice either up or down by, say, a teaspoon, the value

to us of apple juice changes only very slightly, and similarly with the value of orange juice, so the rate at which we are just willing to exchange apple juice for orange juice should be almost exactly the same whether we are giving up a little apple juice in exchange for a little orange juice or giving up a little orange juice in exchange for a little apple juice. This is the sort of relation that is hard to put into words. It should become a little clearer in the next section, where the same argument is repeated in a geometric form, and clearer still to those of you familiar with calculus.

THE CONSUMER II: GEOMETRY AND INDIFFERENCE CURVES

Figure 3-1 shows another way of describing the preferences shown in Table 3-1. The horizontal axis represents apples; the vertical axis represents oranges. Instead of showing utility, we show *indifference curves* U_a , U_b , and U_c . Each indifference curve connects a set of bundles that have the same utility--bundles among which the consumer is indifferent. Higher indifference curves represent preferred bundles. Note, for instance, that point H on U_b is a bundle containing more apples and more oranges than point B on U_a . Since we have assumed that apples and oranges are both goods (you would rather have more than less), you prefer H to B. Since all bundles on U_a are equivalent to B (by the definition of an indifference curve) and all bundles on U_b are equivalent to H (ditto), any bundle on U_b is preferred to any bundle on U_a . Similarly any bundle on U_c is preferred to any bundle on either of the other two indifference curves. This conclusion depends only on assuming that apples and oranges are goods; it does not require us to know the actual utilities of the different bundles.

A table such as Table 3-1 can show only a finite number of bundles; one of the advantages of the geometric approach is that one indifference curve contains an infinite number of points, representing an infinite number of different bundles. Another advantage is that looking at the blank space between the indifference curves shown on a figure such as Figure 3-1 reminds us that the indifference curves we draw, or the bundles on a table such as Table 3-1, are only a tiny selection from an infinite set. Any point on the figure, such as J, K, or L, is a bundle of goods--so many apples, so many oranges--a bundle you prefer to those on the indifference curves below it and to which you prefer those on the indifference curves above it. Through any such point, you could draw a new indifference curve containing all the bundles you regard as equivalent to it.




Figure 3-1

Indifference curves showing your preferences among different bundles of apples and oranges. The slope of an indifference curve shows the value of one good measured in terms of the other. $\Delta O / \Delta A$ is the average slope of indifference curve U_a between F and D. The slope of mL and Ln are almost equal, indicating that it does not matter whether you measure value in terms of a little more of a good or a little less, provided you consider only very small changes.

Preferences: The Shape of Indifference Curves

All of the indifference curves I have drawn have a similar shape--they slope down and to the right, and the slope becomes less steep the farther right you go (this shape is sometimes described as *convex* or *convex to the origin*). Why?

The curves slope down to the right because both apples and oranges are goods. If one bundle (J) has more of both apples and oranges than another (C), so that a line through them would slope up and to the right, both points cannot be on the same indifference curve. You would obviously prefer J , which has more of both goods, to C . But an indifference curve connects bundles among which you are indifferent. So if a bundle (C) has more apples than another on the same indifference curve (D), putting it farther right, it must have fewer oranges--putting it lower. So indifference curves must slope down to the right, up to the left. If, for some (large) quantity of apples, apples became a bad (you have so many that you would prefer fewer to more), the indifference curve would start to slope up; in order to keep you on the same indifference curve, additional apples (a bad) would have to be balanced by additional oranges (a good).

Two different indifference curves cannot intersect. If they did, the point of intersection would represent a bundle that was on both curves, and therefore had two different utilities. A different way of saying the same thing is to point out that if two indifference curves do intersect, they must have the same utility (the utility of the bundle that is in both of them), and are therefore really one indifference curve.

What can we say about the shape of the curve? As you move from point F to point D along U_a , the number of apples increases by ΔA and the number of oranges decreases by ΔO . Since F and D are on the same indifference curve (U_a), you are indifferent between them. That implies that ΔA apples have the same value to you as ΔO oranges; one apple is worth $\Delta O/\Delta A$ oranges.

$\Delta O/\Delta A$ is the value of an apple measured in oranges. It is also (minus) the slope of the line FD --which is approximately equal to the slope of U_a between F and D (more nearly equal the smaller ΔO and ΔA are). The slope gets less steep as you move down and to the right along the indifference curve, because the value of apples measured in oranges becomes less as you have more apples (farther right) and fewer oranges (lower). This is the same pattern we already saw in Table 3-1.

Figure 3-1 also allows us to see geometrically why the meaning of the value of apples becomes less ambiguous the smaller the changes (in quantity of apples and oranges) we consider. Suppose we start at point L on indifference curve U_c . For large changes in both directions, the two ways of calculating the value of an apple (how many oranges would you have to get to make up for losing one apple versus how many oranges would you be willing to lose in exchange for getting one apple) correspond to finding the slopes of the

lines LM and LN , which are substantially different. For small changes, they correspond to finding the slopes of the shorter lines Lm and Ln , which are almost equal. As the change approaches zero, the two slopes approach equality with each other and with the slope of the indifference curve at L .

The indifference curves on one figure in a textbook are usually very similar; sometimes they are simply the same curve shifted to different positions. In part that is because it is easier to draw them that way, in part because for many utility functions indifference curves that are close to each other have similar shapes. It need not, however, always be true.

Numerical Example

In Figure 3-1, point D is a bundle of 4 apples and 3 oranges, and point F is a bundle of 2 apples and 8 oranges. ΔA is 2 apples and ΔO is 5 oranges. The slope of the line connecting D and F is (minus) $5/2$. $5/2$ is also the value of an apple--2 apples are worth 5 oranges, so an apple is worth $5/2$ of an orange.

Finding the Optimal Bundle

In the previous section, we considered an individual who started with a particular bundle of apples and oranges (A) and could trade apples for oranges at a rate of 1 for 1. In this section, we will analyze essentially the same situation, starting out in a slightly different way. We begin by assuming that you have an income (I), which you can use to buy apples and oranges; the price of apples is P_a and the price of oranges is P_o . If you spend your entire income of \$100 on apples at \$0.50 apiece, you can buy I/P_a ($\$100/\$0.50 = 200$) apples and no oranges, putting you at point K on Figure 3-2a. If you spend your entire income on oranges at \$1 apiece, you can buy I/P_o ($\$100/\$1 = 100$) oranges and no apples, putting you at point L . You should be able to convince yourself, by either algebra or trial and error, that the line B connecting L and K (called the *budget line*) represents all of the different combinations of apples and oranges that you could buy, using your entire income. Its equation is $I = a(P_a) + o(P_o)$ where a is the quantity of apples you buy and o is the quantity of oranges. Put in words, that says that the amount you spend on apples and

oranges equals quantity of apples times price of apples plus quantity of oranges times price of oranges--equals your entire income. Remember that at this point, apples and oranges are the only goods that exist.

Numerical Example

Suppose your income is \$100/month; $P_a = \$0.50/\text{apple}$; $P_o = \$1/\text{orange}$. Table 3-2 shows some of the different bundles that you could buy with your \$100 income. Figure 3-2a shows the corresponding budget line.



Table 3-2

Apples	Oranges	Expenditure
200	0	200 apples x \$0.50 + 0 oranges x \$1 = \$100
160	20	160 apples x \$0.50 + 20 oranges x \$1 = \$100
120	40	120 apples x \$0.50 + 40 oranges x \$1 = \$100
100	50	100 apples x \$0.50 + 50 oranges x \$1 = \$100
60	70	60 apples x \$0.50 + 70 oranges x \$1 = \$100
20	90	20 apples x \$0.50 + 90 oranges x \$1 = \$100
0	100	0 apples x \$0.50 + 100 oranges x \$1 = \$100

Indifference curves, such as those of Figure 3-1, show a consumer's preferences. The budget line plus the region below it (bundles that cost less than his income) show the alternatives available to him--his opportunity set. Figure 3-2a shows both.

The bundles on indifference curve U_5 are preferred to those of the other two curves; unfortunately there is no point that is both on U_5 and on (or below) the budget line--no bundle on U_5 that the consumer can buy with his income. There are two points on U_7 that are also on the budget line (M and N), representing two bundles that the consumer could buy; in addition, the portion of U_7 between the two points is below the budget line and

therefore represents bundles that the consumer could buy and still have some money left over. Should the consumer choose one such point? No. Points O and P are on both the budget line and U_2 ; since U_2 is above (hence preferred to) U_1 , the consumer prefers O (or P) to M or N or any other bundle on U_1 .




The solution to the consumer-choice problem for a world of only 2 goods. B is the budget line for a consumer who has \$100 and can buy oranges at \$1 each or apples at \$0.50 each. The optimal bundle is S , where the budget line is tangent to an indifference curve, since there is no point on B that is on a higher indifference curve than U_4 .

Should the consumer choose a bundle represented by O , P , or one of the points in between, such as Q ? Again the answer is no. Remember that the three indifference curves are merely the three I have chosen to draw out of the infinite number needed to describe the consumer's preferences. Consider point R . It represents a bundle containing more of both goods than Q ; hence it is preferable to Q . Since all points on U_2 are equivalent, R must also be superior to O and P . To find out whether it is the best possible bundle, we draw the indifference curve on which it lies-- U_3 on Figure 3-2b. As I have drawn it, there is another point, S , that lies on a still higher indifference curve and is also on the budget line.

Should the consumer choose S ? Yes. Its indifference curve, U_4 , just touches the budget line. Since any higher indifference curve must be above U_4 , it cannot intersect the budget line. S is the optimal bundle.

It appears that the highest indifference curve consistent with the consumer's income is the one that is just tangent to the budget line, and the optimal bundle is at the point of tangency. This is the usual solution; Figures 3-3a and 3-3b show two exceptions. In each case, the budget line is the same as in Figures 3-2 but the indifference curves are different; the figures represent the same income and prices as Figures 3-2 but different preferences.

On Figure 3-3a, the consumer's optimal point is X on indifference curve U_2 . He could move to a still higher indifference curve by moving down and to the right along the budget line--except that to do so, he would have to consume a negative quantity of oranges! Similarly, in Figure 3-2b, in order to do better than point Y , he would have to consume a negative quantity of apples. These are both corner solutions. In the normal case (interior solution), where the optimal bundle contains both apples and oranges, the result of the previous paragraph holds--the optimal bundle is at the point of tangency.



Corner solutions on an indifference curve diagram. X shows a situation in which the consumer's preferred bundle contains only apples; Y shows a situation in which it contains only oranges.

Price = Value

If two lines are tangent, that means that they are touching and their slopes are the same. The budget line runs from the point $(0, I/P_o)$ to the point $(I/P_a, 0)$, so its slope is $-(I/P_o)/(I/P_a) = -P_a/P_o$. The rate at which you can trade apples for oranges (while keeping your total expenditure fixed) is simply the ratio of the price of an apple to the price of an orange. That is the same thing as the price of an apple measured in oranges; if apples cost \$0.50 and oranges \$1, then in order to get one more apple you must give up half an orange. The price of an apple (measured in oranges) is half an orange. So the slope of the budget line is minus the price of an apple measured in oranges.

The slope of the indifference curve, as I showed earlier in this chapter, is minus the value of an apple measured in oranges. So, in equilibrium, the price of an apple measured in oranges (the rate at which you can transform oranges into apples by selling one and buying the other) is equal to the value of an apple measured in oranges (the *marginal rate of substitution*--the rate at which oranges substitute for apples as consumption goods, the number of oranges you are willing to give up in exchange for an apple). This is the same result that I sketched verbally at the end of the first part of the chapter, when I said that you would keep trading until you reached a point where the value of an additional orange (in apples) was equal to its cost (also in apples).

One possible reaction to this result is "that's obvious; of course the value of something is the same as its price." Another is "this is a bunch of meaningless gobbledygook." Both are

wrong.

To see why the first reaction is wrong, consider what we mean by price and value. Price is what you *have* to give up in order to get something. Value is what you *are just barely willing* to give up to get something. Nothing in those two concepts makes it obvious that they are the same.

The second reaction is much more defensible. You have just been bombarded with a considerable junkpile of abstractions; it may take a while to dig yourself out. You may find it useful to go through the argument in each of the four ways it is presented (two so far, two in the next chapter) until you find one that makes sense to your intuition. Once you have done that, you should be able to go back over the other three and make sense out of them too. One of the reasons for using several different languages is that different people learn in different ways.

This equality between relative prices and relative values is one example of a very general pattern that we will see again and again. I will refer to it as the *equimarginal principle*--marginal because the values being compared are values for one more apple, orange, or whatever. It is a statement not about our tastes but about equilibrium--where we are when we stop trading. The same pattern has already appeared several times, in a very different context, in the optional sections of Chapter 1, where we saw that in equilibrium all lines in a supermarket and all lanes on a freeway are equally attractive--provided that the cost of getting to them is the same.

The Invisible World--A Brief Digression

Another response you may have at this point is "Where do all these tables and indifference curves come from, anyway? How can you possibly know what my preferences are? How, for that matter, can I know exactly how many apples I would give for an orange? Are economists people who go around asking people what bundles they are indifferent among--and are they fools enough to believe the answers?"


I shall answer the five questions in order. The tables and figures all came out of my head--I made them up, subject to the requirement that the numbers in the table have a certain pattern and the curves a certain shape. I cannot tell what your preferences are. You do not

know exactly what your preferences are. No. No.

If we cannot find out what indifference curves are, what good are they? The answer is that indifference curves--like much of the rest of economics--are tools used to help us think clearly about human behavior. By using them, we can show that if people have preferences and rationally pursue them (the assumptions that I made and defended in Chapter 1), certain consequences follow. So far in this chapter, I have concentrated on one particular consequence--the equality between relative values and relative prices. I will show others later. Indifference curves and the like are useful as analytical tools; it is a serious error to think of them as things we actually expect to go out and measure.

Income and Substitution Effects

Now that we know what indifference curves are, we shall use them to show how the amount you consume of a good varies with your income and with the price of the good. Figure 3-4 shows what happens as income rises, with price held constant. B_1 is the same budget line we have seen before, corresponding to an income of \$100 and prices of \$1/orange and \$0.50/apple. B_2 is the budget line for the same prices but for an income of \$125, B_3 for an income of \$150. Since relative prices are the same in all three cases, all three budget lines have the same slope, making them parallel to each other. In each case, I have drawn in the indifference curve that is just tangent to the budget line. As income rises, the consumption bundle shifts from X to Y to Z ; in the case illustrated, consumption of both apples and oranges rises with income--they are *normal goods*. The line IEP is the *income expansion path* showing how consumption of apples and oranges changes as the consumer's income increases.



Optimal bundles for three different incomes--2 normal goods. X is the optimal bundle for an income of \$100, Y for an income of \$125, and Z for an income of \$150--as shown by B_1 , B_2 , and B_3 . Consumption of both apples and oranges increases with increasing income.


Figure 3-5 shows the same pattern of income and prices but a different set of indifference curves, corresponding to an individual with different preferences. This time, as income increases, the consumption of oranges increases but the consumption of apples decreases! In such a situation, apples are an *inferior good*--a good of which we consume less the richer we are. Hamburger and beans are both plausible examples of inferior goods, for some ranges of income. As a very poor person becomes less poor, he eats hamburger instead of beans; his consumption of beans goes down as his income goes up, so for that range of incomes beans are an inferior good. As his income becomes still higher, he starts eating steak instead of hamburger. His consumption of hamburger goes down as his income goes up, so for that range of incomes, hamburger is an inferior good.

In describing the budget lines B_1 , B_2 , and B_3 , I gave specific values for income and prices. I could just as easily have told you that income was \$200, \$250, and \$300 and that prices were \$2/orange and \$1/apple; that would have produced exactly the same budget lines. The reason is obvious: If you double your income and simultaneously double the price of everything you buy, your real situation is unchanged--you can buy exactly the same goods as before.

I could also have told you that income was \$100 for all three budget lines and that the price of an orange was \$1 for B_1 , \$0.80 for B_2 , and \$0.66 $\frac{2}{3}$ for B_3 , with corresponding prices (\$0.50, \$0.40, \$0.33 $\frac{1}{3}$) for apples. A drop in the price of everything you consume has the same effect on what you can buy as an increase in income.

It is not obvious when we should describe changes on an indifference curve diagram--or changes in the situations that such diagrams represent--as changes in prices and when we should describe them as changes in income. That is not because there is something wrong with indifference curves but because the distinction between a change in income and a change in price is less clear than it at first seems. We are used to thinking of prices and incomes in terms of money, but money is important only for what it can buy; if all prices go down and my income stays the same, my *real income*--my ability to buy things--has risen in exactly the same way as if prices had stayed the same and my money income had gone up.

If income and prices all change at once, how can we say whether my real income has gone up, gone down, or stayed the same? Income is useful for what it can buy; the value to me of the bundle of goods that I buy is indicated, on an indifference curve diagram, by what indifference curve it is on. It therefore seems natural to say that a change in money income and prices that leaves me on the same indifference curve as before has left my real income unchanged. A change that leaves me on a higher indifference curve has increased my real income; a change that leaves me on a lower indifference curve has lowered my real income.

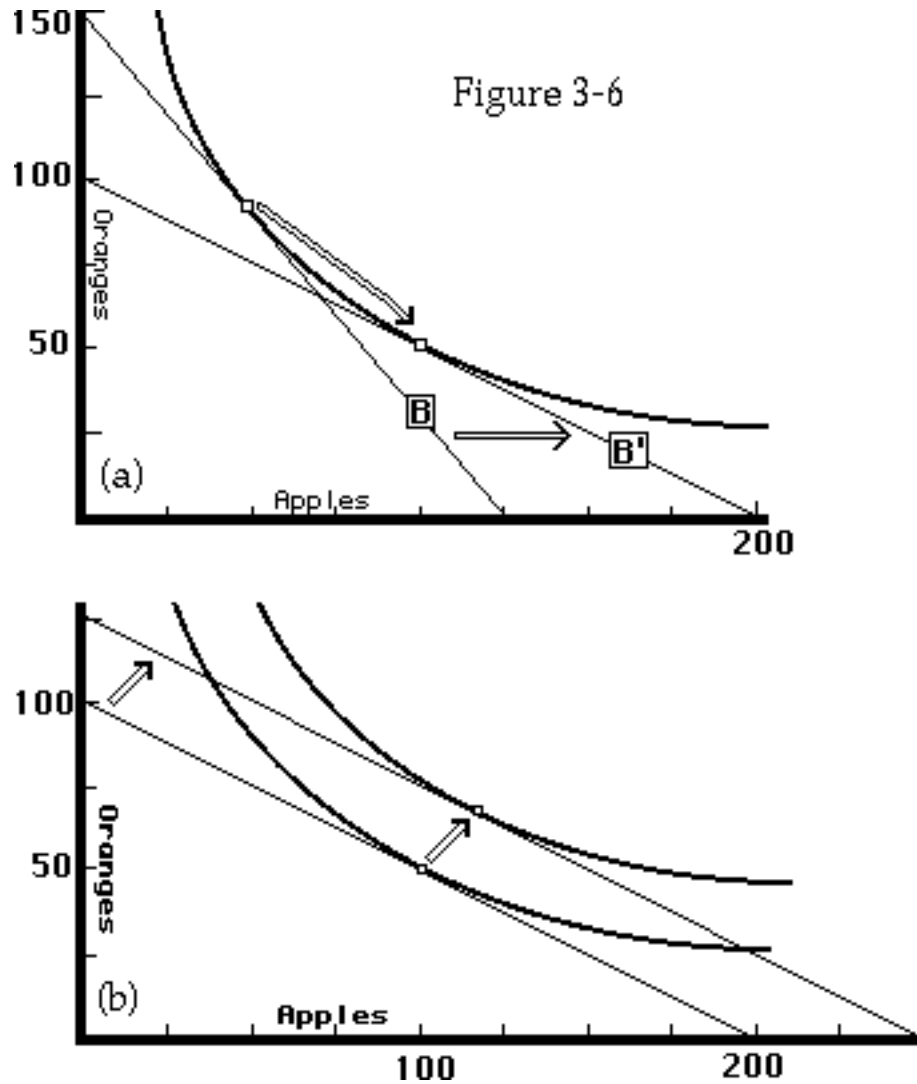


Optimal bundles for three different incomes--a normal good and an inferior good. As income increases, consumption of oranges increases but consumption of apples decreases; so apples are an inferior good. IEP is the income expansion path.

The prices that are important are *relative prices*--how much of one good I must give up to get another. As I showed earlier, the price of one good in terms of another corresponds to (minus) the slope of the budget line. So a change in money income and money prices that alters the slope of the budget line while leaving me on the same indifference curve is a pure change in prices--prices have changed and (real) income has not. A change that leaves the slope of the budget line the same but shifts it so that it is tangent to a different indifference curve is a pure change in income--real income has changed but (relative) prices have not. An example of the former is shown on Figure 3-6a; an example of the latter, on Figure 3-6b.

Figure 3-7 shows the effect of a decrease in the price of apples. B_1 is the same budget line as before; A is the optimal bundle on B_1 . B_2 is a budget line for the same income (\$100) and the same price of oranges (\$1/orange), but for a new and lower price of apples (0.33 1/3/apple). C is the optimal point on that budget line. We can decompose the movement from point A to point C into two parts, as shown in Figure 3-7. A pure change in price with real income fixed would leave us on the same indifference curve, changing the

budget line from B_1 to B' and the optimal bundle from A to B . A pure change in income would keep relative prices (the slope of the budget line) unchanged, while moving us to a different indifference curve. That is the movement from bundle B on budget line B' to bundle C on budget line B_2 ; note that B' and B_2 are parallel to each other. The change in our consumption as we move from A to B is called the *substitution effect* (we substitute apples for oranges because they have become relatively cheaper); the change as we move from B to C is called the *income effect*.



A pure change in price (a) and a pure change in income (b). On Figure 3-6a, relative prices change, but real income does not, since the individual ends up on the same indifference curve after the change. On Figure 3-6b, relative prices stay the same but real income increases.




Figure 3-7

The effect of a fall in the price of apples. When the price of apples falls, the optimal bundle changes from A to C. The movement from A to B is a substitution effect--relative prices change, real income does not. The movement from B to C is an income effect; real income changes, relative prices do not. A further price drop moves the optimal bundle to D. The line PEP, running from A to C to D, is the price expansion path.

A pure substitution effect always increases the consumption of the good that has become relatively cheaper. You can see that by looking at the shape of the indifference curve and imagining what happens as the budget line "rolls along it" (as it does from B_1 to B'). This corresponds to lowering the price of one good while at the same time cancelling out the gain to the consumer by either raising the price of the other good or lowering income. On net, the consumer is neither better off nor worse off. The result is to increase the consumption of the good that has become cheaper. The pure income effect from a decrease in the price of a good (an increase in real income), on the other hand, may either increase or decrease its consumption, according to whether it is a normal or an inferior good.

A drop in the price of one good without any compensating change in income or other prices produces both a substitution effect and an income effect, as shown on Figure 3-7; apples are cheaper than before relative to oranges, and the lower price of apples makes the consumer better off than before. The substitution effect always increases the consumption of the good whose price has fallen; the income effect may increase or decrease it. You can see the net effect by looking at the *price expansion path* (PEP on Figure 3-7), which shows how consumption of both goods changes (from *A* to *C* to *D*) with changes in the price of one good.

This suggests the possibility of a good so strongly inferior that the income effect more than cancels the substitution effect--as its price falls, its consumption goes down. Imagine, for example, that you are spending most of your income on hamburger. If the price of hamburger falls by 50 percent while your income and all other prices remain the same, your real income has almost doubled. Since you are now much richer than before, you may decide to buy some steak and reduce your consumption of hamburger. The substitution effect tends to make you consume more hamburger; at the lower price of hamburger, the money required to buy an ounce of steak would buy twice as much hamburger as before the price change; so steak is more expensive in terms of hamburger than before. But you are now much richer--so you may choose to eat more steak in spite of its higher relative cost.

A good whose consumption goes down instead of up when its price goes down is called a *Giffen good*. It is not clear whether any such goods actually exist. The reason is that most of us consume many different goods, spending only a small part of our income on any one. A drop in the price of one good has a large effect on its relative price (hence a large substitution effect) but only a small effect on our real income. A Giffen good must either consume a large fraction of income or be so strongly inferior that the effect of a small change in income outweighs that of a large change in relative price.

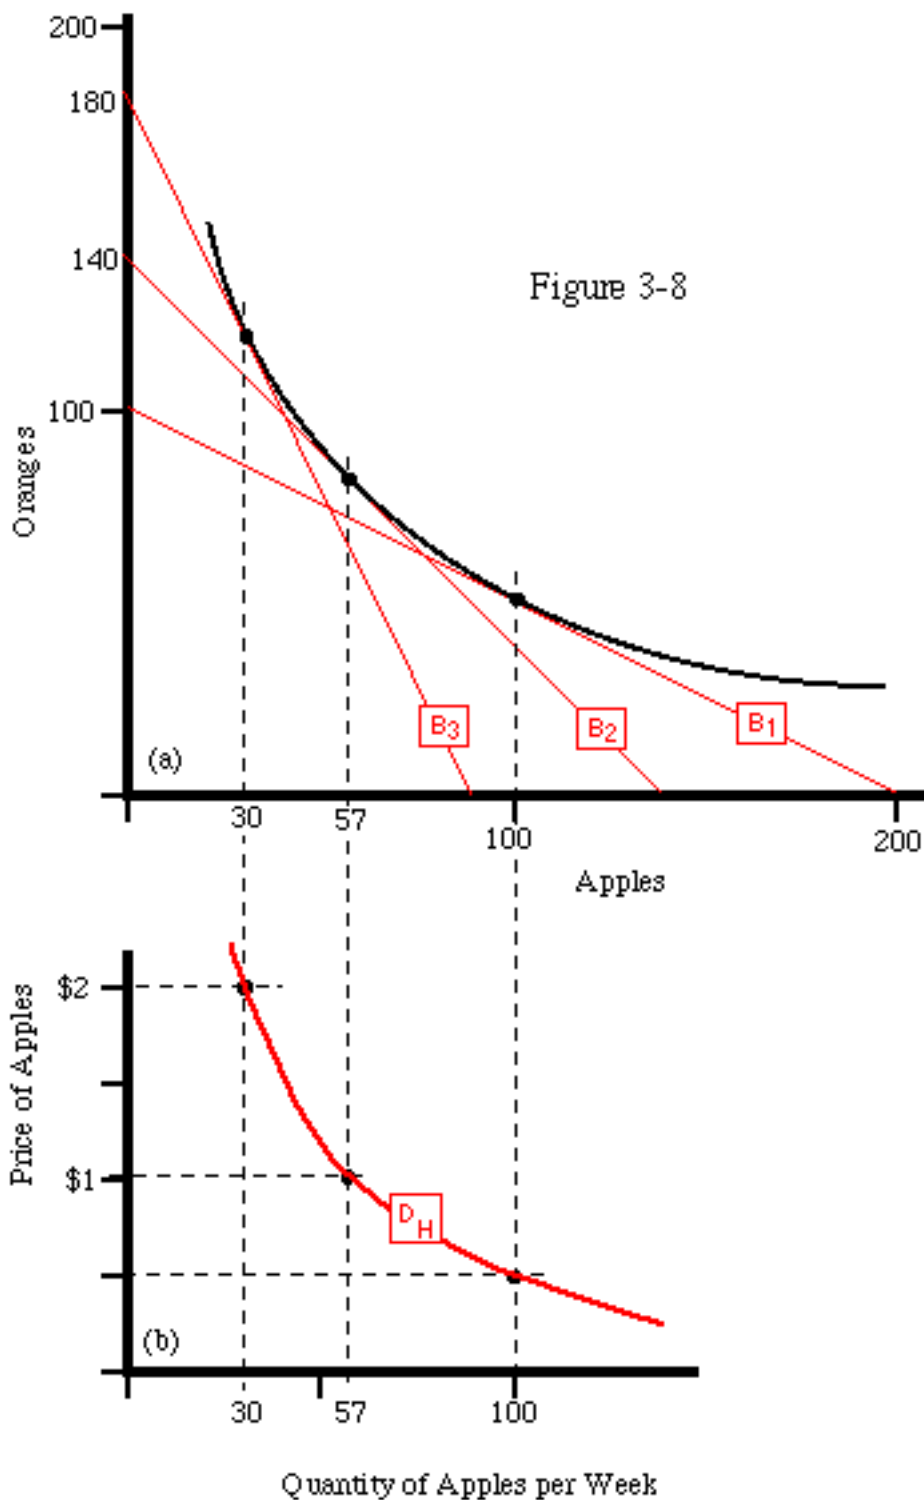
Students frequently confuse the idea of an inferior good with the idea of a Giffen good. An inferior good is a good that you buy less of when *your income* goes up. There are many examples--for some of you, McDonald's hamburgers or bicycles. A Giffen good is a good that you buy less of when *its price* goes down. A Giffen good must be an inferior good, but most inferior goods are not Giffen goods.

If Giffen goods are rare or nonexistent, why have I spent time discussing them? The main reason is that in much of economic analysis (including a good deal of this book), we

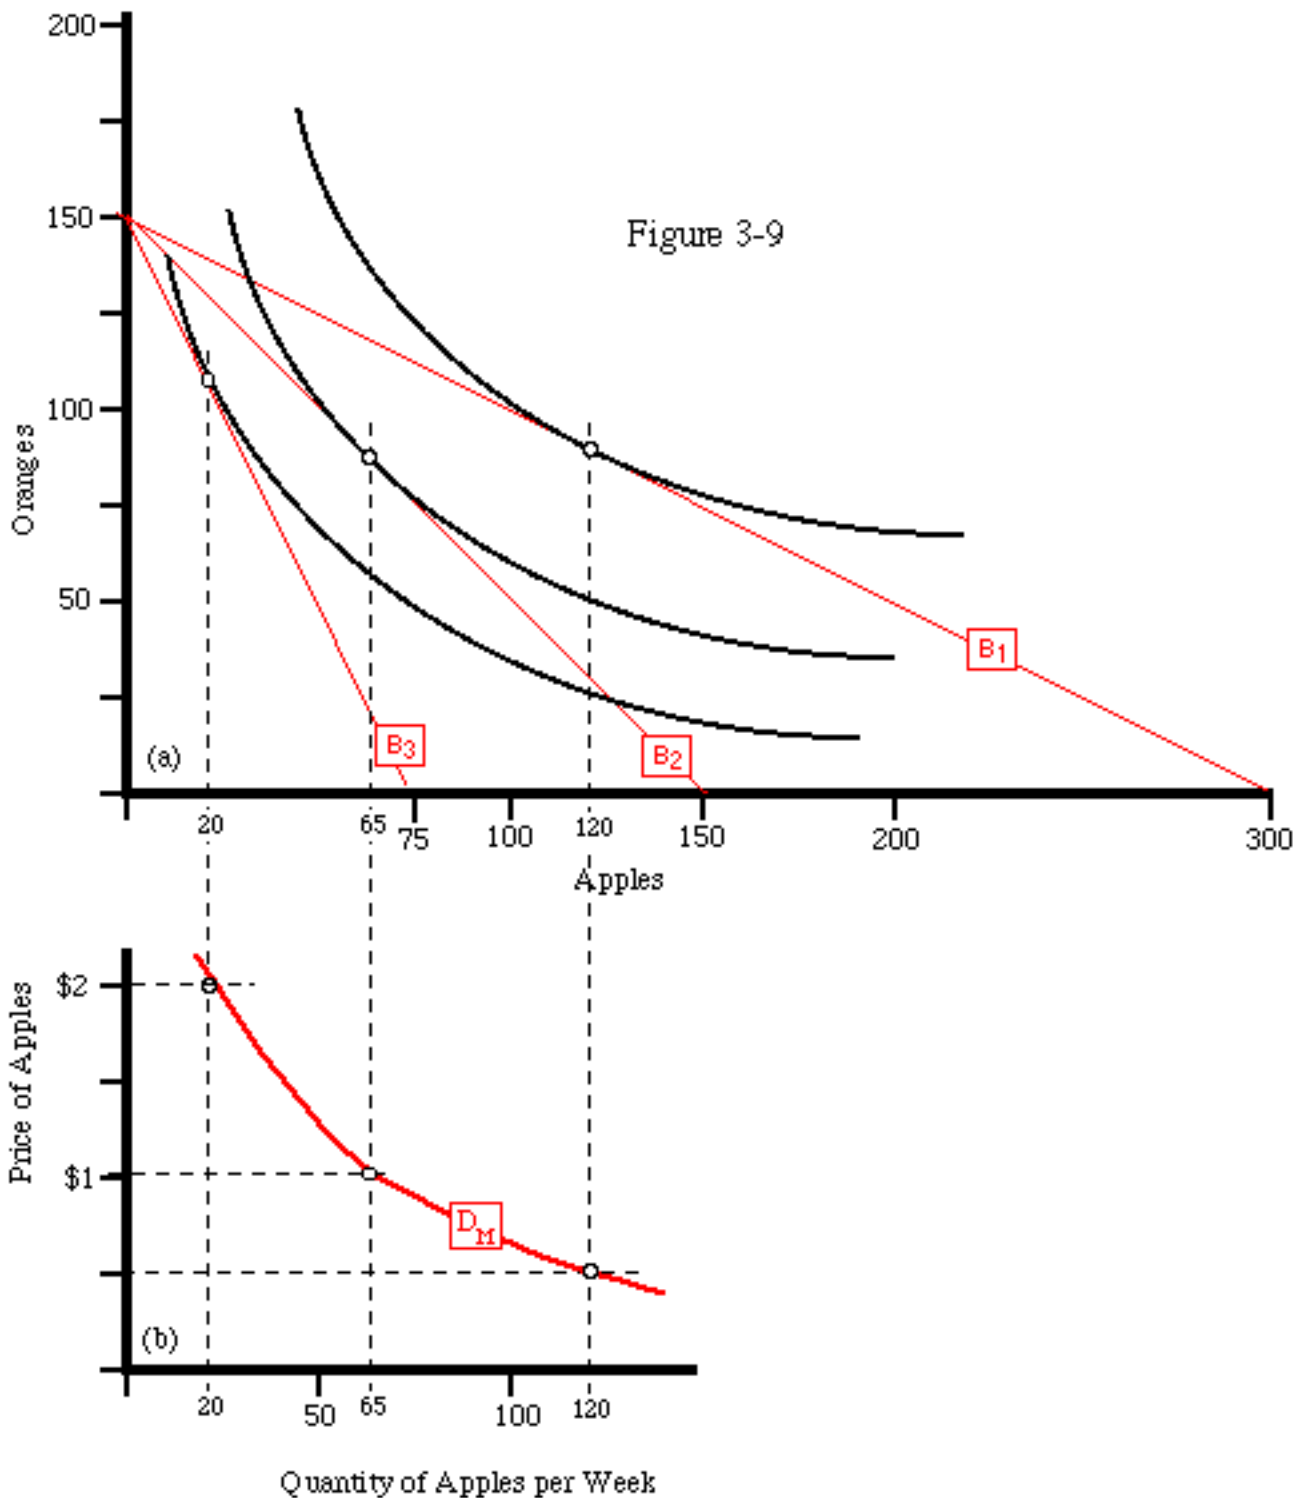
assume that demand curves slope down--that the higher the price of something is, the less of it you buy. If I am going to use that assumption over and over again, it is only fair to give you some idea of how solid it is--by describing the circumstances in which it would be false.

Demand Curves

Figure 3-6a showed the effect on consumption of a pure change in price. Figures 3-8a and b and Table 3-3 show how the same analysis can be used to derive an *income-compensated demand curve* (also known as a *Hicksian* demand curve after economist John Hicks). The budget lines on Figure 3-8a correspond to a series of different prices for apples, from \$0.50/apple to \$2/apple. The price of oranges is held constant at \$1/orange. Table 3-3 shows prices, quantities, and income for each budget line. Figure 3-8b is the resulting demand curve. It shows the relation between price of apples and quantity purchased for the consumer whose preferences are represented on Figure 3-8a. It is an income-compensated demand curve because, as we increase the price of apples, we also increase income by just enough to keep the consumer on the same indifference curve. We thus eliminate the income effect; the change in the quantity purchased is due to the substitution effect alone.



The derivation of an income-adjusted demand curve. Budget lines B_1 , B_2 and B_3 show different combinations of prices and income corresponding to the same real income. D_H is the resulting income-adjusted (Hicksian) demand curve.



The derivation of an ordinary demand curve. Budget lines B_1 , B_2 and B_3 show different prices of apples but the same income and price of oranges. D_M is the ordinary (Marshallian) demand curve.

Figures 3-9a and b and Table 3-4 show the similar derivation of an ordinary demand curve

(called a *Marshallian* demand curve after economist Alfred Marshall). This time, just as on Figure 3-7, the price of apples is changed while both the price of oranges and income are held constant. The higher the price of apples, the worse off the consumer; his dollar income is the same but, since his dollars will buy fewer apples, his real income is lower. So the higher the price of an apple on Figure 3-9a, the lower the indifference curve to which the corresponding budget line is tangent. This time the change in quantity purchased includes both an income and a substitution effect.

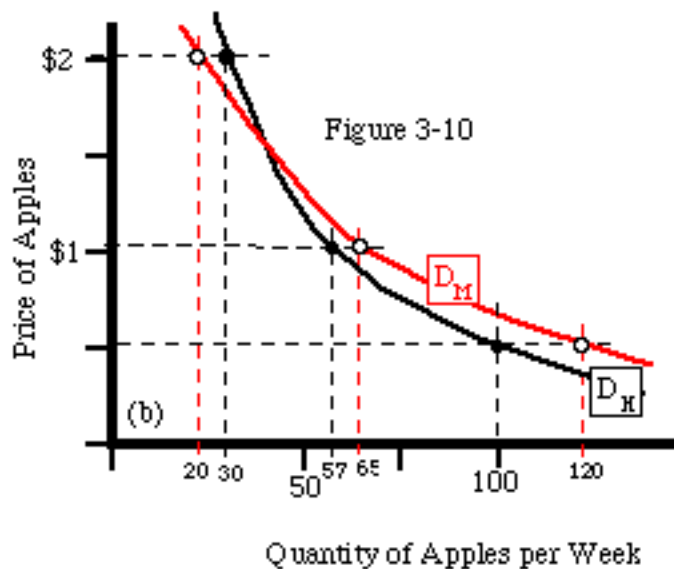
Table 3-3

Budget Line	Price of Apples	Price of Oranges	Income (\$/week)	Quantity of Apples Purchased per Week
B ₁	\$0.50	\$1.00	100	100
B ₂	\$1.00	\$1.00	140	57
B ₃	\$2.00	\$1.00	180	30

Table 3-4

Budget Line	Price of Apples	Price of Oranges	Income (\$/week)	Quantity of Apples Purchased per Week
B ₁	\$0.50	\$1.00	150	120
B ₂	\$1.00	\$1.00	150	65
B ₃	\$2.00	\$1.00	150	20

For most economic problems, the relevant demand curve is the Marshallian one, since there is generally no reason to expect a change in the price of one good to cause a compensating change in income or other prices. Some parts of economic theory however, including consumer surplus, which will be explained in Chapter 4, can be derived rigorously only by using income-compensated demand curves.



Ordinary and income-adjusted demand curves for the same individual. D_M is the ordinary (Marshallian) demand curve; D_H is the income-adjusted (Hicksian) demand curve.

The Marshallian demand curve D_M on Figure 3-9b and the Hicksian demand curve D_H on Figure 3-8b are significantly different, as you can see on Figure 3-10. That is because we are considering a world with only two goods. Since raising the price of one of them makes the consumer significantly worse off, his behavior (the amount of the good he buys) is substantially different according to whether we do or do not compensate him for the change.

But in the real world, as I pointed out earlier, we divide our expenditure among many goods. If I spend only a small fraction of my income on a particular good, a change in its price has only a small effect on my real income. In such a case, the difference between the two demand curves is likely to be very small. For this reason, we will generally ignore the distinction between ordinary and income-compensated demand curves in what follows.

Application: Housing Prices--A Paradox

You have just bought a house. A month after you have concluded the deal, the price of houses goes up. Are you better off (your house is worth more) or worse off (prices are higher) as a result of the price change? Most people will reply that you are better off; you

own a house and houses are now more valuable.

You have just bought a house. A month after you have concluded the deal, the price of houses goes down. Are you worse off (your house is worth less) or better off (prices are lower)? Most people, in my experience, reply that you are worse off. The answers seem consistent, even to those who are not sure what the right answer is. It appears obvious that if a rise in the price of housing makes you better off, then a fall must make you worse off, and if a rise makes you worse off, then a fall must make you better off.

Although it appears obvious, it is wrong. The correct answer is that either a rise or a fall in the price of housing makes you better off!

Before proving this, I will first describe the situation a little more precisely. I am assuming that you have an income (I), part of which went to buy the house. One may imagine either that your income is from a portfolio of stocks and bonds, part of which you sold in order to buy the house, or that you have a salary, part of which must now go for interest on the mortgage. In either case, you have bought housing and, as a result, have less to spend on other goods.

I am also assuming that none of the circumstances determining how much housing you want are ever going to change, except for the price of housing; if the price of housing stayed the same, so would the amount of housing you want. You are not, in other words, planning to have children and move to a bigger house or planning to retire, sell your house, and move to Florida. To simplify the argument, I will ignore all costs of buying, selling, or owning housing other than the price--sales taxes, realtor's commissions, and the like. Finally, I will assume that the change in price was unexpected; when you bought the house you were assuming that the price of housing, like everything else, was going to stay the same forever.

Now that I have described the situation more precisely, you may want to stop and try to figure out how my answer--that a change in either direction benefits you--can be true.

The situation is shown in Figure 3-11. The vertical axis represents housing; the horizontal axis represents expenditure on all other goods. The budget line B_1 shows the different combinations of quantity of housing and quantity of other expenditure you could have chosen at the initial price of housing (\$50/square foot). Point A_1 is the optimal bundle--the amount of housing you bought. It is on indifference curve U_1 .

Line B_2 shows the situation after the price of housing has risen to \$75/square foot. B_2 , your new budget line, must have a slope of (minus) 1 square foot of housing/\$75, since that is the new price of housing--the rate at which you can exchange dollars spent on all other goods for housing, or vice versa. The new budget line must go through point A_1 , since one of the alternatives available to you is to do nothing--to keep the bundle that you had before the price change. You can choose to move away from point A_1 along the budget line either up (buy more housing, trading dollars for housing at a rate of \$75/square foot) or down (sell your house and move to a smaller one--sell some of your housing for money at \$75/square foot). So your new budget line, B_2 , is simply a line with slope $-1/75$ drawn through point A_1 .




Figure 3-11

The effect on a homeowner of a change in the price of housing. B_1 shows the alternatives available at the original price of housing; B_2 shows those available if the price of housing rises after the house is bought; B_3 shows the alternatives available if the price falls. A_1 shows the homeowner's bundle of housing and all other consumption after the house is built and before there is any change in housing prices. The change in the slope of the budget line has been exaggerated to make the effect clearer.

The figure shows what you choose to do; your new optimal point is at A_2 . Since housing is

now more expensive than before, you have chosen to sell your house and buy a smaller one--the gain in income is worth more to you than the reduction in the amount of housing you consume. You are now on indifference curve U_2 , which is above (preferred to) U_1 .

Line B_3 shows the situation if the price of housing goes down rather than up after you buy your house--to \$30/square foot. It is drawn in exactly the same way except that the price ratio is now $1/30$. Again you have the choice of keeping your original house, so the line has to go through A_1 . Your new optimal point is A_3 ; you have adjusted to the lower price of housing by selling your house and buying a bigger one. You are now on U_3 --which is above U_1 ! The drop in the price of your house has made you better off.

By looking at the figures, you should be able to convince yourself that the result is a general one; whether housing prices go up or down after you buy your house, you are better off than if they had stayed the same. The same argument can be put in words as follows:

What matters to you is what you consume--how much housing and how much of everything else. Before the price change, the bundle you had chosen--your house plus whatever you were buying with the rest of your income--was the best bundle of those available to you. If prices had not changed, you would have continued to consume that bundle. After prices change, you can still choose to consume the same bundle. The house belongs to you, so as long as you choose to keep it, the amount of money you have to spend on other things is unaffected by the price of the house.

You cannot be worse off as a result of the price change--at worst you continue to consume the same bundle (of housing and other goods) as before. But since the optimal combination of housing and other goods depends, among other things, on the price of housing, it is unlikely that the old bundle is still optimal. If it is not, that means there is now some more attractive alternative, so you are now better off; a new alternative exists that you prefer to the best alternative (the old bundle) that you had before.

This seemingly paradoxical result is interesting in part for what it shows us about the relative virtues of our different languages. In solving the problem geometrically, the drawing tells us the answer. All we have to do is look at Figure 3-11 in order to see that any budget line that goes through A_1 with a different slope than B_1 has to intersect some indifference curve higher than U_1 --whether the slope is steeper (lower price of housing) or shallower (higher price of housing). What the drawing does not tell us is *why* it is true. When we solve the problem verbally, we are likely to get the wrong answer (as at the beginning of this section, where I concluded that a fall in the price should make you worse off). But once we do get the right answer (possibly with some help from the figures), we not only know what is true, we also understand why.

I have ignored the transaction costs associated with buying and selling houses--realtor's commissions, sales taxes, the time spent finding a satisfactory house, and so on. If such costs are included, the result is that small changes in housing prices have no effect at all on you--it is not worth paying the transaction costs necessary to increase or decrease your consumption of housing. Large changes in either direction benefit you.

If you still find the result puzzling, the reason may be that you are confusing two quite different questions--whether a change in price makes you better off, given that you have bought a house, and whether having bought a house made you better off, given that the price is going to change. I have been discussing the first question. I asked whether, given that you had bought a house, a subsequent change in price made you better or worse off. The conclusion was that it made you better off, whether the price went up or down. That does not mean that buying the house was a good idea; if the price is going to go down, you would have been still better off if you had waited until it did so before you bought. The alternatives we have been comparing are "buy a house and have the price go down (or up)" versus "buy a house and have the price stay the same," not "buy a house and have the price go down (or up)" versus "have the price go down (or up) and *then* buy a house."


Application: Subsidies

Figure 3-12 shows your preferences between potatoes and expenditure on all other goods. You have an income of \$150/week; the price of potatoes is \$3/pound. If you spend all your income on potatoes, you can consume 50 pounds per week of potatoes and nothing else. If you spend nothing on potatoes, you have \$150/week left to spend on all other

goods. Line B is your budget line; point D is the bundle you choose.

The potato lobby convinces the government that potatoes are good for you and should therefore be subsidized. For every \$3 you spend on potatoes, the government gives you \$1. So for each pound of potatoes you buy, you have \$2 less (instead of \$3 less) to spend on other goods--the cost of potatoes to you is now only \$2/pound instead of \$3/pound.

If you choose to buy no potatoes, you are unaffected by the subsidy and can spend your entire income of \$150/week on other goods. If you choose to spend your entire income on potatoes, you can now buy 75 pounds per week. B_1 is your new budget line. Your optimal bundle is D_1 . Your consumption of potatoes has risen. Since you are on a higher indifference curve than before-- U_4 instead of U_2 --you are better off than before. You are happier (and, if the potato farmers are right, healthier); the potato farmers are selling more potatoes; all is well with the world.




The effect of a potato subsidy that someone else pays for. B is the initial budget line, B_1

is the budget line after the government announces that it will pay one third of the cost of the potatoes you buy.

There is one problem. At point D_I you are consuming 40 pounds per week of potatoes (if that seems unreasonable, you may assume that some of the potatoes are converted into vodka before you consume them). Each pound costs \$3, of which you pay only \$2; the other dollar is provided by the subsidy. So the total subsidy is \$40/week. Some taxpayers somewhere are paying for that subsidy. Before we conclude that the potato subsidy is a complete success, we should include their costs in our calculations.

To do so, I will assume that consumers and taxpayers are the same people. For simplicity I will also assume that everyone has the same income and the same preferences, as shown on Figures 3-12 and 3-13. Each individual has a pretax income of $I = \$150/\text{week}$ and an aftertax income of $I - T$, where T is the amount of tax paid. To find T on the figure, note that a consumer who buys no potatoes has $I - T$ (income minus tax) available to spend on everything else, so $I - T$ is the vertical intercept of the budget line. While we do not yet know what T is, we do know that the total amount collected in taxes must be the same as the amount paid out in subsidy (we ignore the cost of collecting taxes and administering the subsidy).

For the population as a whole, tax collected equals subsidy paid, and the amount of subsidy paid depends on how many pounds of potatoes people buy. But from the standpoint of each individual in a large population, the quantity of potatoes he buys has a negligible effect on the total subsidy and hence on his taxes. So each individual takes T as given and finds his optimal bundle, as shown on Figure 3-13. Since the effective price of potatoes is still \$2/pound (pay \$3 and get \$1 back as subsidy), the corresponding budget line (B') has the same slope as B_I on Figure 3-12.



The effect of a potato subsidy that you pay for. T is the tax that pays for the subsidy; B' is the budget line for a consumer who pays the tax and can receive the subsidy. At D' , the optimal point on B' , the consumer pays exactly as much in tax as he receives in subsidy—and is worse off than he would be, at D , if neither tax nor subsidy existed.

How do we know that the budget line is B' , instead of some other line with the same slope? B' is the only budget line for which the tax collected from each taxpayer ($T = I - (I - T) = \$150/\text{week} - \$120/\text{week} = \$30/\text{week}$) is exactly equal to the subsidy paid to each taxpayer ($\$1/\text{lb} \times 30 \text{ lb}/\text{week}$ at point D')—as it must be, since everyone is identical and total taxes paid must equal total subsidy received.

When Is a Wash Not a Wash? D' , the bundle you choose to consume, lies on both B' , your budget line given the tax and subsidy, and B , your original budget line. This is not an accident. In the simple case I have described, everyone buys the same amount of potatoes, receives the same amount of subsidy, and pays the same amount of taxes; so taxes and subsidy must be equal not only for the population as a whole but for each individual

separately. If you pay as much in taxes as you receive in subsidy, tax and subsidy cancel; the bundle (potatoes plus expenditure on all other goods) that you purchase is one you could have purchased from your original income if there had been neither tax nor subsidy. So it must be on B , your initial budget line.

That, in fact, is how I found B' in the first place. I knew that B' had to be parallel to B . I also knew that its optimal point, where it was tangent to an indifference curve, had to occur where it intersected B . B' was the only line that met both conditions.

D' is on a lower indifference curve than D --the combination of tax and subsidy makes you worse off. This is not accidental. Since D' is on your original budget line, it is a bundle that you could have chosen to consume if there had been no subsidy and no tax. In that situation, you chose D instead, so you must prefer D to D' . So the combination of a subsidy and a tax that just pays for the subsidy must make you worse off.

In accounting, a transaction that results in two terms that just cancel--a \$1,000 gain balanced by a \$1,000 loss--is referred to as a wash. Your first reaction on reading the previous few paragraphs may be that the sort of tax/subsidy combination I have described is a wash; since you are getting back just as much as you are paying, there is no net effect at all.

In one sense, that is true; in another and more important sense, it is not. The total dollar value of your consumption bundle is the same with or without the tax/subsidy combination; in that sense, there is no effect. But, as you can see on Figure 3-13, the bundle you choose is different in the two cases; with the tax and subsidy, you end up choosing a less attractive consumption bundle--one on a lower indifference curve--than without it.

The reason for the difference goes back to a point I made earlier--that although the amount of the tax is determined for the population as a whole by how many pounds of potatoes are consumed, each individual will and should treat the amount of the tax as a given when deciding how many potatoes to buy. Given what everyone else is doing, your budget line (with the tax and subsidy) is B' , not B . Since B' does not include D , you do not have the option of choosing that bundle. All of us, acting together, could choose D ; each of us, rationally responding to the subsidy and the rational behavior of everyone else, chooses D' . This seemingly paradoxical result--that in some situations, rational behavior by every individual leaves each individual worse off--is not new. We encountered it before when

we were explaining why armies run away and traffic jams.

Where You Are Going, Not How You Get There. Students faced with something like the potato subsidy problem often make the mistake of trying to solve it in stages. First they draw the budget line representing the subsidy (B_1). From that they calculate how many potatoes the consumer buys, then from that they calculate the amount of the tax necessary to pay for the subsidy. The problem with this approach is that imposing the tax shifts the budget line, which changes the number of potatoes consumed, which changes the amount of subsidy paid out, which changes the amount of tax needed to pay for the subsidy! You are caught in an infinite loop; each time you solve one part of the problem another part changes.

The solution is to ignore the series of successive approximations by which someone trying to find the tax that just paid for the subsidy would grope his way towards the solution, and simply ask what the solution must look like when he has finally reached it. That is what we did on Figure 3-13. A subsidy of \$1/lb implies a budget line parallel to B_1 . A tax that just pays for the subsidy implies a budget line whose optimal point (where it is tangent to an indifference curve) occurs where it intersects B_1 --meaning that consumers with that budget line buy a quantity of potatoes such that the tax just pays for the subsidy. B' is the only budget line you can draw on Figure 3-13 that meets both of those conditions, so it must be the solution.

Fine Point. One assumption implicit throughout this discussion is that the tax/subsidy does not affect the market price of potatoes; that was always assumed to be \$3/pound. The assumption is a reasonable one if we imagine that the subsidy and tax apply to only a small part of the population--say, a single town. Changes in the potato consumption of Podunk are unlikely to have much effect on the world market price of potatoes. It is less reasonable if we consider a program applying to the entire population of the United States. One effect of the subsidy is to increase the demand for potatoes, which should produce an increase in their price. That is one of the reasons why the potato farmers are in favor of the subsidy.

This raises a second question. So far, in analyzing the problem, we have only considered the interests of the consumers and the taxpayers; what about the producers? Is it possible that if we take them into account as well, the net effect of the subsidy is positive?

Insofar as we can answer that question--insofar, in other words, as we have a way of

adding up different people's gains and losses--the answer is no. Even including the effect on the producers, the net effect of the subsidy is negative. You will have to wait until Chapter 17 to learn why.

Other Constraints

The same techniques that we have been using to analyze the constraint imposed upon a consumer by his limited income could just as easily be used to analyze other sorts of constraints. Consider, for example, someone on a thousand calorie/day diet. He faces a calorie constraint. Each food has a price in calories per ounce; he must choose a bundle of foods whose total cost is no more than a thousand calories. If he is considering only two alternative foods the thousand calorie bundles will lie along a budget line; his optimal bundle will be where that budget line is tangent to an indifference curve.

There is another constraint that applies to everyone, even those fortunate enough not to have to diet. Most things we do, including earning money and spending it, require time. Each of us must allocate his limited budget of 24 hours a day among a variety of uses--work, play, consumption, rest. If we consider only two alternatives, holding the rest fixed, we again have a choice that can be represented by an indifference curve diagram.

OPTIONAL SECTION

UTILITY FUNCTIONS

Utility

Utility and the utility function were important ideas in the development of economics and remain useful as tools for thinking about rational behavior. The idea of utility grows out of the attempt to understand all of an individual's choices in terms of a single thing he is trying to maximize--happiness, pleasure, or something similar. We call this his utility. Utility is observed only in choices. The statement "The utility to you of a Hawaiian

vacation is greater than the utility to you of a moped" is equivalent to the statement "Given the choice between a Hawaiian vacation and a moped, you would choose the vacation." It does not mean "A vacation is more useful to you than a moped." Used as a technical term in economics, utility does not have the same meaning as in other contexts.

A *utility function* is a way of describing your preferences among different bundles of goods. Suppose we consider only two goods--apples and pears. The statement "Your utility function is 3 x (number of pounds of apples) + 2 x (number of pounds of pears)," which we write mathematically as

$$u(a,p) = 3a + 2p,$$

means that if you have to choose between two bundles of apples and pears, you will choose the bundle for which that function is greater. You will prefer four pounds of apples plus three pounds of pears (total utility = 18) to three pounds of apples plus four pounds of pears (total utility = 17).

If you are not familiar with functions, you may find the expression $u(a,p)$ confusing. All it means is "utility, which depends on a (the number of pounds of apples) and p (the number of pounds of pears)." The form of the dependence is then shown on the other side of the equality sign.

Several things are worth noting about such functions. The first is that we are very unlikely to know what someone's utility function actually is--we would have to know his preferences among all possible bundles of goods. The purpose of utility functions is to clarify our thinking by allowing us to build simplified pictures of how people act. Such *models* are not attempts to describe reality; they are attempts to set up a simplified situation with the same logical structure as the much more complicated reality in order to use the former to understand the latter. You should not confuse such models with large-scale econometric models--complicated sets of equations used (not very successfully) to try to predict the behavior of some real-world economy.

The second point to note is that the same pattern of behavior can be described by many different utility functions. In the example given above, suppose the utility function had been not u but

$$v(a,p) = 6a + 4p = 2 \times u(a,p).$$

The second function (v) is just twice the first (u); if the first is larger for one combination of apples and oranges than for another, so is the second. An individual always chooses the bundle that has higher utility, so the two utility functions imply exactly the same behavior.

So far, we have assumed that your utility depends on only two goods. More generally we can write $u(\mathbf{x})$, where \mathbf{x} is a bundle of goods. In the simple two-good case, \mathbf{x} is the number of apples and of pears; we could write $\mathbf{x} = (2,3)$ to describe a bundle of 2 apples and 3 pears. In the more general case of n goods, \mathbf{x} is a longer list, describing how much of each good is in the particular bundle being considered. If we call the first good X_1 and the amount of the first good x_1 , the second good X_2 and the amount of the second good x_2 , and so on, and if the price of the first good is P_1 and similarly for the other goods, your income constraint--the requirement that the total bundle you purchase is worth no more than your total income--is the equation

$$I \geq P_1x_1 + P_2x_2 + \dots + P_nx_n,$$

where the right-hand side of the equation is the amount you have to spend to buy that quantity of the first good (the quantity times its price--3 pounds of apples at \$1/pound equals \$3), plus the amount for the quantity you are buying of the second good, plus

The point I made above about equivalent utility functions can be made more general by observing that if there are two functions, $u(x)$, $v(x)$, and if for any two bundles of goods, x_a , x_b , whenever $u(x_a) > u(x_b)$ then $v(x_a) > v(x_b)$ and vice versa, then the two utility functions describe exactly the same behavior and are equivalent. The purpose of a utility function is to tell which bundle of goods I prefer (the one for which the utility function gives a higher utility). Two different functions that always give the same answer to that question are equivalent--they imply exactly the same preferences.

My income and the prices of the goods I want define the alternatives from which I can choose; my utility function defines my preferences. Mathematically speaking, the problem of consumption is simply the problem of choosing the bundle of goods that maximizes your utility, subject to the income constraint--the requirement that the bundle you choose cost no more than your income. This, of course, is what we were doing earlier in the chapter. The utility function simply provides a more mathematically precise way of talking about it.

Calculus

We have a utility function $u(x,y,z, \dots)$ depending on the amount consumed of goods X, Y, Z , etc. We assume that the quantities x, y, z, \dots can be continuously varied; that for $x, y, z, \dots > 0$, u is a continuous function with continuous first derivatives, and that u is an increasing function of all its arguments (since they are goods, utility increases with increased consumption). u obeys the principle of declining marginal utility: du/dx decreases as x increases (and similarly for y, z , etc.), so $\frac{\partial^2 u}{\partial x^2} < 0$. Our problem is to maximize u subject to the income constraint:

$$I \geq xP_x + yP_y + zP_z + \dots$$

The general approach to solving such a problem (a constrained maximization) uses a mathematical device called a Lagrange multiplier, with which you may already be familiar. In this particular case, we can use a simpler and (to me) more intuitive approach. To begin with, note that \geq in the income constraint can be replaced by $=$; since the only thing money is good for is buying goods and since more goods are always preferred to fewer goods, there is never any reason to spend less than your entire income.

We now consider varying x and y , while holding fixed the quantities of all other goods. If utility is at a maximum, an infinitesimal increase in x combined (because of the income constraint) with an infinitesimal decrease in y (such that total expenditure on x and y is unchanged) must leave utility unchanged, or in other words:

$$0 = du(x, y, z, \dots) = \frac{\partial u}{\partial x} dx + \frac{\partial u}{\partial y} dy$$

From which it follows that:

$$0 = du(x, y, z, \dots) / dx = \left[\frac{\partial u}{\partial x} + \frac{\partial u}{\partial y} \frac{dy}{dx} \right] \text{ (Eqn. 1)}$$

To find dy/dx we solve the income constraint for y in terms of x then take the derivative, thus:

$$y = (I - zP_z - \dots - xP_x) / P_y$$

$$dy/dx = -(P_x / P_y).$$

Substituting this into Equation 1, we have

$$0 = \frac{\partial u}{\partial x} - (P_x / P_y) \frac{\partial u}{\partial y}$$

Rearranging this gives us

$$\frac{\partial u}{\partial x} / P_x = \frac{\partial u}{\partial y} / P_y$$

$$\frac{\frac{\partial u}{\partial x}}{\frac{\partial u}{\partial y}} = \frac{P_x}{P_y}$$

which is the same relation that we derived earlier in the chapter, when we concluded that the price of an apple measured in oranges (P_a/P_o) is equal to the value of an apple measured in oranges. $\frac{\partial u}{\partial x}$ is the marginal utility of x and $\frac{\partial u}{\partial y}$ the marginal utility of y ; their ratio is the value of x measured in y --the marginal rate of substitution. If a pound of X has a marginal utility of 3 and a pound of Y has a marginal utility of 1, then on the margin a pound of X is worth 3 pounds of Y . We could have made the same argument for X and Z instead of X and Y , or for any other pair of goods (holding consumption of everything else constant), so the equimarginal principle holds for all goods we consume.

It does not hold for goods we do not consume. As you may remember from a calculus course, the normal condition for a maximum, which is that the derivative is zero, does not apply if the maximum occurs at one end or the other of the variable's range. The situation is shown in Figure 3-14; $f(x)$, which is only defined for $x > 0$, has its maximum value at $x = 0$. Its derivative there is negative, but we cannot find higher values of f at lower values of x because there are no lower values of x . This is a corner solution; the maximum occurs at the corner (point A on Figure 3-11), where the function runs into the barrier at $x = 0$.




Figure 3-14

A corner solution. At A, $f(x)$ is maximum, but df/dx is not zero.

A corner solution arises in consumption if there is a good X such that your maximum utility occurs when you are consuming none of it: $x = 0$. Since it is a corner solution, the derivative of utility need not be zero even though utility is at a maximum; so Equation 1 need not hold. Put in words, that means that utility is still increasing as you decrease consumption of the good (and spend the money on other goods) up to the point where your consumption of X reaches 0. The marginal utility of X is less per dollar than the marginal utility of other goods, but you cannot increase your utility by consuming a dollar less of X and a dollar more of something else since you have already reduced your consumption of X to zero. So the equimarginal principle does not apply to goods you do not consume. This is the same point I made earlier and illustrated on Figures 3-3a and 3-3b. The picture of the corner is different, since it involves a utility function here and an indifference curve there, but the situation is the same.

Indifference Curves and Utility Functions

Next let us look at indifference curve analysis in terms of the utility function. Since we have only two dimensions, we will limit ourselves to a utility function with only two goods. Even in that case, showing two goods uses up the two dimensions we have available, leaving no place to show the utility function itself. With a third dimension, we could draw it as a surface, letting the height of the surface above any point (x,y) represent $u(x,y)$. Unfortunately this book is written on two-dimensional paper; Figure 3-15a is an attempt to overcome that limitation.

This is not a new problem; mapmakers face it whenever they try to represent a three-dimensional landscape on two-dimensional paper. The solution is a contour map. A contour map has one line through all points 100 feet above sea level, another through all points 200 feet above, and so on; by looking at the map you can, with practice, figure out the shape of the land in the third dimension. Where it is rising steeply, the contours are close together (the land rises 100 feet in only a short horizontal distance); where it is gently sloped, they are far apart.

The economist's equivalent of the contour on a topographical map is an indifference curve;

it represents all of the points among which you are indifferent--or in other words, all of the bundles that give you the same utility. Figure 3-15a shows indifference curves U_1 , U_2 , and U_3 , each labeled by its utility. Since U_1 is less than U_2 , which in turn is less than U_3 , points on U_3 are preferred to points on U_2 , which are in turn preferred to points on U_1 . The X-Y plane of Figure 3-15a corresponds to the indifference curve diagrams done earlier in the chapter.

Indifference curves do not completely describe the utility function from which they come. A curve is not labeled; it does not say "utility equals 9" on it. All the indifference curves tell us is which bundles we are indifferent among and which we prefer to which. They are in this sense less informative than the lines on a contour map, which tell us not only where the contour is but which contour (how many feet above sea level) it is. Hence one set of indifference curves may correspond to many different utility functions. The fact that, in spite of that, we can analyze consumption completely in terms of indifference curves corresponds to a point I made earlier--that different utility functions may describe exactly the same behavior.

[I have omitted Figure 3-15 from the webbed version of this chapter; my publisher holds copyright on the figures that he had drawn and I'm not up to doing 3-D perspective. I will try to get permission to scan it in Real Soon Now. Use your imagination.]

As we move along an indifference curve, utility stays the same. Suppose (x,y) and $(x + dx, y + dy)$ are two points on the same indifference curve. We have:

$$u(x,y) = u(x + dx, y + dy) \cong u(x,y) + dx \frac{\partial u}{\partial x} + dy \frac{\partial u}{\partial y}$$

As $dy, dx \rightarrow 0$, their ratio becomes the slope of the indifference curve, and the approximate equality becomes an equality. In that case,

$dx \frac{\partial u}{\partial x} + dy \frac{\partial u}{\partial y} = 0$, and

$-(\frac{\partial u}{\partial x}) / \frac{\partial u}{\partial y} = dy/dx = \text{slope of the indifference curve.}$

This is equivalent, as you should be able to show, to the conclusion we reached earlier--that minus the slope of the indifference curve was equal to the value of apples (X) measured in oranges (Y).

If in Figure 3-15b, like the indifference curves discussed earlier, slopes down to the right; its slope is negative. To keep utility constant, a reduction in the amount of one good must be balanced by an increase in the amount of another. Indifference curves sloping the other way would describe your preferences between two things, one of which is a good and one a bad--something for which $\frac{\partial u}{\partial y} > 0$. If this seems an odd thing to graph, consider representing your utility as a function of number of hours worked and number of dollars of income, the first a bad and the second a good, and deducing how many hours you will work at any given wage. Or consider the situation where production of a good results in undesirable waste products.

The slope of an indifference curve is usually negative because we are usually representing preferences between two goods. Its curvature, the fact that the slope of the indifference curves becomes shallower (i.e., less negative) as you move right or down on the diagram and steeper as you move left or up, is suggested by the principle of declining marginal utility but is not, strictly speaking, implied by it. Imagine that you move from *A* to *B* on Figure 3-15b. Quantity of *Y* stays the same and quantity of *X* increases, so $\frac{\partial u}{\partial x}$ must decrease. The slope of the indifference curve is $-(\frac{\partial u}{\partial x}) / (\frac{\partial u}{\partial y})$, so the slope of the indifference curve through *B* is shallower than the slope of the indifference curve through *A*--unless $\frac{\partial u}{\partial y}$ decreases even faster than $\frac{\partial u}{\partial x}$ as *x* increases. There is no obvious reason why it should, but nothing in our assumptions makes it impossible. Similarly, as you move from *C* to *D*, *y* increases, *x* stays the same, $\frac{\partial u}{\partial y}$ decreases, and the slope of the indifference curves becomes steeper--unless, for some reason, an increase in the quantity of *Y* decreases the marginal utility of *X* even faster than it decreases the marginal utility of *Y*.

Here again, as several times before, our analysis is complicated by the possibility that consumption of one good may affect the utility of another. In most real-world situations, we would not expect such effects to be very large--we consume many different goods, most of which have little to do with each other. The exceptions are pairs of closely related goods--cars and bicycles, bread and butter, bananas and peanut butter. In some of these cases (*substitutes*, such as cars and bicycles) the more we have of one good the less we value the other; in other cases (*complements*, such as bread and butter or gasoline and automobiles) the more we have of one the more we value the other.

In such cases, we may expect indifference curves to be oddly shaped--some examples are given in the problems at the end of this chapter. In most other cases, we assume the *principle of declining marginal rate of substitution*--which means that the slope of the indifference curves becomes shallower as we move to the right on the diagram and steeper as we move up. As you can see from the previous discussion, this is close enough to the principle of declining marginal utility that for most practical purposes we may think of them as the same.

We have now derived the equimarginal principle directly from utility functions and shown the connection between utility functions and indifference curves. It is worth noting that although the argument was made in terms of money income and money prices, money has nothing essential to do with it. We could just as easily have started with a bundle of goods (x, y, \dots), and allowed you to exchange X for Y at a price (of Y in terms of X) of P_y/P_x , for Z at a price of P_z/P_x , and so on. Here, as elsewhere in price theory, the use of money and money prices simplifies exposition but does not affect the conclusions.

PROBLEMS

1. Near the beginning of the chapter, I gave some examples of bads. Do you agree with them? If not, is one of us necessarily wrong? Discuss.

2. Suppose my preferences with regard to hamburger and pens are as shown.

Options	Hamburgers (pounds/year)	Pens/year	Utility (utiles/year)
A	100	30	50

B	108	29	50
C	118	28	50
D	200	30	75
E	?	29	75

- a. What is the value of a pound of hamburger to me (between points *A* and *B*)?
 - b. In choosing between bundles *A* and *B*, which do I prefer? Between *C* and *D*?
 - c. About how much hamburger should be in *E* to make me indifferent between it and *D*? Explain briefly.
3. Figure 3-16 shows your preferences between brandy and champagne. Which (if any) of the bundles shown do you prefer to point *A*? To which is *A* preferred? Which are equivalent to *A*? For which bundles can you not tell whether they are equivalent, better, or worse than *A*?
 4. Answer the same questions for point *B*.
 5. Figure 3-17 shows your indifference curves for cookies and bananas. You have an income of \$100, the price of cookies is \$1, and the price of bananas is \$0.25. How many of each do you choose to consume?
 6. Figure 3-18*a* shows a set of indifference curves; Figure 3-18*b* shows a set of budget lines. Your income is \$12/week, the price of good *X* is \$2, and the price of good *Y* is \$4.
 - a. Which line on Figure 3-18*b* is your budget line?
 - b. Which point on Figure 3-18*a* do you prefer, among those available to you? In other words, how much of *X* and of *Y* do you choose to consume?
-




Figure 3-16




Figure 3-17





Figure 3-18

7. Figures 3-19a, b, c, and d show four different sets of indifference curves; in each case, points on U_3 are preferred to points on U_2 , and points on U_2 are preferred to points on U_1 . Describe verbally the pattern of preferences illustrated in each case. Yes, they are odd.

8. Figure 3-19e shows your preferences with regard to two goods--left shoes and right shoes. Explain why the indifference curves have the shape shown.

9. Draw a possible set of indifference curves for two things that are close, but not perfect, complements. An example might be bread and butter, if you much prefer your bread with butter but are willing to eat bread without butter (or with less than your preferred amount of butter).

10. Draw a possible set of indifference curves for two things that are perfect substitutes--butter and margarine for someone who cannot tell them apart. Draw another set for two things that are close, but not perfect, substitutes. An example might be chicken and turkey, if for some recipes you mildly prefer one or the other.

11. Figure 3-20 shows an indifference curve map and a budget line.

a. What is your marginal rate of substitution at points *A*, *B*, *C* ?

b. What is the slope of the budget line at points *A*, *B*, *C* ?







Figure 3-19





12. Use Figure 3-4 to derive an income-adjusted demand curve for Apples; B_I on the figure should be one of your budget lines.
13. Use Figure 3-4 to derive an ordinary demand curve for Apples; assume that your income is \$100 and the price of oranges is \$1.
14. William's income is \$3/day; apples cost \$0.50/apple.
- Draw William's budget line, showing the choice between apples and expenditure on all other goods.
 - In order to reduce medical expenditures, the government decides to subsidize apples; for every dollar William spends on apples, he will be given \$0.25 back. William pays no taxes. Draw William's budget line .
- C. Instead of a subsidy, the government decides to use a voucher. The government provides William with \$0.50/day which can only be spent on apples; any part of it that he does not spend on apples must be returned. William still pays no taxes. Draw his budget line. Be careful; it will not look like the other budget lines we have drawn.
15. The situation is the same as in the previous problem. Figure 3-21 shows William's indifference curves. How many apples a day does William consume:

- a. With neither subsidy nor voucher?
- b. With subsidy?
- c. With voucher?



16. Suppose that instead of subsidizing potatoes, as discussed in the text, we tax them; for every \$2 you spend on potatoes, you must give an additional \$1 to the government. The tax collected is then returned to the consumers as a *demogrant*: everyone gets a fixed number of dollars to add to his income. We assume that everyone has the same income and the same tastes.

Would people be better or worse off than if there were no tax (and no subsidy)? Prove your answer.

The following problems refer to the optional section:

17. What testable proposition is suggested by the statement "A has more utility than B to me?"

18. Do each of a-d, *both* geometrically (you need not be precise) and using calculus. There are only two goods; x is the quantity of one good and y of the other. Your income is I . $u(x,$

$$y) = xy + x + y.$$

- a. $P_x = \$1$; $P_y = \$1$; $I = \$10$. Suppose P_y rises to $\$2$. By how much must I increase in order that you be as well off as before?
- b. In the case described in part (a), assuming that I does not change, what quantities of each good are consumed before and after the price change? How much of each change is a substitution effect? How much is an income effect?
- c. $P_x = \$1$; $I = \$10$. Graph the amount of Y you consume as a function of P_y , for values of P_y ranging from $\$0$ to $\$10$ (your ordinary demand curve for Y).
- d. With both prices equal to $\$1$, show how consumption of each good varies as I changes from $\$0$ to $\$100$.

19. Answer the following questions for the utility function:

$$u(x,y) = x - 1/y$$

- a. $P_x = \$1$; $I = \$10$. Draw the demand curve for good Y , $\$1 < P_y < \100 .
- b. $P_x = \$1$; $I = \$10$. P_y increases from $\$1$ to $\$2$. Show the old and the new equilibria. The income effect could be eliminated either by changing I or by changing P_x and P_y while keeping their ratio fixed. What would the necessary change in I be? What would the necessary change in the prices be? Diagram both.
- c. $P_x = \$1$. Draw the income-compensated demand curve for good Y . $\$1 < P_y < \100 . Start with $P_y = \$1$ and $I = \$10$.
- d. $P_x = \$1 = P_y$. Graph y against I for $\$0 < I < \10 .

[Next Chapter](#)

[Table of Contents](#)