

## 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



## 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.