CHAPTER 1

THE MARKET

The conventional first chapter of a microeconomics book is a discussion of the "scope and methods" of economics. Although this material can be very interesting, it hardly seems appropriate to *begin* your study of economics with such material. It is hard to appreciate such a discussion until you have seen some examples of economic analysis in action.

So instead, we will begin this book with an *example* of economic analysis. In this chapter we will examine a model of a particular market, the market for apartments. Along the way we will introduce several new ideas and tools of economics. Don't worry if it all goes by rather quickly. This chapter is meant only to provide a quick overview of how these ideas can be used. Later on we will study them in substantially more detail.

1.1 Constructing a Model

Economics proceeds by developing **models** of social phenomena. By a model we mean a simplified representation of reality. The emphasis here is on the word "simple." Think about how useless a map on a one-to-one scale would be. The same is true of an economic model that attempts to describe every aspect of reality. A model's power stems from the elimination of irrelevant detail, which allows the economist to focus on the essential features of the economic reality he or she is attempting to understand.

Here we are interested in what determines the price of apartments, so we want to have a simplified description of the apartment market. There is a certain art to choosing the right simplifications in building a model. In general we want to adopt the simplest model that is capable of describing the economic situation we are examining. We can then add complications one at a time, allowing the model to become more complex and, we hope, more realistic.

The particular example we want to consider is the market for apartments in a medium-size midwestern college town. In this town there are two sorts of apartments. There are some that are adjacent to the university, and others that are farther away. The adjacent apartments are generally considered to be more desirable by students, since they allow easier access to the university. The apartments that are farther away necessitate taking a bus, or a long, cold bicycle ride, so most students would prefer a nearby apartment ... if they can afford one.

We will think of the apartments as being located in two large rings surrounding the university. The adjacent apartments are in the inner ring, while the rest are located in the outer ring. We will focus exclusively on the market for apartments in the inner ring. The outer ring should be interpreted as where people can go who don't find one of the closer apartments. We'll suppose that there are many apartments available in the outer ring, and their price is fixed at some known level. We'll be concerned solely with the determination of the price of the inner-ring apartments and who gets to live there.

An economist would describe the distinction between the prices of the two kinds of apartments in this model by saying that the price of the outer-ring apartments is an **exogenous variable**, while the price of the inner-ring apartments is an **endogenous variable**. This means that the price of the outer-ring apartments is taken as determined by factors not discussed in this particular model, while the price of the inner-ring apartments is determined by forces described in the model.

The first simplification that we'll make in our model is that all apartments are identical in every respect except for location. Thus it will make sense to speak of "the price" of apartments, without worrying about whether the apartments have one bedroom, or two bedrooms, or whatever.

But what determines this price? What determines who will live in the inner-ring apartments and who will live farther out? What can be said about the desirability of different economic mechanisms for allocating apartments? What concepts can we use to judge the merit of different assignments of apartments to individuals? These are all questions that we want our model to address

1.2 Optimization and Equilibrium

Whenever we try to explain the behavior of human beings we need to have a framework on which our analysis can be based. In much of economics we use a framework built on the following two simple principles.

The optimization principle: People try to choose the best patterns of consumption that they can afford.

The equilibrium principle: Prices adjust until the amount that people demand of something is equal to the amount that is supplied.

Let us consider these two principles. The first is *almost* tautological. If people are free to choose their actions, it is reasonable to assume that they try to choose things they want rather than things they don't want. Of course there are exceptions to this general principle, but they typically lie outside the domain of economic behavior.

The second notion is a bit more problematic. It is at least conceivable that at any given time peoples' demands and supplies are not compatible, and hence something must be changing. These changes may take a long time to work themselves out, and, even worse, they may induce other changes that might "destabilize" the whole system.

This kind of thing can happen ... but it usually doesn't. In the case of apartments, we typically see a fairly stable rental price from month to month. It is this *equilibrium* price that we are interested in, not in how the market gets to this equilibrium or how it might change over long periods of time.

It is worth observing that the definition used for equilibrium may be different in different models. In the case of the simple market we will examine in this chapter, the demand and supply equilibrium idea will be adequate for our needs. But in more general models we will need more general definitions of equilibrium. Typically, equilibrium will require that the economic agents' actions must be consistent with each other.

How do we use these two principles to determine the answers to the questions we raised above? It is time to introduce some economic concepts.

1.3 The Demand Curve

Suppose that we consider all of the possible renters of the apartments and ask each of them the maximum amount that he or she would be willing to pay to rent one of the apartments.

Let's start at the top. There must be someone who is willing to pay the highest price. Perhaps this person has a lot of money, perhaps he is very lazy and doesn't want to walk far ... or whatever. Suppose that this person is willing to pay \$500 a month for an apartment.

If there is only one person who is willing to pay \$500 a month to rent an apartment, then if the price for apartments were \$500 a month, exactly one apartment would be rented—to the one person who was willing to pay that price.

Suppose that the next highest price that anyone is willing to pay is \$490. Then if the market price were \$499, there would still be only one apartment rented: the person who was *willing* to pay \$500 would rent an apartment, but the person who was willing to pay \$490 wouldn't. And so it goes. Only one apartment would be rented if the price were \$498, \$497, \$496, and so on ... until we reach a price of \$490. At that price, exactly two apartments would be rented: one to the \$500 person and one to the \$490 person.

Similarly, two apartments would be rented until we reach the maximum price that the person with the *third* highest price would be willing to pay, and so on.

Economists call a person's maximum willingness to pay for something that person's **reservation price**. The reservation price is the highest price that a given person will accept and still purchase the good. In other words, a person's reservation price is the price at which he or she is just indifferent between purchasing or not purchasing the good. In our example, if a person has a reservation price p it means that he or she would be just indifferent between living in the inner ring and paying a price p and living in the outer ring.

Thus the number of apartments that will be rented at a given price p^* will just be the number of people who have a reservation price greater than or equal to p^*. For if the market price is p^*, then everyone who is willing to pay at least p^* for an apartment will want an apartment in the inner ring, and everyone who is not willing to pay p^* will choose to live in the outer ring.

We can plot these reservation prices in a diagram as in Figure 1.1. Here the price is depicted on the vertical axis and the number of people who are willing to pay that price or more is depicted on the horizontal axis.

Another way to view Figure 1.1 is to think of it as measuring how many people would want to rent apartments at any particular price. Such a curve is an example of a **demand curve**—a curve that relates the quantity demanded to price. When the market price is above \$500, zero apartments will be rented. When the price is between \$500 and \$490, one apartment will be rented. When it is between \$490 and the third highest reservation price, two apartments will be rented, and so on. The demand curve describes the quantity demanded at each of the possible prices.

The demand curve for apartments slopes down: as the price of apartments decreases more people will be willing to rent apartments. If there are many people and their reservation prices differ only slightly from person to

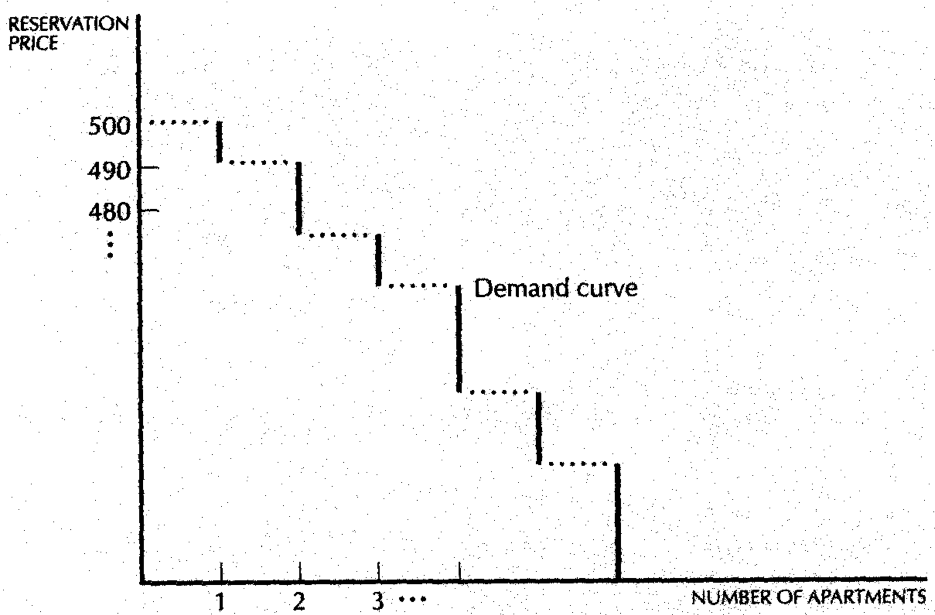

The demand curve for apartments. The vertical axis measures the market price and the horizontal axis measures how many apartments will be rented at each price.

Figure 1.1

person, it is reasonable to think of the demand curve as sloping smoothly downward, as in Figure 1.2. The curve in Figure 1.2 is what the demand curve in Figure 1.1 would look like if there were many people who want to rent the apartments. The "jumps" shown in Figure 1.1 are now so small relative to the size of the market that we can safely ignore them in drawing the market demand curve.

1.4 The Supply Curve

We now have a nice graphical representation of demand behavior, so let us turn to supply behavior. Here we have to think about the nature of the market we are examining. The situation we will consider is where there are many independent landlords who are each out to rent their apartments for the highest price the market will bear. We will refer to this as the case of a **competitive market**. Other sorts of market arrangements are certainly possible, and we will examine a few later.

For now, let's consider the case where there are many landlords who all operate independently. It is clear that if all landlords are trying to do the best they can and if the renters are fully informed about the prices the landlords charge, then the equilibrium price of all apartments in the inner ring must be the same. The argument is not difficult. Suppose instead that there is some high price, p_h, and some low price, p_l, being charged

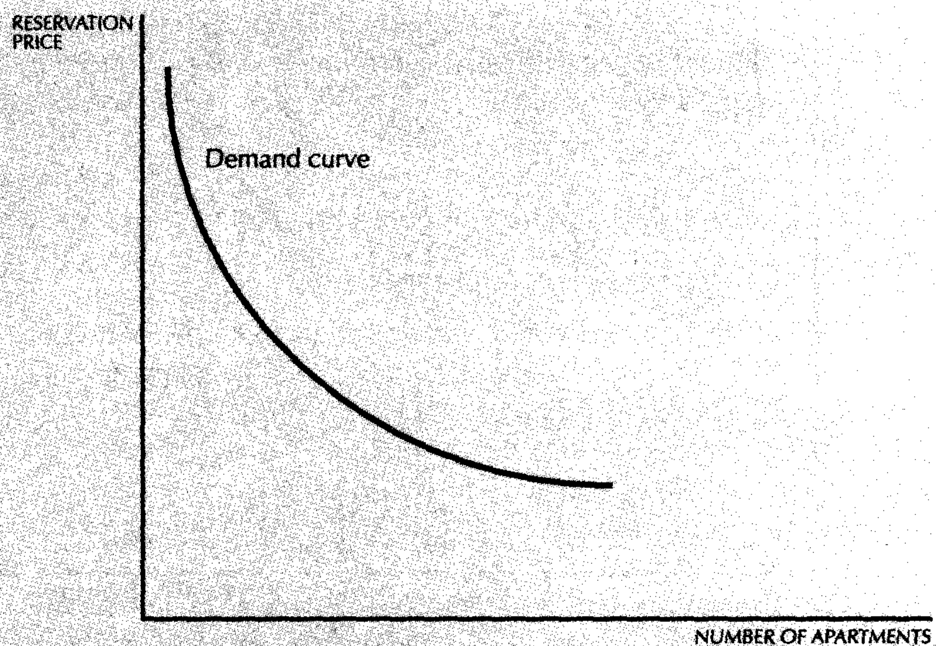

Demand curve for apartments with many demanders. Because of the large number of demanders, the jumps between prices will be small, and the demand curve will have the conventional smooth shape.

for apartments. The people who are renting their apartments for a high price could go to a landlord renting for a low price and offer to pay a rent somewhere between p_h and p_l. A transaction at such a price would make both the renter and the landlord better off. To the extent that all parties are seeking to further their own interests and are aware of the alternative prices being charged, a situation with different prices being charged for the same good cannot persist in equilibrium.

But what will this single equilibrium price be? Let us try the method that we used in our construction of the demand curve: we will pick a price and ask how many apartments will be supplied at that price.

The answer depends to some degree on the time frame in which we are examining the market. If we are considering a time frame of several years, so that new construction can take place, the number of apartments will certainly respond to the price that is charged. But in the "short run"—within a given year, say—the number of apartments is more or less fixed. If we consider only this short-run case, the supply of apartments will be constant at some predetermined level.

The **supply curve** in this market is depicted in Figure 1.3 as a vertical line. Whatever price is being charged, the same number of apartments will be rented, namely, all the apartments that are available at that time.

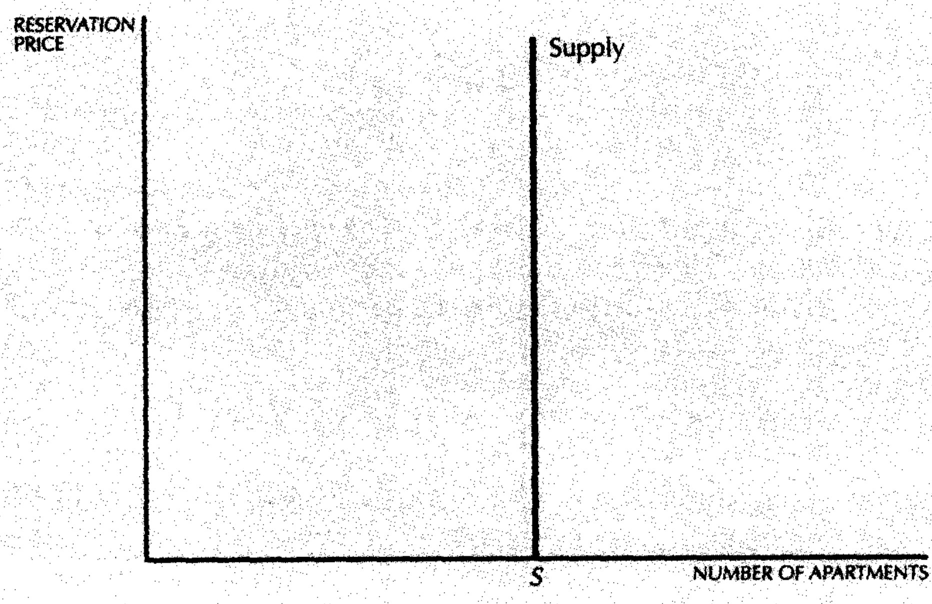

Short-run supply curve. The supply of apartments is fixed in the short run.

1.5 Market Equilibrium

We now have a way of representing the demand and the supply side of the apartment market. Let us put them together and ask what the equilibrium behavior of the market is. We do this by drawing both the demand and the supply curve on the same graph in Figure 1.4.

In this graph we have used p^* to denote the price where the quantity of apartments demanded equals the quantity supplied. This is the **equilibrium price** of apartments. At this price, each consumer who is willing to pay at least p^* is able to find an apartment to rent, and each landlord will be able to rent apartments at the going market price. Neither the consumers nor the landlords have any reason to change their behavior. This is why we refer to this as an *equilibrium:* no change in behavior will be observed.

To better understand this point, let us consider what would happen at a price other than p^*. For example, consider some price $p < p^*$ where demand is greater than supply. Can this price persist? At this price at least some of the landlords will have more renters than they can handle. There will be lines of people hoping to get an apartment at that price; there are more people who are willing to pay the price p than there are apartments. Certainly some of the landlords would find it in their interest to raise the price of the apartments they are offering.

Similarly, suppose that the price of apartments is some p greater than p^*.

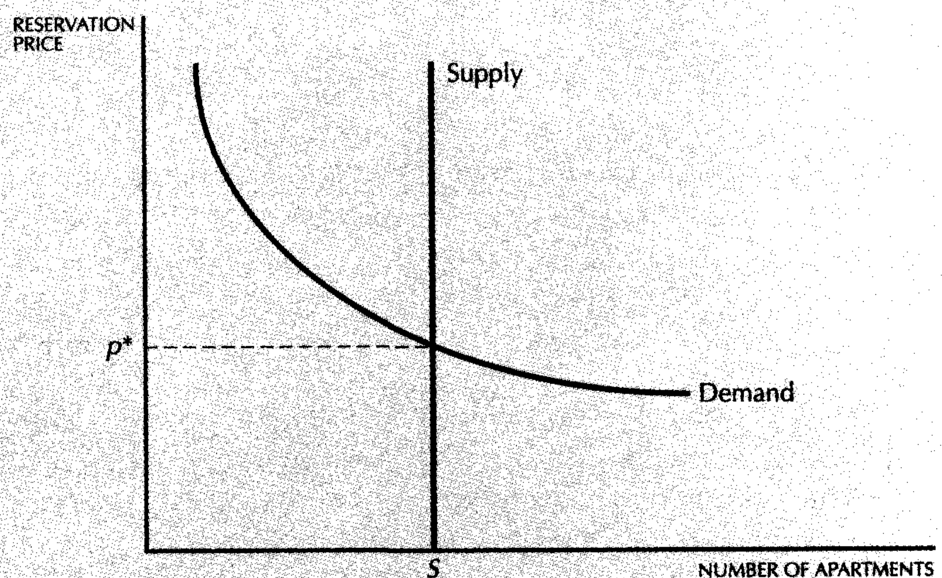

Equilibrium in the apartment market. The equilibrium price, p^*, is determined by the intersection of the supply and demand curves.

Then some of the apartments will be vacant: there are fewer people who are willing to pay p than there are apartments. Some of the landlords are now in danger of getting no rent at all for their apartments. Thus they will have an incentive to lower their price in order to attract more renters.

If the price is above p^* there are too few renters; if it is below p^* there are too many renters. Only at the price of p^* is the number of people who are willing to rent at that price equal to the number of apartments available for rent. Only at that price does demand equal supply.

At the price p^* the landlords' and the renters' behaviors are compatible in the sense that the number of apartments demanded by the renters at p^* is equal to the number of apartments supplied by the landlords. This is the equilibrium price in the market for apartments.

Once we've determined the market price for the inner-ring apartments, we can ask who ends up getting these apartments and who is exiled to the farther-away apartments. In our model there is a very simple answer to this question: in the market equilibrium everyone who is willing to pay p^* or more gets an apartment in the inner ring, and everyone who is willing to pay less than p^* gets one in the outer ring. The person who has a reservation price of p^* is just indifferent between taking an apartment in the inner ring and taking one in the outer ring. The other people in the inner ring are getting their apartments at less than the maximum they would be willing to pay for them. Thus the assignment of apartments to renters is determined by how much they are willing to pay.

1.6 Comparative Statics

Now that we have an economic model of the apartment market, we can begin to use it to analyze the behavior of the equilibrium price. For example, we can ask how the price of apartments changes when various aspects of the market change. This kind of an exercise is known as **comparative statics**, because it involves comparing two "static" equilibria without worrying about how the market moves from one equilibrium to another.

The movement from one equilibrium to another can take a substantial amount of time, and questions about how such movement takes place can be very interesting and important. But we must walk before we can run, so we will ignore such dynamic questions for now. Comparative statics analysis is only concerned with comparing equilibria, and there will be enough questions to answer in this framework for the present.

Let's start with a simple case. Suppose that the supply of apartments is increased, as in Figure 1.5.

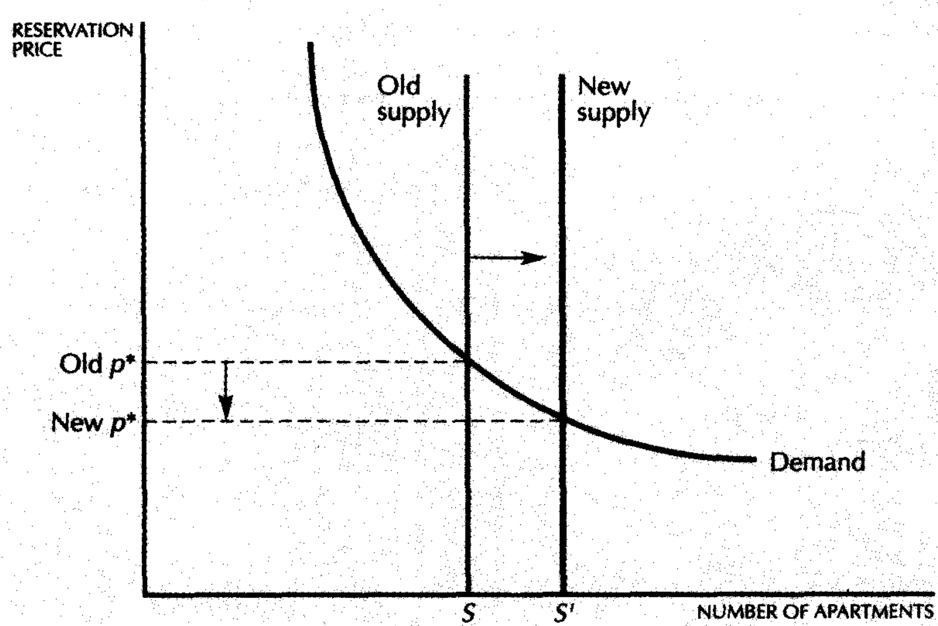

Increasing the supply of apartments. As the supply of apartments increases, the equilibrium price decreases.

It is easy to see in this diagram that the equilibrium price of apartments will fall. Similarly, if the supply of apartments were reduced the equilibrium price would rise.

Let's try a more complicated—and more interesting—example. Suppose that a developer decides to turn several of the apartments into condominiums. What will happen to the price of the remaining apartments?

Your first guess is probably that the price of apartments will go up, since the supply has been reduced. But this isn't necessarily right. It is true that the supply of apartments to rent has been reduced. But the *demand for apartments* has been reduced as well, since some of the people who were renting apartments may decide to purchase the new condominiums.

It is natural to assume that the condominium purchasers come from those who already live in the inner-ring apartments—those people who are willing to pay more than p^* for an apartment. Suppose, for example, that the demanders with the 10 highest reservation prices decide to buy condos rather than rent apartments. Then the new demand curve is just the old demand curve with 10 fewer demanders at each price. Since there are also 10 fewer apartments to rent, the new equilibrium price is just what it was before, and exactly the same people end up living in the inner-ring apartments. This situation is depicted in Figure 1.6. Both the demand curve and the supply curve shift left by 10 apartments, and the equilibrium price remains unchanged.

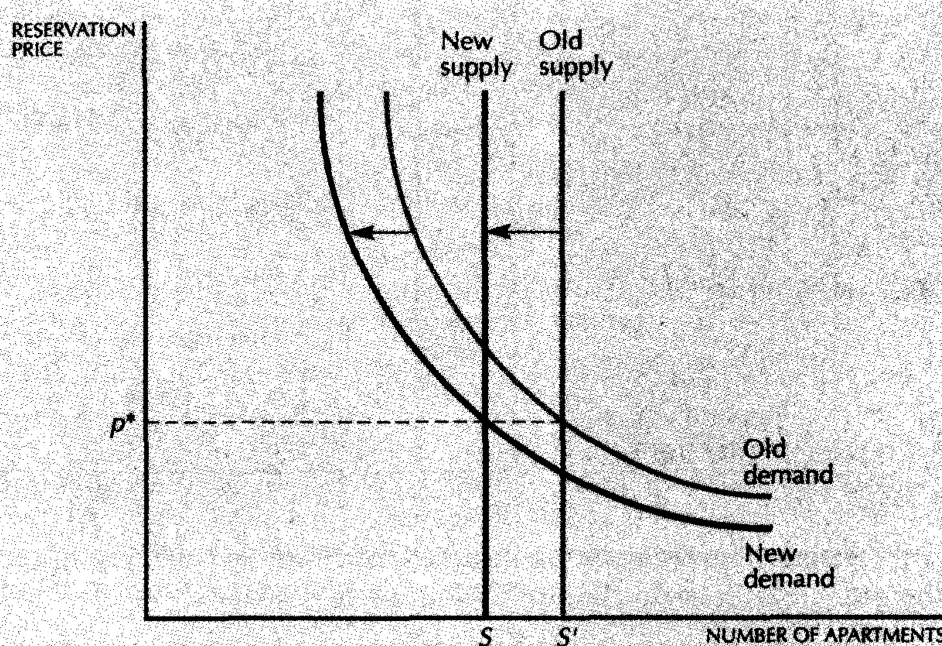

Effect of creating condominiums. If demand and supply both shift left by the same amount the equilibrium price is unchanged.

Most people find this result surprising. They tend to see just the reduction in the supply of apartments and don't think about the reduction in demand. The case we've considered is an extreme one: *all* of the condo purchasers were former apartment dwellers. But the other case—where none of the condo purchasers were apartment dwellers—is even more extreme.

The model, simple though it is, has led us to an important insight. If we want to determine how conversion to condominiums will affect the apartment market, we have to consider not only the effect on the supply of apartments but also the effect on the demand for apartments.

Let's consider another example of a surprising comparative statics analysis: the effect of an apartment tax. Suppose that the city council decides that there should be a tax on apartments of \$50 a year. Thus each landlord will have to pay \$50 a year to the city for each apartment that he owns. What will this do to the price of apartments?

Most people would think that at least some of the tax would get passed along to apartment renters. But, rather surprisingly, that is not the case. In fact, the equilibrium price of apartments will remain unchanged!

In order to verify this, we have to ask what happens to the demand curve and the supply curve. The supply curve doesn't change—there are just as many apartments after the tax as before the tax. And the demand curve doesn't change either, since the number of apartments that will be rented at each different price will be the same as well. If neither the demand curve nor the supply curve shifts, the price can't change as a result of the tax.

Here is a way to think about the effect of this tax. Before the tax is imposed, each landlord is charging the highest price that he can get that will keep his apartments occupied. The equilibrium price p^* is the highest price that can be charged that is compatible with all of the apartments being rented. After the tax is imposed can the landlords raise their prices to compensate for the tax? The answer is no: if they could raise the price and keep their apartments occupied, they would have already done so. If they were charging the maximum price that the market could bear, the landlords couldn't raise their prices any more: none of the tax can get passed along to the renters. The landlords have to pay the entire amount of the tax.

This analysis depends on the assumption that the supply of apartments remains fixed. If the number of apartments can vary as the tax changes, then the price paid by the renters will typically change. We'll examine this kind of behavior later on, after we've built up some more powerful tools for analyzing such problems.

1.7 Other Ways to Allocate Apartments

In the previous section we described the equilibrium for apartments in a competitive market. But this is only one of many ways to allocate a resource; in this section we describe a few other ways. Some of these may sound rather strange, but each will illustrate an important economic point.

The Discriminating Monopolist

First, let us consider a situation where there is one dominant landlord who owns all of the apartments. Or, alternatively, we could think of a number of individual landlords getting together and coordinating their actions to act as one. A situation where a market is dominated by a single seller of a product is known as a **monopoly**.

In renting the apartments the landlord could decide to auction them off one by one to the highest bidders. Since this means that different people would end up paying different prices for apartments, we will call this the case of the **discriminating monopolist**. Let us suppose for simplicity that the discriminating monopolist knows each person's reservation price for apartments. (This is not terribly realistic, but it will serve to illustrate an important point.)

This means that he would rent the first apartment to the fellow who would pay the most for it, in this case \$500. The next apartment would go for \$490 and so on as we moved down the demand curve. Each apartment would be rented to the person who was willing to pay the most for it.

Here is the interesting feature of the discriminating monopolist: *exactly the same people will get the apartments as in the case of the market solution,* namely, everyone who valued an apartment at more than p^*. The last person to rent an apartment pays the price p^*—the same as the equilibrium price in a competitive market. The discriminating monopolist's attempt to maximize his own profits leads to the same allocation of apartments as the supply and demand mechanism of the competitive market. The amount the people *pay* is different, but who gets the apartments is the same. It turns out that this is no accident, but we'll have to wait until later to explain the reason.

The Ordinary Monopolist

We assumed that the discriminating monopolist was able to rent each apartment at a different price. But what if he were forced to rent all apartments at the same price? In this case the monopolist faces a tradeoff: if he chooses a low price he will rent more apartments, but he may end up making less money than if he sets a higher price.

Let us use $D(p)$ to represent the demand function—the number of apartments demanded at price p. Then if the monopolist sets a price p, he will rent $D(p)$ apartments and thus receive a revenue of $pD(p)$. The revenue that the monopolist receives can be thought of as the area of a box: the height of the box is the price p and the width of the box is the number of apartments $D(p)$. The product of the height and the width—the area of the box—is the revenue the monopolist receives. This is the box depicted in Figure 1.7.

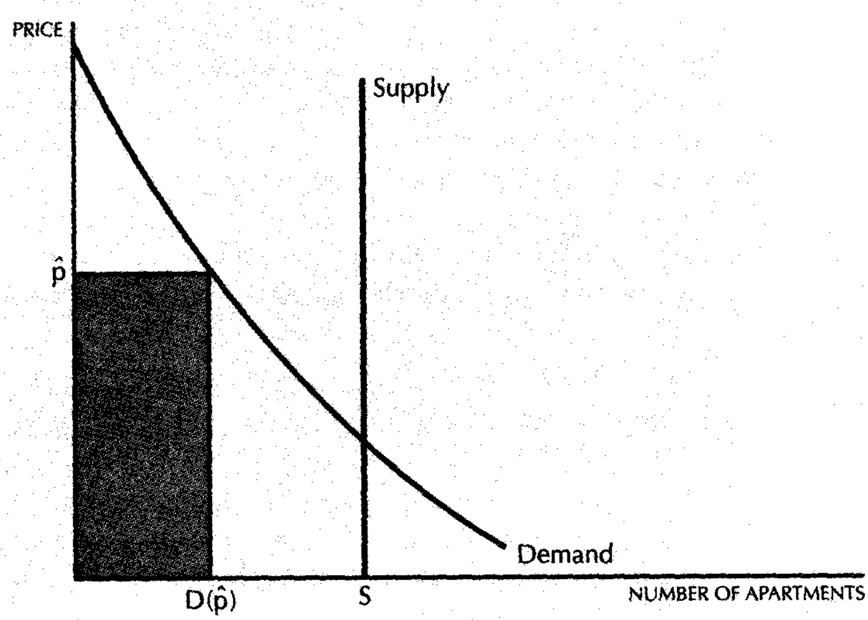

Revenue box. The revenue received by the monopolist is just the price times the quantity, which can be interpreted as the area of the box illustrated.

If the monopolist has no costs associated with renting an apartment, he would want to choose a price that has the largest associated revenue box. The largest revenue box in Figure 1.7 occurs at the price $\hat{p}$. In this case the monopolist will find it in his interest *not* to rent all of the apartments. In fact this will generally be the case for a monopolist. The monopolist will want to restrict the output available in order to maximize his profit. This means that the monopolist will generally want to charge a price that is higher than the equilibrium price in a competitive market, p^*. In the case of the ordinary monopolist, fewer apartments will be rented, and each apartment will be rented at a higher price than in the competitive market.

Rent Control

A third and final case that we will discuss will be the case of rent control. Suppose that the city decides to impose a maximum rent that can be charged for apartments, say p_{max}. We suppose that the price p_{max} is less than the equilibrium price in the competitive market, p^*. If this is so we would have a situation of **excess demand**: there are more people who are willing to rent apartments at p_{max} than there are apartments available. Who will end up with the apartments?

The theory that we have described up until now doesn't have an answer to this question. We can describe what will happen when supply equals demand, but we don't have enough detail in the model to describe what will happen if supply doesn't equal demand. The answer to who gets the apartments under rent control depends on who has the most time to spend looking around, who knows the current tenants, and so on. All of these things are outside the scope of the simple model we've developed. It may be that exactly the same people get the apartments under rent control as under the competitive market. But that is an extremely unlikely outcome. It is much more likely that some of the formerly outer-ring people will end up in some of the inner-ring apartments and thus displace the people who would have been living there under the market system. So under rent control the same number of apartments will be rented at the rent-controlled price as were rented under the competitive price: they'll just be rented to different people.

1.8 Which Way Is Best?

We've now described four possible ways of allocating apartments to people:

- The competitive market.
- A discriminating monopolist.
- An ordinary monopolist.
- Rent control.

These are four different economic institutions for allocating apartments. Each method will result in different people getting apartments or in different prices being charged for apartments. We might well ask which economic institution is best. But first we have to define "best." What criteria might we use to compare these ways of allocating apartments?

One thing we can do is to look at the economic positions of the people involved. It is pretty obvious that the owners of the apartments end up with the most money if they can act as discriminating monopolists: this would generate the most revenues for the apartment owner(s). Similarly the rent-control solution is probably the worst situation for the apartment owners.

What about the renters? They are probably worse off on average in the case of a discriminating monopolist—most of them would be paying a higher price than they would under the other ways of allocating apartments.

Are the consumers better off in the case of rent control? Some of them are: the consumers *who end up getting the apartments* are better off than they would be under the market solution. But the ones who didn't get the apartments are *worse off* than they would be under the market solution.

What we need here is a way to look at the economic position of all the parties involved—all the renters *and* all the landlords. How can we examine the desirability of different ways to allocate apartments, taking everybody into account? What can be used as a criterion for a "good" way to allocate apartments taking into account *all* of the parties involved?

1.9 Pareto Efficiency

One useful criterion for comparing the outcomes of different economic institutions is a concept known as Pareto efficiency or economic efficiency.[1] We start with the following definition: if we can find a way to make some people better off without making anybody else worse off, we have a **Pareto improvement**. If an allocation allows for a Pareto improvement, it is called **Pareto inefficient**; if an allocation is such that no Pareto improvements are possible, it is called **Pareto efficient**.

A Pareto inefficient allocation has the undesirable feature that there is some way to make somebody better off without hurting anyone else. There may be other positive things about the allocation, but the fact that it is Pareto inefficient is certainly one strike against it. If there is a way to make someone better off without hurting anyone else, why not do it?

The idea of Pareto efficiency is an important one in economics and we will examine it in some detail later on. It has many subtle implications that we will have to investigate more slowly, but we can get an inkling of what is involved even now.

Here is a useful way to think about the idea of Pareto efficiency. Suppose that we assigned the renters to the inner- and outer-ring apartments randomly, but then allowed them to sublet their apartments to each other. Some people who really wanted to live close in might, through bad luck, end up with an outer-ring apartment. But then they could sublet an inner-ring apartment from someone who was assigned to such an apartment but who didn't value it as highly as the other person. If individuals were assigned randomly to apartments, there would generally be some who would want to trade apartments, if they were sufficiently compensated for doing so.

For example, suppose that person A is assigned an apartment in the inner ring that he feels is worth \$200, and that there is some person B in the outer ring who would be willing to pay \$300 for A's apartment. Then there is a

[1] Pareto efficiency is named after the nineteenth-century economist and sociologist Vilfredo Pareto (1848–1923) who was one of the first to examine the implications of this idea.

"gain from trade" if these two agents swap apartments and arrange a side payment from B to A of some amount of money between \$200 and \$300. The exact amount of the transaction isn't important. What is important is that the people who are willing to pay the most for the apartments get them—otherwise, there would be an incentive for someone who attached a low value to an inner-ring apartment to make a trade with someone who placed a high value on an inner-ring apartment.

Suppose that we think of all voluntary trades as being carried out so that all gains from trade are exhausted. The resulting allocation must be Pareto efficient. If not, there would be some trade that would make two people better off without hurting anyone else—but this would contradict the assumption that all voluntary trades had been carried out. An allocation in which all voluntary trades have been carried out is a Pareto efficient allocation.

1.10 Comparing Ways to Allocate Apartments

The trading process we've described above is so general that you wouldn't think that anything much could be said about its outcome. But there is one very interesting point that can be made. Let us ask who will end up with apartments in an allocation where all of the gains from trade have been exhausted.

To see the answer, just note that anyone who has an apartment in the inner ring must have a higher reservation price than anyone who has an apartment in the outer ring—otherwise, they could make a trade and make both people better off. Thus if there are S apartments to be rented, then the S people with the highest reservation prices end up getting apartments in the inner ring. This allocation is Pareto efficient—anything else is not, since any other assignment of apartments to people would allow for some trade that would make at least two of the people better off without hurting anyone else.

Let us try to apply this criterion of Pareto efficiency to the outcomes of the various resource allocation devices mentioned above. Let's start with the market mechanism. It is easy to see that the market mechanism assigns the people with the S highest reservation prices to the inner ring—namely, those people who are willing to pay more than the equilibrium price, p^*, for their apartments. Thus there are no further gains from trade to be had once the apartments have been rented in a competitive market. The outcome of the competitive market is Pareto efficient.

What about the discriminating monopolist? Is that arrangement Pareto efficient? To answer this question, simply observe that the discriminating monopolist assigns apartments to exactly the same people who receive apartments in the competitive market. Under each system everyone who is willing to pay more than p^* for an apartment gets an apartment. Thus the discriminating monopolist generates a Pareto efficient outcome as well.

Although both the competitive market and the discriminating monopolist generate Pareto efficient outcomes in the sense that there will be no further trades desired, they can result in quite different distributions of income. Certainly the consumers are much worse off under the discriminating monopolist than under the competitive market, and the landlord(s) are much better off. In general, Pareto efficiency doesn't have much to say about distribution of the gains from trade. It is only concerned with the *efficiency* of the trade: whether all of the possible trades have been made.

What about the ordinary monopolist who is constrained to charge just one price? It turns out that this situation is not Pareto efficient. All we have to do to verify this is to note that, since all the apartments will not in general be rented by the monopolist, he can increase his profits by renting an apartment to someone who doesn't have one at *any* positive price. There is some price at which both the monopolist and the renter must be better off. As long as the monopolist doesn't change the price that anybody else pays, the other renters are just as well off as they were before. Thus we have found a **Pareto improvement**—a way to make two parties better off without making anyone else worse off.

The final case is that of rent control. This also turns out not to be Pareto efficient. The argument here rests on the fact that an arbitrary assignment of renters to apartments will generally involve someone living in the inner ring (say Mr. In) who is willing to pay less for an apartment than someone living in the outer ring (say Ms. Out). Suppose that Mr. In's reservation price is $300 and Ms. Out's reservation price is $500.

We need to find a Pareto improvement—a way to make Mr. In and Ms. Out better off without hurting anyone else. But there is an easy way to do this: just let Mr. In sublet his apartment to Ms. Out. It is worth $500 to Ms. Out to live close to the university, but it is only worth $300 to Mr. In. If Ms. Out pays Mr. In $400, say, and trades apartments, they will both be better off: Ms. Out will get an apartment that she values at more than $400, and Mr. In will get $400 that he values more than an inner-ring apartment.

This example shows that the rent-controlled market will generally not result in a Pareto efficient allocation, since there will still be some trades that could be carried out after the market has operated. As long as some people get inner-ring apartments who value them less highly than people who don't get them, there will be gains to be had from trade.

1.11 Equilibrium in the Long Run

We have analyzed the equilibrium pricing of apartments in the **short run**—when there is a fixed supply of apartments. But in the **long run** the supply of apartments can change. Just as the demand curve measures the number of apartments that will be demanded at different prices, the supply curve measures the number of apartments that will be supplied at different prices.

The final determination of the market price for apartments will depend on the interaction of supply and demand.

And what is it that determines the supply behavior? In general, the number of new apartments that will be supplied by the private market will depend on how profitable it is to provide apartments, which depends, in part, on the price that landlords can charge for apartments. In order to analyze the behavior of the apartment market in the long run, we have to examine the behavior of suppliers as well as demanders, a task we will eventually undertake.

When supply is variable, we can ask questions not only about who gets the apartments, but about how many will be provided by various types of market institutions. Will a monopolist supply more or fewer apartments than a competitive market? Will rent control increase or decrease the equilibrium number of apartments? Which institutions will provide a Pareto efficient number of apartments? In order to answer these and similar questions we must develop more systematic and powerful tools for economic analysis.

Summary

1. Economics proceeds by making models of social phenomena, which are simplified representations of reality.

2. In this task, economists are guided by the optimization principle, which states that people typically try to choose what's best for them, and by the equilibrium principle, which says that prices will adjust until demand and supply are equal.

3. The demand curve measures how much people wish to demand at each price, and the supply curve measures how much people wish to supply at each price. An equilibrium price is one where the amount demanded equals the amount supplied.

4. The study of how the equilibrium price and quantity change when the underlying conditions change is known as comparative statics.

5. An economic situation is Pareto efficient if there is no way to make some group of people better off without making some other group of people worse off. The concept of Pareto efficiency can be used to evaluate different ways of allocating resources.

REVIEW QUESTIONS

1. Suppose that there were 25 people who had a reservation price of \$500, and the 26th person had a reservation price of \$200. What would the demand curve look like?

2. In the above example, what would the equilibrium price be if there were 24 apartments to rent? What if there were 26 apartments to rent? What if there were 25 apartments to rent?

3. If people have different reservation prices, why does the market demand curve slope down?

4. In the text we assumed that the condominium purchasers came from the inner-ring people—people who were already renting apartments. What would happen to the price of inner-ring apartments if all of the condominium purchasers were outer-ring people—the people who were not currently renting apartments in the inner ring?

5. Suppose now that the condominium purchasers were all inner-ring people, but that each condominium was constructed from two apartments. What would happen to the price of apartments?

6. What do you suppose the effect of a tax would be on the number of apartments that would be built in the long run?

7. Suppose the demand curve is $D(p) = 100 - 2p$. What price would the monopolist set if he had 60 apartments? How many would he rent? What price would he set if he had 40 apartments? How many would he rent?

8. If our model of rent control allowed for unrestricted subletting, who would end up getting apartments in the inner circle? Would the outcome be Pareto efficient?