
5.1 Introduction: Crisis in a Cup

Mohammed Ali Idris is a coffee grower from Ethiopia. Interviewed in 2002, he had this to say about what was happening to his livelihood and life as a result of changes in the coffee market:

Five to seven years ago, I was producing seven sacks of red cherry (unprocessed coffee) and this was enough to buy clothes, medicines, services and to solve so many problems. But now even if I sell four times as much, it is impossible to cover all my expenses. I had to sell my oxen to repay the loan I previously took out to buy fertilizers and improved seed for my corn, or face prison.

Medical expenses are very high as this is a malaria-affected area. At least one member of my household has to go to hospital each year for treatment. It costs US \$6 per treatment. We also need to buy teff, salt, sugar, soap kerosene for lighting. We have to pay for schooling. Earlier we could cover expenses, now we can't. . . Three of the children can't go to school because I can't afford the uniform. We have stopped buying teff and edible oil. We are eating mainly corn. The children's skin is getting gray and they are showing signs of malnutrition.¹

In the early years of the century, plunging world coffee prices created an economic and humanitarian crisis in much of the developing world. Coffee is the world's second largest commodity trade behind oil; more than 25 million people are employed growing it, mostly on small farms. Coffee grows only in tropical and near-tropical regions, so the major producers are in Africa, Latin America and Southeast Asia. These countries are not rich to begin with. As the price small coffee growers can get in the market plummeted, malnutrition, disease and illiteracy were the results.

Figure 5.1 shows the long-term trend in prices earned by growers for Arabica, the most commonly traded type of coffee, during the period leading up to and following the crisis. (The other is Robusta, a lower quality bean that accounts for

¹Gresser, Charis, & Sophia Tickell. (2002). *Mugged: Poverty in your coffee cup* (p. 10). Oxfam. www.maketradefair.com.

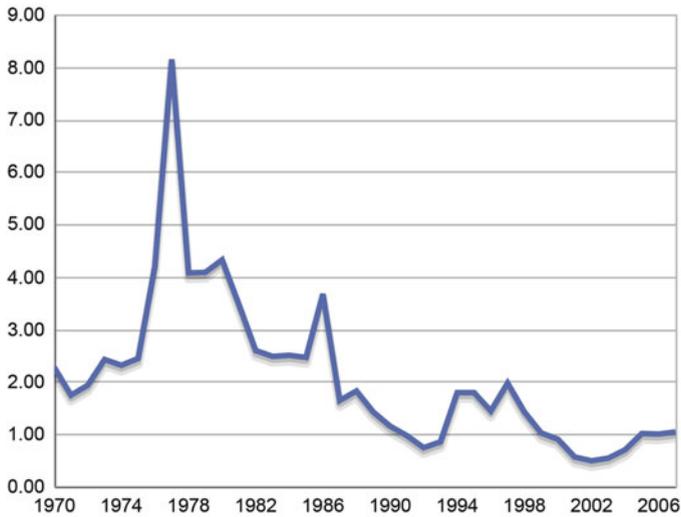


Fig. 5.1 NY price for Brazilian Arabica, 1970–2007, per pound in 2005 \$US (Source: USDA 2008)

20 % of the market.) Specifically, it displays the price per pound received by shippers in New York City for Brazilian Arabica, adjusted for inflation. (The prices in each year were converted to an equivalent number of 2005 dollars.) Other coffees were somewhat higher or lower, but the general trend was virtually the same. The high point was a spike in the mid 1970s, when a drought in Brazil led to a global shortage. Since then the overall movement was down, and by 2002 the price had reached an all-time low—so low, in fact, that farmers could no longer even recover their cost of production, much less earn an income. After this the price recovered, but only slightly.

This disaster affected whole countries. Particularly in Africa, it was common for coffee to account for a large share of a country's export earnings. In Burundi, for instance, an astonishing 80 % of all foreign sales were for coffee; the number was over 50 % in Ethiopia and close to that level in Uganda. When coffee earnings collapsed, so did the ability of these countries to earn the precious foreign exchange (dollars and euros) with which they could keep up payments on their foreign debt and perhaps purchase a few imports. So the coffee crisis gripped entire populations and pushed countries further back in their pursuit of economic and social development.

Those of us who are fortunate enough to have witnessed this catastrophe from a distance should also be disturbed. We are the ones who drink the brew purchased at the cost of so much hardship. Moreover, we have the resources to make a difference in how the world coffee market operates—if we understand it. But what exactly is the problem? Why did prices drop to such a low level? Unless we can answer these questions, anything we might try runs the risk of doing more harm than good.

There are many *possible* reasons why coffee prices might fall. Perhaps it has become cheaper to produce the bean in some regions, and this puts pressure on everyone else. Or maybe consumers are buying less of it. Or maybe there is too much coffee being produced, or not the right kind. Or maybe the price is being manipulated by powerful special interests. How would we go about trying to figure out what the true story is? What information should we look for, and how can we analyze to separate reality from fantasy?

Economics offers a helpful tool, supply and demand analysis. This is a fairly simple but highly flexible way of depicting how a market works, and it guides users toward answers to questions like those we have asked about coffee. It doesn't take long to learn, and it has many "add-ons" that permit it to tackle more complex issues, such as those involving labor, natural resources and technology. This chapter is devoted to presenting this tool and explaining how to use it. After a brief look at the assumptions necessary to apply the supply and demand model, the following section introduces the model itself. After this, we will look more closely at the three main building blocks, supply, demand and equilibrium. We then put this model through some practice exercises and conclude by seeing what it has to say about the causes of the coffee crisis.

5.2 Some Simplifying Assumptions

Models require that we impose simplifying assumptions on complex, messy real-world situations. What follows are the most important ones for using supply and demand analysis. Some just make it easier to use, while others are necessary if the model is to be used at all. The line that separates convenience from necessity is not fixed; it shifts as economic theorists delve ever deeper into the logical underpinnings of the relationships that make up the model. There is a lot of prestige attached to being able to demonstrate that a model holds under slightly less restrictive assumptions than previously believed. The assumptions of interest to us here are:

1. There is a single homogeneous type of good or service that can be the subject of a market. You can have a market for running shoes, garden hoes, or haircuts, but first you have to assume that there is a generic item we can identify as a running shoe, a garden hoe or a haircut. More sophisticated models of markets can accommodate qualitative differences, but only in certain specific ways. Most supply and demand models you will see, in this text and elsewhere, simply assume that all goods in the market are identical. From a historical standpoint, this is an extraordinary assumption—one that could only have been made in relatively recent times. For most of human history, production methods were artisanal: each item was individually distinctive, and it wouldn't have occurred to people to treat them as identical. (There are some exceptions; for instance, evidence exists that grain, wine and other commodities were viewed as standardized goods in Roman times.) With the rise of standardization and especially mass production, thinking that similar goods are essentially the

same comes more naturally. We will have more to say about standardization in Chap. 7.

2. Producers produce goods with the intention of selling them, and those who want to use these goods expect to buy them. This should be obvious, but again it is a product of social evolution and has hardly been the case through most of human history. People have to accept and adopt the role of buyer or seller. One of the difficulties with establishing markets in human organs, for instance, is that there is great resistance to adopting these roles.
3. Market participants are self-interested. Sellers want to receive the highest price they can get, and buyers want the lowest price. Neither sacrifices her own interests in the market for the sake of others.
4. Market participants are rational. Here rationality is used in the sense of Chap. 3: people maximize their self-interest through the choices they make. In particular, sellers refuse to sell unless the price they receive is sufficient to cover their costs, and buyers refuse to buy unless the price they pay is less than or equal to the value they get from their purchases. No one is made worse off by entering into a market transaction; no one makes mistakes.
5. A corollary to the previous assumption is the “law of one price”: in every market there can be only a single price at any moment in time. No one would agree to buy at a higher price if a lower price were available, and no one would agree to sell at a lower price if a higher price were available. Notice that this “law” is often broken in practice.
6. Other than their market interconnection, individual preferences are independent of one another. What I want to buy and the price I am willing to pay do not depend on what you buy and vice versa, except insofar as your choice affects the market price. The same goes for sellers. This assumption allows us to simply add up the independent demand curves of individuals to get the overall market demand curve, and the independent supply curves of different producers to get the market supply curve. It also guarantees that the supply and demand curves will be independent of one another, so that we can track changes on one side of the market (supply or demand) while holding the other constant.
7. There is a one-time-only interaction between buyers and sellers. Historical time as such does not exist in a supply and demand analysis; rather, we assume that the model represents a single, unrepeatable instance. The choices of buyers and sellers may well depend on their expectations of the future, but their participation in the market is instantaneous. In making their choices, they do not take into account any effects their current decisions may have on future market transactions. Each moment in the market is its own universe, with no past or future. (In the language of game theory, a market is a one-shot game, not a repeated game.)

These may seem like very stringent and unrealistic assumptions. As in most aspects of economics, however, it pays to give the matter some benefit of the doubt. The simplifications necessary to employ market analysis are often not too far removed from reality, and experience shows that we can often learn a lot by organizing our thinking in terms of this model. As you will see, for instance, it

definitely helps us understand the ups and downs of coffee prices. Moreover, the model provides a baseline, a reasonable starting point for analysis. When we encounter an aspect of real life that forces us to question one of these assumptions, we can then think about how this will alter the conclusions a simple supply and demand analysis would otherwise impose on us. You will see examples of this in later chapters. For now, we will make the assumptions we need to make in order for supply and demand to work as intended—but we will do it *knowing that we are making these assumptions*.

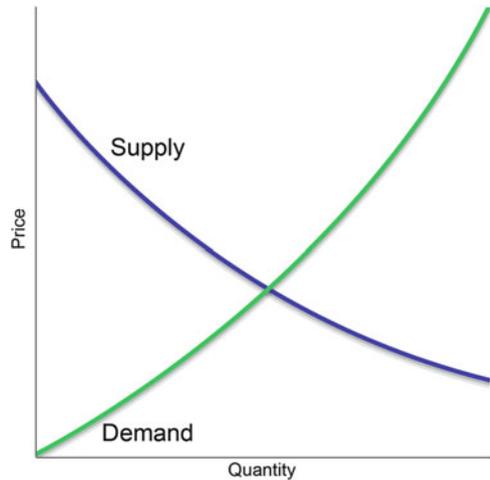
5.3 A First Look

Here is a “naked” supply and demand diagram. It is stripped to its bare essentials. What is being depicted is a market, which, in non-technical terms, consists of all the buying and selling of some particular good or service. Individuals do not appear in this diagram; rather, it represents everyone involved in this market collectively. At this extremely general level, only two kinds of information are recorded, the possible prices people may buy or sell at, and the possible amounts they may buy or sell. All other aspects of the situation—the thoughts or feelings of these people, their relationships to one another, the qualities of the goods being considered—are ignored. Note, incidentally, that there are few labels and no numbers in this diagram. We don’t know what goods are being traded, how much or at what price. All we see are two curved lines intersecting in two-dimensional space, bounded on the left and below by measurement axes.

So this is it: the basis for nearly all ground-level economic reasoning. In a nutshell, you would read it in this way: The straight lines represent two axes; the vertical axis is price and the horizontal axis is quantity. Thus, the further “north” we go in this diagram, the higher the prices, and the further “east” we go the higher the quantities of goods being offered for sale or purchase. The supply curve represents the amount individuals wish to sell at various prices. (One assumption behind this curve, and the demand curve as well, is that there will be only one price for all the transactions taking place in this market at any moment in time. The price may change, but only over time, not between individual goods or people.) It is upward-sloping, which is to say that it travels from the southwest to the northeast. At lower prices there are fewer goods offered for sale, and as the price rises more goods are placed on the market. The demand curve represents the amount individuals wish to buy at different prices. It is downward-sloping, which means that at high prices there is less desire to buy, but more at lower prices. Here in a nutshell is the framework for thinking about markets (including coffee markets). Don’t worry if Fig. 5.2 is still mysterious at this point: we are just making an initial acquaintance.

Before going further, consider an important point: even though we are drawing curves in a two-dimensional diagram, the relationships they represent are nearly always invisible in real life. That is, in general you don’t see supply and demand curves. At any moment in time you see a specific price, or maybe a small range of prices, and a quantity of goods being bought and sold, but you don’t see the full

Fig. 5.2 A basic supply and demand diagram



range of *possible* prices and quantities represented by these two curves. True, you might see more price-quantity combinations over a period of time or across different regions, but, strictly speaking, these would not tell you about supply or demand curves at a moment in time and a specific place. (You would also have to use somewhat imprecise statistical techniques to infer separate supply and demand curves from observable trading data.) The closest you can come to actually seeing such curves in real life would be to perform a survey of buyers and sellers, asking them how much they would buy or sell at various prices. Market researchers sometimes do this, but in most cases to which supply and demand analysis is applied, there do not exist “real” curves corresponding to the ones we draw in our diagrams. What this means is that such curves are not real objects we can observe, but intellectual constructs that help us understand how markets work.

Having clarified that point, let’s look under the hood. Where do these curves come from, and what shapes can they reasonably take?

(A) The supply curve. The amount of something that people are willing to sell depends on many things. The cost of acquiring this good is an obvious consideration. Perhaps the sellers are making this item themselves, in which case their supply will depend on the cost of materials, the cost and characteristics of labor, the technology, etc. Or they may be middlemen, like wholesalers or retailers, in which case they have to pay attention to the price they have to pay as buyers in some other market. The capacity of the sellers may also be a factor: how many of them are there, and how much do they have in the way of investments, like buildings, land, materials, equipment, etc.? Another influence might be the expectations that sellers have of the future—whether they expect prices to go up or down, for instance, which will persuade them to sell now or wait until later. And of course sellers will be interested in the price they can get if they sell today. For any given situation, the possibilities are almost endless. But the trick employed in economics is to suppose that all of these factors are assumed to be constant, fixed and unchanging, except

just one: the current price buyers are offering to pay. This is the meaning of the supply curve. It says that, with all other factors held constant, there is a one-to-one relationship between the going price and the amount sellers wish to sell. (Remember that in a supply and demand diagram there are no individuals, just sellers as a group and buyers as a group.) Each point on the supply curve represents one such combination, and movement along the curve means seeing how a change in the price affects the amount offered for sale. If you go up (northeast) along the curve, the price is rising along with the quantity; the opposite holds for movement down (southwest) along the curve. The story is Pavlovian: flash a price and the sellers will respond with a perfectly predictable amount they want to sell.

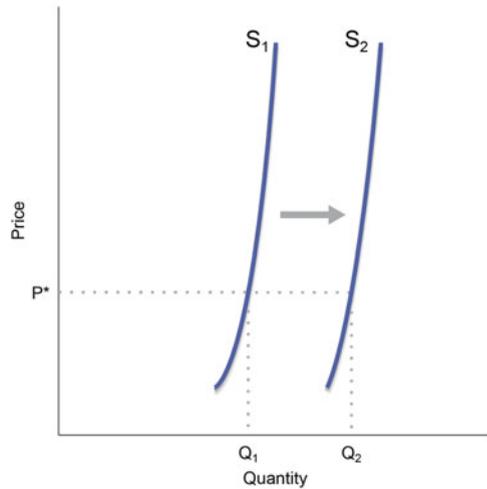
There is one critical point to bear in mind: this supposedly predictable relationship between price and quantity depends entirely on the initial assumption that every other factor is being held constant. This assumption is important enough to merit its own name, *ceteris paribus*. The words are Latin, for “things [being] equal”. In econospeak, people will say things like, “This relationship holds *ceteris paribus*.” They are simply invoking the common, but sweeping, assumption that nothing else will change that might interfere with the one relationship being examined. In the case of the supply curve, each point along the curve, and the curve as a whole, depends on everything else in the whole world being exactly what it is and staying that way. If anything of significance changes, the whole curve moves. Read that sentence again: if anything of significance changes, the whole curve moves. Understanding why and how this happens is fundamental to understanding the use (and possible misuse) of the supply and demand apparatus.

Let’s consider our original example, the global market for coffee. The supply side of this market consists of producers and processors who make coffee beans available to the ultimate consumers, such as most of us. One well-publicized event that occurred was that, during the decade of the 1990s, Vietnam became a major producer. Let us suppose (falsely of course) that this is the only thing that happened on the supply side of the market during the 1990s. We might represent this in the diagram on the following page:

Before Vietnam entered the market the curve might be drawn as S_1 , after as S_2 . The whole curve has shifted to the right. For any given price, say P^* , the world’s suppliers, including Vietnam, will produce more: Q_2 instead of Q_1 . This is true for any possible price: the new supply curve is completely to the right of the old one. This shift in the supply curve underlines the original point: it was the assumption of *ceteris paribus* that enabled you to draw the curve in the first place. If this assumption is broken—if some factor of significance other than the market price changes—the whole curve must be redrawn. In this case, the change is rather obvious; the addition of Vietnam to the ranks of major producers leads to a rightward shift in the curve as a whole.

To summarize, we have seen two potential ways prices or quantities can change. Either can change as a result of a change in the other, as we saw represented by movement *along* a supply curve. But either can also change despite *no* change in the

Fig. 5.3 Vietnam expands its coffee production; the supply curve shifts



other, as represented in movement *of* a supply curve. Knowing the difference between these two possibilities is 90 % of what you need to know about this topic.

Take a moment to consider some other hypothetical possibilities. (1) Because of a drought during a crucial phase of the growing season, there is a change in the amount supplied. This will look like a mirror-image of Fig. 5.3. Now the supply curve will shift to the left. Once again there has been a change in one of the factors normally held constant under the *ceteris paribus* assumption, in this case the productivity of coffee growers. At any potential price they will bring less to the market. (2) Caffeine is found to be a major cause of brain damage. This will *not* cause any shift in the supply curve, because none of the *ceteris paribus* factors have changed. From the supplier's perspective, the only thing that's changed is the amount consumers are willing to buy, which is to say the amount they will be able to charge if they want to sell their harvest. The supply curve stays put, but there is movement along it—in this case down and to the left. (3) Fair Trade importers offer growers a higher price than the going market rate. Once again the supply remains fixed, since it is the price sellers can get which is changing. The movement is up and to the right along the supply curve, at least for the fortunate growers who are able to qualify under Fair Trade rules.

If this makes sense to you, you are ready for the other 10 %. Take another look at Fig. 5.3. Notice how vertical the supply curve looks in this diagram, compared to the one in Fig. 5.2. Why did I draw it this way? The assumption behind the artwork (I'm being charitable) is that agricultural commodities like coffee have more vertical supply curves than average. The reason for this is that, in the relatively short run (within the growing and harvesting period of a year), the amount that will be produced is more or less what it is, no matter what the going price. Coffee plants don't yield more beans just because the price goes up. The decision to plant, combined with the weather and a few other factors, predetermines the harvest. On

the other hand, the curve is not completely vertical, because there is still some discretion on the part of growers. They can harvest more or less intensively, put more or fewer resources into storage and processing in order to cut down on waste, etc. When prices are high, they will squeeze a few more beans out of their operation somehow; the reverse when prices are low. So the supply curve still slants, but only a little.

We have a name for highly vertical supply curves: they are called **inelastic**. There is a formula for calculating the elasticity of a supply curve:

$$\text{elasticity of supply} = \frac{\% \text{change in quantity supplied}}{\% \text{change in price}}$$

We speak of a supply curve as being elastic if the formula has a value greater than one—if the percentage change in quantity is greater than the percentage change in price. It is inelastic if it is less than one. The nearly vertical curves in Fig. 5.3 are highly inelastic. Imagine going from a low price to a high one on either curve. There would be a very large percentage change in the price, perhaps more than 100 % (double), but only a very small percentage change in the quantity, say 5 % or 10 %. So the value of the fraction, with quantity on top and price on the bottom, would be close to zero. A very horizontal curve would have an extremely high elasticity; it could easily be 10 or 20 or 100 (if it were very flat). Here is a tip for making sense of the word “elasticity”: think of something elastic like a rubber band. It is highly stretchable; you pull on it and it gets longer. This is the same as the quantity supplied in an elastic supply curve. If the price goes up, the quantity stretches in response. The opposite is true for an inelastic supply curve. Price changes can pull and pull on it, but quantity stretches hardly at all. This may not be the most scientific approach to understanding elasticity, but you might find it useful anyway.

Visually, a supply curve is inelastic if it points downward toward the horizontal axis, which is the case in Fig. 5.3. If it points toward the vertical axis, it is elastic. Of course, unless the supply curve is a straight line, it will point in different directions depending on what part of the curve you look at. Recall from geometry that the slope of a curve at any point is determined by the straight line tangent to it at that point. We could then say that the supply curve is elastic at one point but inelastic at another, depending on the slopes of the corresponding tangents.

Why bother with this terminology of elasticity? One reason is that it gives us a language to talk about the geometry of curves without actually having to draw them. Simply by saying that the supply of coffee is inelastic, I am alluding to a curve such as the one(s) drawn in Fig. 5.3. Moreover, if the supply of coffee really is inelastic, this tells us something about the nature of the coffee industry. In the time frame represented by the diagram, the amount of coffee producers put on the market is not very responsive to the price: the amounts produced at a high price will be very similar to the amounts produced at a low price. This piece of information will prove useful later on when we attempt to unravel the mystery of the coffee crisis.

(B) The same sort of analysis can be applied to the demand side of the market. Many factors affect the amount that consumers want to buy—their desire for the product, the amount of income available to them, the prices of other goods they might buy instead, their expectations about future prices and availability, and certainly the price currently being charged for the good in question. In order to produce a demand curve, we make the assumption that all these factors, except the current price, are held constant, and then we can consider the relationship between price and quantity purchased. This relationship depends on the *ceteris paribus* assumption, just as the supply relationship did; change one of the factors being held constant and the whole relationship changes.

Let's imagine how this would work for the world coffee market. Suppose the buyers we are interested in are the volume coffee purchasers, the “middlemen” who buy from the actual growers and then resell to the companies that sell coffee to people like you and me. These buyers must keep in mind the amount of coffee consumers may be willing to drink, costs of marketing and distribution, and the need to stockpile supplies in times of low prices or draw down stockpiles if prices are expected to rise. But we can also assume that they will buy in larger quantities when current prices are low, and in smaller quantities when they are high. This gives rise to a demand curve such as we see in Fig. 5.4 on the next page.

As you can see from this diagram, at high prices buyers purchase somewhat less; as the price falls, the amount they purchase goes up. Nevertheless, the quantity varies just a little, even if price changes are substantial. That suggests that demand is **inelastic**, just as supply was. The formula for the price elasticity of demand is virtually the same as that for supply:

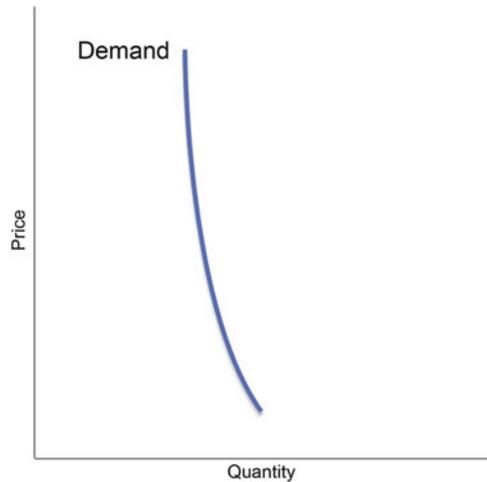
$$\text{price elasticity of demand} = \frac{\% \text{ change in quantity demanded}}{\% \text{ change in price}}$$

In this case, very large percentage changes in price have very little effect on the percentage change in the quantity demanded; hence the choice of an inelastic demand curve. Is this realistic? Both logic and experience suggest it is. The price of coffee beans is a relatively small part of the total cost of a cup of coffee, so coffee drinkers will tend to buy about the same amount whatever the price of beans. The number of coffee drinkers is generally independent of prices anyway, and wholesalers have little difficulty passing along increased costs.

Note that no change is assumed to take place on the demand side of the market. No events have transpired in the world to change the relationship between the price charged for coffee and the quantity purchased. The demand curve doesn't tell us what the price or the quantity demanded will be, just that if we know one we can deduce the other by finding the corresponding point on the curve.

(C) Equilibrium. So here we have two pieces of information, the supply curve and the demand curve. What do they tell us about events in the coffee market, or any other market? To answer this question, we need to add more assumptions about how buyers and sellers respond to one another.

Fig. 5.4 World demand for coffee

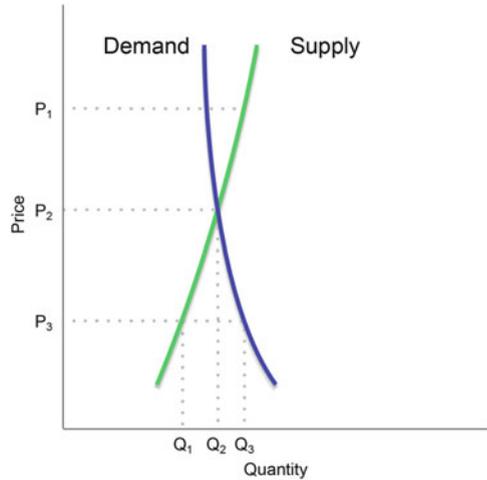


Let us assume that buyers always prefer to be at some point on their demand curve rather than off the curve, and that sellers always prefer to be at some point on their supply curve. If this is true, then both can do this simultaneously if and only if their curves cross. The reason is that there can be only one set of prices and quantities for both groups: whatever price buyers are paying is also the price sellers are receiving, and the amount buyers are buying is also the amount sellers are selling. We can see this graphically in Fig. 5.5 on the next page.

At any point along the D curve buyers would be acting in accordance with their intentions. At any point along the S curve sellers would be acting in accordance with their intentions. They are both able to do this simultaneously at price P_2 and quantity Q_2 . At any other price this would not be possible. Consider price P_3 , for instance. At this price, buyers want to purchase an amount of coffee equal to Q_3 , while sellers would prefer to sell Q_1 . Both cannot be satisfied. The most likely outcome is that sellers will make Q_1 available, and that is all buyers will be able to acquire. The difference between what they want to buy and what they are able to buy, Q_3 minus Q_1 , represents **excess demand**. Those who place their orders first may be able to make their purchases, but there will be other buyers who will be told that all the supplies are gone. An opposite situation would occur at P_1 . In this case the amount that buyers wish to buy is less than the amount sellers wish to sell. Some suppliers would manage to make sales at this higher price, but others would be left with unsold stock. This would be a condition of **excess supply**.

Economists generally assume that both excess demand and excess supply are unstable. If either buyers or sellers are not able to make the transactions they wish (as represented by their demand and supply curves), they will have an incentive to change their response to the market. In the case of excess demand, suppliers who find themselves quickly selling out will be tempted to raise prices, and they may find buyers among shoppers who are trying to avoid being frozen out by a shortage. In the case of excess supply, suppliers who find their inventories piling up may try

Fig. 5.5 Equilibrium in the coffee market



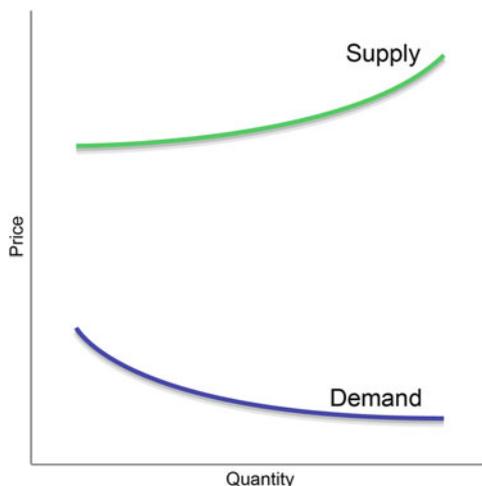
lowering their prices, and clever consumers will bargain aggressively to take advantage of the situation. Thus excess demand will tend to lead to an increase in prices, and excess supply to a decrease. As long as the price is below P_2 there will be pressure for the price to go up, and vice versa. Only at P_2 is the price at least temporarily stable.

This analysis explains why the term **equilibrium** is used to describe prices and quantities like P_2 and Q_2 . A situation is in equilibrium if there are no forces internal to it that would lead to a change; change can come only from the outside. If the situation is in **disequilibrium**, however, change is likely to occur even if no outside factors play a role. In our analysis of demand and supply, the line that separates inside and outside is the *ceteris paribus* assumption. If this assumption holds—if all the factors that enter into the demand and supply relationships remain constant—then nothing is transpiring on the outside to alter the market. It will remain as it is if it is in equilibrium, which is to say if the market price is P_2 . If this is not the price there will be pressure for readjustment, pushing the price toward P_2 . The pressure will end only when equilibrium is restored.

It should be apparent that equilibrium is a powerful concept for explaining and predicting how markets will function—the purpose of positive theorizing in economics. Once we have the information that enables us to draw supply and demand curves, we can say with some confidence what prices we expect to see, and how much will be produced and sold. This is true even if the initial situation is not in equilibrium, since the pressures exerted by excess demand and supply will tend to move the market in the direction of equilibrium.

To see the power of this concept, consider a good that is *not* currently bought or sold on the market, such as individual solar-powered helicopters. The technology probably exists to produce this gizmo, and no doubt there are some people who would like to have one. Nevertheless, we don't see them for sale at the mall. Why not?

Fig. 5.6 Demand and supply for individual solar-powered helicopters



The market in all likelihood looks like that depicted in Fig. 5.6.

There are prices at which consumers would buy and prices at which producers would produce and sell, but there is no overlap between them: the cheapest price at which such a device can be made is still too expensive for the most rabid consumer. There is no equilibrium price in this market, and the only equilibrium quantity is zero. If the supply-and-demand model of the market, with its predictive notion of equilibrium, is correct, then *every* non-produced good must look something like this. Thus, equilibrium analysis can not only explain the prices we expect to find in the market, but also why there are no markets at all for many potential goods.

At this point it is only fair to mention three important limitations of equilibrium reasoning in economics. First, in the real world few markets are actually in equilibrium at any moment in time. If there are strong pressures pushing markets toward equilibrium, it is also true that the flow of outside events never stops, and this leads to changes in the equilibrium even before the market can get there. At best, equilibrium analysis is approximate; it is not a precise reflection of how markets function. Second, the assumptions we made about excess demand and supply do not always hold. Buyers and sellers may respond as we suggested, but they may not. There are many markets in which excess supply or demand can persist for months or years, with no apparent effect on prices. In those cases the notion of equilibrium as an intersection of demand and supply curves (no excess demand or supply) may be too simplistic. There are more intricate analytical devices (such as those offered by game theory) that can be used to represent more complex forms of equilibrium. (See Box 5.1.) Finally, equilibrium is a positive concept—it helps us explain or predict—but it has no necessary normative significance. Although we already considered this in Chap. 3, it is such an important point that it deserves another paragraph all to itself.

Language can play tricks on us. Words often mean different things in different contexts, and we can go wrong by failing to recognize the distinctions. Equilibrium

can give rise to exactly this type of confusion. In normal speech, equilibrium is a desirable state of affairs. The word has connotations of balance and harmony. It is only reasonable to extend this warm, fuzzy feeling to economic equilibrium—reasonable but wrong. In fact, there is no presumption in economics that a market equilibrium is better than a disequilibrium. It all depends on the market. Here is an extreme example: the market in nuclear weapons. It is altogether possible that there are private individuals who have access to nuclear weapons at the present time. Some may have escaped military control during the collapse of the Soviet Union; it may also be the case that countries with secret programs have permitted a few of the weapons to enter the black market. No doubt there are other individuals, many of them terrorists or criminals, who would like to purchase such weapons. Insofar as sellers are motivated by the money they can make selling weapons and buyers are constrained by how much they can afford to pay, there are supply and demand curves, and therefore also a market equilibrium. At some price, the number of nuclear warheads underground arms dealers want to sell is equal to the number potential mass murderers want to buy. With luck, however, the market will *not* reach an equilibrium, and these transactions will never take place.

This is clearly an extreme example, and yet it is not so different in principle from more commonplace economic threats. Many of the goods produced in modern economies are harmful to the environment and to human health. Much of the work performed to manufacture and distribute these goods is hazardous, degrading or oppressive. We can explain why these problems exist using the apparatus of market equilibrium, but we should be clear at all times that explanation is not justification. Equilibrium is a positive, not normative, concept. In the next chapter we shall consider the conditions under which there may (repeat: may) be a connection between equilibrium and human betterment, but for now we should view them as entirely separate and distinct phenomena.

Box 5.1: Excess Demand, Excess Supply and Everyday Life in Market Economies

The standard representation of market equilibrium is where the demand and supply curves intersect: the amount sellers wish to sell is exactly equal to the amount buyers wish to buy at the equilibrium price. If there is either excess demand or excess supply, the price is out of equilibrium, and economists expect it to adjust.

If you think a bit about this, you will realize it does not capture the normal experience we have living in a market economy. The easiest way to see this is to consider what life is like in an economy *not* governed by markets. Much of the world was like this, in fact, until 1989, when the Soviet Union began to collapse and its allies in eastern Europe shed their Communist parties and embraced capitalism. In the old Soviet model, sellers were in a position of power and buyers had to struggle to be able to make a purchase. Stores

(continued)

Box 5.1 (continued)

seldom had enough inventory to meet demand, and there were frequently lines that formed so that scarce items could be distributed on a first-come, first-serve basis. There was no attempt to advertise goods, and sometimes you couldn't even locate where they could be found. The stores themselves often kept most of their wares in a back room: you went to the counter and asked the clerk to see if what you wanted was in stock. With shortages the norm, store employees could play favorites, deciding who to serve on the basis of friendship, gifts or sheer whim.

It is exactly the other way around in a market economy. In most cases there is chronic excess supply: sellers carry more inventory than they can expect to sell, and they knock themselves out trying to convince us to buy more of it. An extreme example is provided by grocery stores. A recent study found that about 10 % of all food items stocked by American groceries is never sold at all, but simply thrown away—in a world in which hundreds of millions of people do not get enough to eat. But this is just a more visible instance of a pattern that holds in most parts of the economy. Hardware stores stock more tools than they can sell, and office supply stores stock more paper. In most restaurants there are more tables than are likely to be filled on all but a few nights, and the remainder bins of bookstores remind us that more books are printed than readers are usually willing to buy.

The point is that much more effort goes into sellers courting consumers than consumers seeking sellers. The advertising industry is built on this fact. It is so fundamental to how we live in market economies that it is easy to lose sight of. Nevertheless, it is worth pondering: why is excess supply so typical of most markets? Is it because of too little competition or competition that is especially intense? Why is it profitable to offer a larger quantity of goods and services than you can expect to sell? One way to approach these questions is to think about markets where excess supply is *not* the norm, such as sales of new cars (where there are wait times for popular models) and medical services (where patients are much more likely to wait for doctors than doctors are to wait for patients). How and why are these examples different?

Whatever answers we give, the fact remains that the simple supply and demand model, with its image of equilibrium where excess supply is zero, is at best an approximate representation of the real world. Economic life would be very different if sellers needed buyers no more than the other way around.

Before leaving the topic of equilibrium, we should consider what the concept means in a world of many markets, all of them interlinked. There are markets for shoes and markets for socks, markets for cars and buses, markets for the goods farmers buy and the goods they sell. What happens in one of these affects the others. When the pressure to get to equilibrium leads to adjustment in one part of the economy, this is experienced as an “outside” force in other markets. A chain

reaction is set off, and the effects may ultimately rebound on the market that set the process in motion in the first place. We can make economic analysis much simpler by focusing on just one market at a time, but to do so is to overlook the interconnectedness of real-world economies.

To address this problem, economists have developed the concept of **general equilibrium**. This is a situation in which all markets are in equilibrium simultaneously, each taking the others into account. Economists spent decades investigating the mathematical basis for the hypothesis that economies might have such a general equilibrium. By the mid 1950s it had been proven that this was a logical possibility, although more recently theorists have found new grounds for skepticism. They have shown that economies typically have many potential general equilibria, a problem from the perspective of explanation and prediction. (If an analysis shows that, say, 12 combinations of prices and quantities could all be in equilibrium, how do we know which one is likely to occur?) They have also shown that the adjustment process sparked by excess demand and supply can shift the equilibrium itself, making it more difficult to use the supply-and-demand model to make predictions. We will have more to say about general equilibrium in the final chapter; for now, the main point is that, in using demand and supply curves, it is important to keep in the back of one's mind the possibility that there may be important effects that extend beyond the confines of a single diagram.

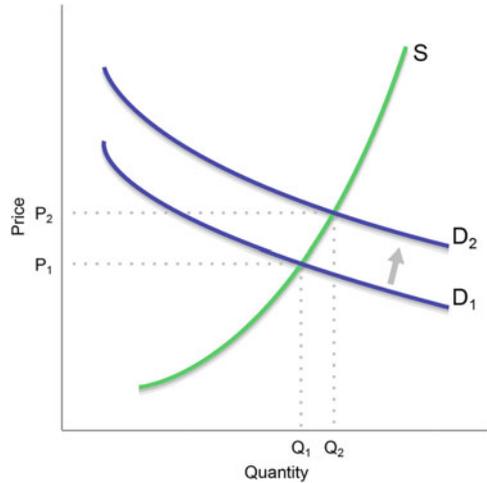
5.4 Using Supply and Demand

The best way to get a feel for the supply and demand apparatus is to use it. Let's imagine various situations that might arise in an economy and see how they could be approached with the three building blocks of supply, demand and equilibrium.

(A) Ice cream. Suppose you open a home-made ice cream stand specializing in vegetable flavors—spinach, zucchini, turnip, etc. After the initial burst of demand (because people have been waiting all their lives to try these new taste sensations), you settle into a predictable level of sales. We might ask how various events would affect the amount of ice cream you sell per week and the amount you are able to charge for it. For instance, suppose that global warming produces a month-long heat wave during the late spring. People are looking for cooling, refreshing snacks, like cucumber-cilantro swirl. The result may be depicted as in Fig. 5.7 on the following page.

By shifting the demand curve from D_1 to D_2 , we are reflecting the increased desire of consumers to buy ice cream at any price. Even after demand shifts, it is still the case that more ice cream will be bought as the price falls (the curve is downward-sloping), but the quantity is greater. If you pick one price in particular (identified as a vertical distance up the P axis), you can find out what the sales will be by tracing a horizontal line over to the demand curve. First you will reach D_1 , which tells you how much will be bought if this is the demand curve, and then you will reach D_2 . Since D_2 is to the right of D_1 , more ice cream will be bought during

Fig. 5.7 The effect of a heat wave on ice cream sales. Before the heat wave, the demand curve was D_1 , resulting in price P_1 and quantity Q_1 . After the heat wave, the demand curve shifts to D_2 , producing the new equilibrium price P_2 and quantity Q_2 .



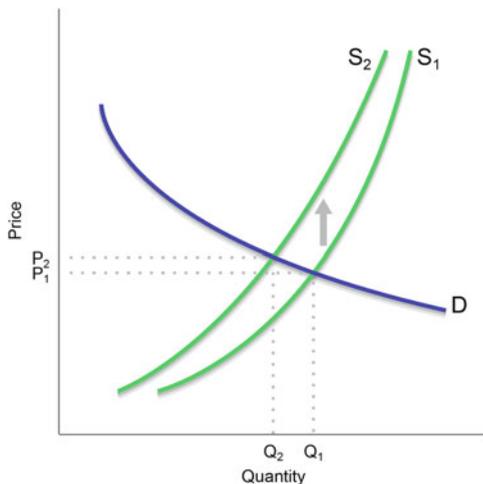
the heat wave than before it. Alternatively, you could think of D_2 as being *above* D_1 : for any potential quantity of sales, a higher price can be charged during the heat wave. Note the thought process that led us to shift the demand curve. Because of the higher temperatures, consumer preferences, which are one of the factors normally held constant in conjunction with the *ceteris paribus* assumption, have changed, and this shifts the entire curve. On the other hand, nothing has transpired to change the *ceteris paribus* factors on the supply side; this curve remains where it was.

Applying the concept of equilibrium permits us to make a prediction. We expect that, in the absence of any outside change (like the weather), the market would settle at a price of P_1 and a sales level of Q_1 . Due to the heat wave, these will change to P_2 and Q_2 . In other words, the price will go up *and* so will the sales. The process has been instigated by a change in the demand curve. Although there has also been a change in the amount of ice cream sold, there has been no change in the supply curve. *We see movement of the demand curve and movement along the supply curve.* Knowing which curve to move and which to keep in place is 90 % of the art of applying the supply-and-demand framework.

Now suppose that, rather than a change in the weather, we see a change in the tax laws. What happens if the city decides to raise money by placing a tax on ice cream vendors for each scoop they sell? We can picture the result in Fig. 5.8 as shown on the following page.

If ice cream sellers such as yourself have to pay a tax, this increases your costs and means that you now need to charge a higher price than before. In other words, your supply curve will shift upward (or to the left). On the other hand, nothing has changed to alter the relationship between the price and the quantity consumers want to buy, so the demand curve stays put. The equilibrium quantity sold will fall from Q_1 to Q_2 , while the equilibrium price will rise from P_1 to P_2 . Note that the increase in price per scoop is substantially less than the amount of the tax per scoop. (You can read the size of the tax increase from the vertical distance between S_2 and S_1 ,

Fig. 5.8 The effect of a tax increase on the ice cream market. The original supply curve is S_1 , resulting in equilibrium price P_1 and quantity Q_1 . After a tax is placed on ice cream sellers, the new supply curve is S_2 , with equilibrium price P_2 and quantity Q_2



which indicates how much the cost per scoop has risen for sellers.) This means that sellers are able to pass along some of the tax increase, but not all of it. By looking closely at this diagram, you can see that the portion that can be passed on to consumers depends on the slope of the demand curve, loosely related to the price elasticity of demand. As drawn, the demand curve is rather elastic: relatively small changes in price induce consumers to make relatively large changes in their ice-cream eating habits. What if the D curve were inelastic—more nearly straight up and down? If you try this out on a piece of scratch paper, you should find that the price increase becomes larger and the quantity decrease smaller. At the limit, with a perfectly inelastic (vertical) demand curve, all of the tax increase could be passed along to the consumer. At the other limit, with a perfectly elastic (horizontal) demand curve, none of the tax could be passed along, and all would be absorbed by the seller.

This gives us one clue toward why economists think about the elasticity of demand and supply. As demand becomes more inelastic, suppliers gain more power over buyers; they can increase prices with relatively little concern about lower sales. But what determines the elasticity of demand? Many factors are relevant, but the most important ones are tied to a single word: substitutes. If consumers have ample opportunity to substitute other goods for the one being considered, their demand will be elastic. This was the assumption that led to the fairly elastic demand curve in Figs. 5.7 and 5.8; presumably most consumers of ice cream have alternative ways to spend their money that are almost as satisfying, and so they will be sensitive to relatively small price changes. An opposite case would be cigarette smokers, whose tobacco consumption is relatively insensitive to price shifts. (This is not to say that prices have no effect at all, just less than for ice cream, because of the addictive quality of nicotine.)

(B) Housing in a college town. Here our example will be the market for apartment rentals in a college town. To simplify matters, we will overlook all the differences in size, quality and location that affect prices and assume that there is a

single “standard” apartment available at the same rent everywhere. (This is only for the purpose of keeping the analysis within a single diagram; in principle we could incorporate all these differences if we were willing to draw a separate diagram for each type and location.) What does the demand curve look like? It is probably moderately inelastic. There are substitutes for renting an apartment—doubling up in existing units, living with parents or other relatives, or finding some other town to live in and commuting—but these are not always convenient. (The reality of homelessness in our communities demonstrates that many of our neighbors can find neither an affordable apartment nor a satisfactory alternative.) As for the supply curve, it is almost completely inelastic in the short run. In other words, within the next few months the supply of housing is nearly fixed; it can only be augmented or diminished slightly by decisions involving repairs to marginal units or potential subdivision. Only in the long run, over a horizon of several years, is it possible to greatly increase the amount of housing available through new building or to remove a large number of units through demolition or conversion to new uses. So let us stick with the short run for now. As a first exercise, imagine how the market will change if the local college expands to take on more students. This is shown in Fig. 5.9 on the next page.

As more students move to town to attend the college, the housing market becomes saturated. If no measures are taken to anticipate this, the supply will increase only marginally, from Q_1 to Q_2 ; meanwhile, the equilibrium rent will shoot up from P_1 to P_2 . Only in the long run, as more apartments are built and the supply curve shifts to the right, will rents moderate.

We could well imagine that, if rents skyrocket as in Fig. 5.9, there will be pressure to hold them back politically. One way to do this is rent control. Rent control laws can take many forms, but they all have in common a legislative prohibition of rent increases above a certain percentage or in the absence of certain types of investment. If rent control is adopted in our hypothetical town, the immediate result might be a situation like that in Fig. 5.10 on the following page.

A new controlled rent, P_c , might be established, higher than the original level (P_1) but below the level that would otherwise result in the market after the influx of students (P_2). Fortunately, in the short run, there would be almost no visible effect on the number of apartment units available for rental. (There isn’t room in the diagram to depict a Q_c between Q_1 and Q_2 .) The danger, however, is that a poorly drafted rent control law could have damaging effects in the long run. This can be seen in Fig. 5.11 on p. 89.

Once again, rent control reduces the rent students and others have to pay, but now there is a bigger difference in the number of apartments available. The more the demand curve shifts to the right, the greater the gap, over time, between the apartment supply at market rents and at controlled rents. Meanwhile, with a shortage of apartments, opportunities open up for black market-type activities. Renters can sublet their apartments for a considerable profit, and landlords can charge “finder’s fees” and other dubious charges to take advantage of the scarcity. These and other stratagems are commonplace in cities with rent control laws. On the other hand, a well designed law can mitigate most of these effects if it

Fig. 5.9 The effect of a college expansion on the local apartment market. The initial demand curve is D_1 , producing the equilibrium price P_1 and quantity Q_1 . After the college accepts more students, the demand curve shifts to D_2 , and the new equilibrium price is P_2 and quantity Q_2 .

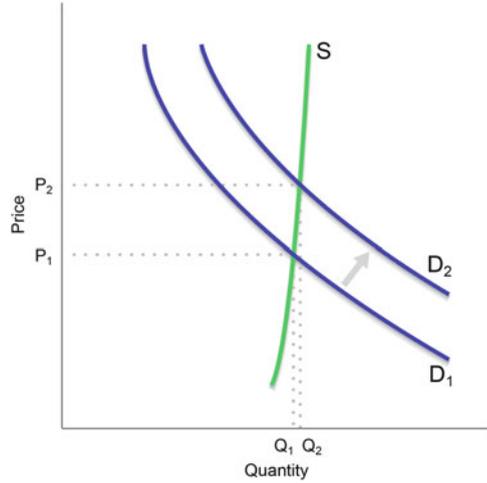
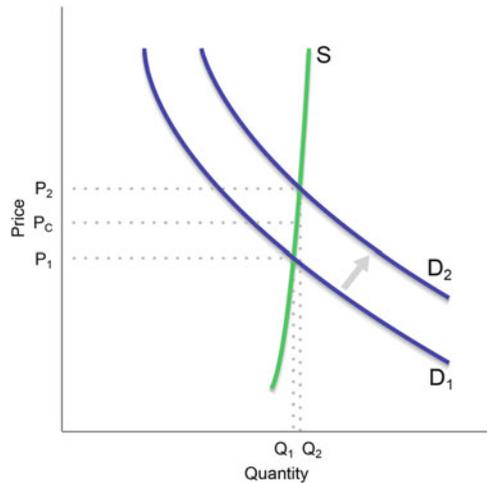


Fig. 5.10 The short run effect of rent control on a local apartment market. In the absence of rent control, a shift from D_1 to D_2 would result in an increase in the equilibrium rental price from P_1 to P_2 . To prevent this, a rent control is imposed, restricting the rise in price to P_C . The result is that the price rise has been limited, with hardly any effect on the number of apartments offered—in the short run.

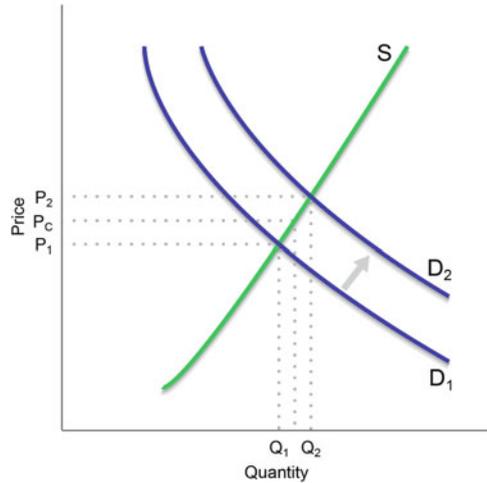


incorporates incentives for new building. It is scarcity, not intervention in the market per se, that leads to the worst aspects of rent control.

5.5 Another Cup of Coffee

Keeping the supply and demand apparatus in mind helps us decipher the changes that roiled the world coffee market. We have already considered the reasons why both the supply and demand curves for coffee are likely to be highly inelastic. (The shapes of these curves could be measured using real-world data, but the technical

Fig. 5.11 The long run effect of rent control on a local apartment market. In the long run, supply is more elastic. Now the effect of imposing rent control at P_C is to reduce the number of apartments supplied to the market well below Q_2 , the number that would be supplied with demand curve D_2 and without rent control



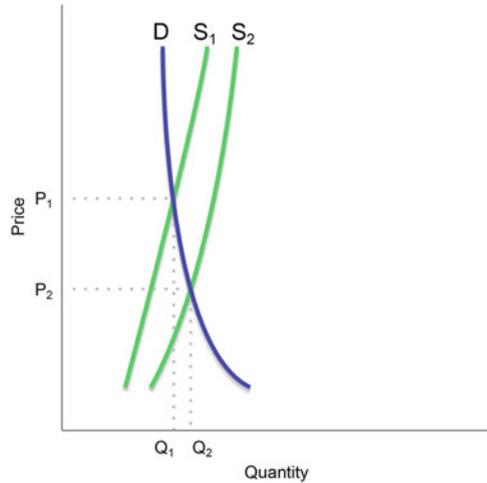
aspects of that procedure would take us away from the subject matter of this chapter.) How about the fluctuations of price and quantity?

The most important single fact to know about coffee is that, until 1989, the major producing countries collaborated on an International Coffee Agreement. This arrangement established export quotas, limiting supply in order to bolster prices. The agreement fell apart in that year, and each country was then free to produce as much as it wanted. It takes about 5 years from the time of initial planting to the first harvest, so it was not until the mid 1990s that the full effect of this change was felt. In 1994/95 world production was about 90 million bags; this total rose to about 115 million bags in 2001/02. The biggest increase came from a new kid on the block, Vietnam. Partly in response to World Bank advice (backed up by loans), Vietnam increased its production from 1.5 million bags in 1990 (when export quotas were lifted) to 15 million in 2000. Brazil was another source of added supply, due to large investments in acreage and new technology. Meanwhile prices for the highest quality coffee, arabica, fell from a high of \$2.00 a pound in 1980 to \$1.30 in 1995 and just over \$.50 in 2002 (all in 1990 dollars). The combination of modest production increases and drastic price decreases tells us what we most need to know about the causes of the coffee crisis. Figure 5.12 on the next page, which is only slightly different from Fig. 5.4, captures these numbers vividly.

With demand so inelastic, it takes only small increases in supply to result in large drops in the market price. Each producing country, by trying to increase its exports so as to make up in quantity what it is losing in value, makes the problem that much worse for all of them. The International Coffee Agreement, which kept the supply curve modestly to the left, was welcomed by consumers but sorely missed by producers.

What can we learn about the uses of supply and demand from this example? The most important lesson is that the supply and demand model of markets does not provide any answers in itself; it is a convenient framework for sorting out

Fig. 5.12 Contours of the coffee crisis. With the ending of export quotas, supply of coffee rose from S_1 to S_2 . This led to a decline in the equilibrium price from P_1 to P_2



information and organizing our thinking. In this respect it is like a language: it provides a syntax that makes it easier to produce certain complex ideas—at the cost (which is also true of language) of making it harder to produce other types of ideas. In this case, we saw that we were able to disentangle information about production and demand. Asking about elasticity drew our attention to certain features of the coffee market that are highly relevant to understanding the crisis, but which might have been overlooked otherwise. Above all, it gave us a simple way to imagine the relationship between relatively small quantity effects—a 7 year rise of 28 %—and very large price effects—a corresponding decline of 62 %.

This illustrates the general point about models. If we are willing to make a number of simplifying assumptions, we can construct potential scenarios for real-world events. The power of models lies in their ability to highlight logical interconnections that might be difficult to see without them. The weakness stems from all the assumptions we have to make, some of which might blind us to important aspects of the problems we care about. What were some of those assumptions in our coffee analysis? We assumed that nothing important happens between the sale of coffee by growers and its ultimate purchase by consumers; all the complications of middlemen, marketing arrangements and brand identity are simply ignored. Also, the international character of this market—its production and distribution across national boundaries, requiring the conversion of different national currencies—is put to the side. (The market is treated as if it all took place at a single location.) Finally, specific sub-markets, such as those for organic, fair-traded and “gourmet”, are not taken into consideration. In other words, the coffee market is not nearly as homogeneous as the model implies. Some of these excluded complexities could be reintroduced through more detailed use of supply and demand analysis, but others are hard to squeeze into this model.

In addition, we should be clear about what we get from using economic models. These devices do not “prove” anything; they assist and illustrate. Moreover, useful

models are flexible. They don't give predetermined answers to questions, but help us gather and organize information. It wasn't the supply and demand model that told us about the importance of the International Coffee Agreement, but instead the information we dug up, prompted by the needs of the model. In another market we might well come to a completely different understanding of the key forces at work; in fact, this will be the case when we revisit the coffee crisis in Chap. 14.

Models, because they purchase their insights at the price of their assumptions, should be used with care. This warning applies with special force to economic models, because the problems they are directed at are extremely complex, and because economics is just one of many perspectives that people have found to be helpful.

The Main Points

1. The following simplifying assumptions make it possible to employ the supply and demand analysis presented in this chapter: (a) each market corresponds to a single category of homogeneous goods, (b) producers and consumers adopt the roles of sellers and buyers, (c) all market participants are self-interested and (d) rational, (e) the law of one price holds, (f) each participant's desired choices are independent of everyone else's except through their effects on the market price, and (g) market "time" is instantaneous and not repeated.
2. The supply curve registers the amount producers of a good or service wish to offer to the market across a range of potential prices. Because it is about what they *want* to offer, it is a behavioral relationship: a given price induces a given desired supply. Of course, the amount offered will be a function of many factors, including the cost of inputs, the technology that transforms inputs into outputs that can be sold, the number and capacity of potential suppliers, expectations of future supply conditions, and the price buyers are willing to pay. To draw a supply curve, all of these factors must be held constant *except* the sale price, so that it is possible to determine the quantity offered at each price. If none of the factors being held constant (the *ceteris paribus* factors) change, there can be only movement *along* the supply curve as demand conditions change. If one or more of the *ceteris paribus* factors changes, there is movement *of* the supply curve: a different quantity will be offered to the market even at the original sale price. The supply curve is normally upward-sloping: an increase in the sale price will typically induce suppliers to offer more of their goods or services to the market.
3. The elasticity of supply tells us how responsive the quantity of market supply is to the price at which the supply can be sold. The formula is:

$$\text{elasticity of supply} = \frac{\% \text{ change in quantity supplied}}{\% \text{ change in price}}$$

If the elasticity of supply is greater than one, we say that supply is "elastic". If it is less than one, we say it is "inelastic". An elastic supply curve is relatively flat

and points toward the vertical axis; an inelastic supply curve is relatively steep and points toward the horizontal axis.

4. The demand curve registers the amount consumers of a good or service wish to buy across a range of potential prices; it too is a behavioral relationship. The amount buyers desire to buy is a function of many factors, such as their personal preferences, their incomes, the price and availability of substitute goods, expectations of future conditions, and the purchase price. We draw the demand curve by holding all other factors constant (*ceteris paribus*) except for the price. If none of these other factors change, there can be only movement along the demand curve as the price changes. If one or more of them does change, we will see movement of the demand curve. The demand curve is normally downward-sloping: an increase in the purchase price will normally lead to a reduction in the amount buyers wish to buy.
5. The price elasticity of demand tells us how responsive the quantity of market demand is to the price buyers have to pay. The formula is:

$$\text{price elasticity of demand} = \frac{\% \text{ change in quantity demanded}}{\% \text{ change in price}}$$

If the (absolute value of the) price elasticity is greater than one, we say that demand is price-elastic. If it is less than one we say it is price-inelastic. An elastic demand curve is relatively flat and situated more toward the NW (upper-left) portion of the price-quantity quadrant; the opposite is true for an inelastic demand curve.

6. Market equilibrium occurs when both behavioral conditions—the supply and demand curves—hold at the same price and quantity. That is, the equilibrium price is the one at which buyers and sellers both want to buy/sell the same quantities, and the equilibrium quantity is the one which would be purchased/sold at the same equilibrium price. In the simple models of this chapter, this is where the two curves intersect. Note, however, that equilibrium is a general concept in economics—a state in which all parties are acting simultaneously according to their preferences, such that their choices are mutually consistent—and it applies to “supply = demand” only under particular conditions. Later we will see examples where equilibrium occurs when supply does not equal demand.
7. A second aspect of equilibrium has to do with the tendency for a situation to return to equilibrium if it departs from it for some reason. In the case of a simple supply and demand model, this results from assumptions about excess demand and supply. We assume that, if there is excess demand, potential buyers who are unable to make purchases will offer to pay higher prices, and sellers will agree to do this: in this way, the price rises until excess demand is eliminated. If there is excess supply, some sellers are unable to find buyers, so we assume they offer to lower their sales price. This continues until the price falls to the point at which excess supply disappears.

8. Many goods that could be produced and sold do not actually exist. Supply and demand theory tells us that there is no price at which at least some producers wish to sell for which there are at least some buyers willing to buy.
9. There is no presumption in economics that equilibrium situations are “good”. The concept of equilibrium is used for predictive and explanatory (positive) purposes only; it has no normative content. A situation could be horrendous for society, and even for most direct participants, and still be an equilibrium.

► Terms to Define

Ceteris paribus

Demand curve

Elastic vs inelastic supply/demand

Elasticity of supply

Excess demand

Excess supply

General equilibrium

Market disequilibrium

Market equilibrium

Movement of vs movement along a curve

Price elasticity of demand

Supply curve

Questions to Consider

1. Suppose you are given the job of estimating the global supply curve for coffee, so that your supply and demand curve can be based on real numbers. How would you go about doing this? First, clarify in your mind exactly the information you will need; then suggest a research strategy. (You can assume you have an unlimited budget at your disposal.)
2. Suggest a good whose supply curve, like coffee, is likely to be highly inelastic. Select another you think will be highly elastic. Explain your reasoning.
3. The amount of oil under the surface of the earth is fixed; no more can be created within a time scale meaningful to human beings. Does this mean that its supply is inelastic?
4. In recent years a vigorous international market has developed in kidneys. Impoverished but healthy people in developing countries are paid to undergo surgical removal of their kidneys, which are then sold to wealthier kidney patients who need a kidney transplant but can't get one through legitimate channels. What would it mean for this market to be in equilibrium? Do you think it should be in equilibrium? If not, what would you propose to do about this problem?
5. For hundreds of years there have been programs, and proposals for programs, to stabilize the supply of agricultural commodities like coffee, building up reserves

after abundant harvests and drawing them down after poor ones. How would you show this process using supply and demand curves? How does this differ from the quota system of the International Coffee Agreement? Which, if either, would you favor as a response to the coffee crisis? Why?