

CHAPTER 31

EXCHANGE

Up until now we have generally considered the market for a single good in isolation. We have viewed the demand and supply functions for a good as depending on its price alone, disregarding the prices of other goods. But in general the prices of other goods *will* affect people's demands and supplies for a particular good. Certainly the prices of substitutes and complements for a good will influence the demand for it, and, more subtly, the prices of goods that people sell will affect the amount of income they have and thereby influence how much of other goods they will be able to buy.

Up until now we have been ignoring the effect of these other prices on the market equilibrium. When we discussed the equilibrium conditions in a particular market, we only looked at part of the problem: how demand and supply were affected by the price of the particular good we were examining. This is called **partial equilibrium** analysis.

In this chapter we will begin our study of **general equilibrium** analysis: how demand and supply conditions interact in several markets to determine the prices of many goods. As you might suspect, this is a complex problem, and we will have to adopt several simplifications in order to deal with it.

First, we will limit our discussion to the behavior of competitive markets, so that each consumer or producer will take prices as given and optimize

accordingly. The study of general equilibrium with imperfect competition is very interesting but too difficult to examine at this point.

Second, we will adopt our usual simplifying assumption of looking at the smallest number of goods and consumers that we possibly can. In this case, it turns out that many interesting phenomena can be depicted using only two goods and two consumers. All of the aspects of general equilibrium analysis that we will discuss can be generalized to arbitrary numbers of consumers and goods, but the exposition is simpler with two of each.

Third, we will look at the general equilibrium problem in two stages. We will start with an economy where people have fixed endowments of goods and examine how they might trade these goods among themselves; no production will be involved. This case is naturally known as the case of **pure exchange**. Once we have a clear understanding of pure exchange markets we will examine production behavior in the general equilibrium model.

31.1 The Edgeworth Box

There is a convenient graphical tool known as the **Edgeworth box** that can be used to analyze the exchange of two goods between two people.¹ The Edgeworth box allows us to depict the endowments and preferences of two individuals in one convenient diagram, which can be used to study various outcomes of the trading process. In order to understand the construction of an Edgeworth box it is necessary to examine the indifference curves and the endowments of the people involved.

Let us call the two people involved A and B and the two goods involved 1 and 2. We will denote A's consumption bundle by $X_A = (x_A^1, x_A^2)$, where x_A^1 represents A's consumption of good 1 and x_A^2 represents A's consumption of good 2. Then B's consumption bundle is denoted by $X_B = (x_B^1, x_B^2)$. A *pair* of consumption bundles, X_A and X_B , is called an **allocation**. An allocation is a **feasible allocation** if the total amount of each good consumed is equal to the total amount available:


$$x_A^1 + x_B^1 = \omega_A^1 + \omega_B^1$$

$$x_A^2 + x_B^2 = \omega_A^2 + \omega_B^2.$$

A particular feasible allocation that is of interest is the **initial endowment allocation**, (ω_A^1, ω_A^2) and (ω_B^1, ω_B^2) . This is the allocation that the consumers start with. It consists of the amount of each good that consumers bring to the market. They will exchange some of these goods with each other in the course of trade to end up at a **final allocation**.

¹ The Edgeworth box is named in honor of Francis Ysidro Edgeworth (1845–1926), an English economist who was one of the first to use this analytical tool.

The Edgeworth box shown in Figure 31.1 can be used to illustrate these concepts graphically. We first use a standard consumer theory diagram to illustrate the endowment and preferences of consumer A. We can also mark off on these axes the *total* amount of each good in the economy—the amount that A has plus the amount that B has of each good. Since we will only be interested in feasible allocations of goods between the two consumers, we can draw a box that contains the set of possible bundles of the two goods that A can hold.



An Edgeworth box. The width of the box measures the total amount of good 1 in the economy and the height measures the total amount of good 2. Person A's consumption choices are measured from the lower left-hand corner while person B's choices are measured from the upper right.

Note that the bundles in this box also indicate the amount of the goods that B can hold. If there are 10 units of good 1 and 20 units of good 2, then if A holds (7,12), B must be holding (3,8). We can depict how much A holds of good 1 by the distance along the horizontal axis from the origin in the lower left-hand corner of the box and the amount B holds of good 1 by measuring the distance along the horizontal axis from the upper right-hand corner. Similarly, distances along the vertical axes give the amounts of good 2 that A and B hold. Thus the points in this box give us both the

bundles that A can hold and the bundles that B can hold—just measured from different origins. The points in the Edgeworth box can represent all feasible allocations in this simple economy.

We can depict A's indifference curves in the usual manner, but B's indifference curves take a somewhat different form. To construct them we take a standard diagram for B's indifference curves, turn it upside down, and "overlay" it on the Edgeworth box. This gives us B's indifference curves on the diagram. If we start at A's origin in the lower left-hand corner and move up and to the right, we will be moving to allocations that are more preferred by A. As we move down and to the left we will be moving to allocations that are more preferred by B. (If you rotate your book and look at the diagram, this discussion may seem clearer.)

The Edgeworth box allows us to depict the possible consumption bundles for both consumers—the feasible allocations—and the preferences of both consumers. It thereby gives a complete description of the economically relevant characteristics of the two consumers.

31.2 Trade

Now that we have both sets of preferences and endowments depicted we can begin to analyze the question of how trade takes place. We start at the original endowment of goods, denoted by the point W in Figure 31.1. Consider the indifference curves of A and B that pass through this allocation. The region where A is better off than at her endowment consists of all the bundles above her indifference curve through W . The region where B is better off than at his endowment consists of all the allocations that are above—from his point of view—*his* indifference curve through W . (This is *below* his indifference curve from *our* point of view . . . unless you've still got your book upside down.)

Where is the region of the box where A and B are *both* made better off? Clearly it is in the intersection of these two regions. This is the lens-shaped region illustrated in Figure 31.1. Presumably in the course of their negotiations the two people involved will find some mutually advantageous trade—some trade that will move them to some point inside the lens-shaped area such as the point M in Figure 31.1.

The particular movement to M depicted in Figure 31.1 involves person A giving up $|x_A^1 - \omega_A^1|$ units of good 1 and acquiring in exchange $|x_A^2 - \omega_A^2|$ units of good 2. This means that B acquires $|x_B^1 - \omega_B^1|$ units of good 1 and gives up $|x_B^2 - \omega_B^2|$ units of good 2.


There is nothing particularly special about the allocation M . Any allocation inside the lens-shaped region would be possible—for every allocation of goods in this region is an allocation that makes each consumer better off than he or she was at the original endowment. We only need to suppose that the consumers trade to *some* point in this region.

Now we can repeat the same analysis at the point M . We can draw the two indifference curves through M , construct a new lens-shaped “region of mutual advantage,” and imagine the traders moving to some new point N in this region. And so it goes . . . the trade will continue until there are no more trades that are preferred by both parties. What does such a position look like?

31.3 Pareto Efficient Allocations

The answer is given in Figure 31.2. At the point M in this diagram the set of points above A’s indifference curve doesn’t intersect the set of points above B’s indifference curve. The region where A is made better off is disjoint from the region where B is made better off. This means that any movement that makes one of the parties better off necessarily makes the other party worse off. Thus there are no exchanges that are advantageous for both parties. There are no mutually improving trades at such an allocation.

An allocation such as this is known as a **Pareto efficient** allocation. The idea of Pareto efficiency is a very important concept in economics that arises in various guises.



A Pareto efficient allocation. At a Pareto efficient allocation such as M , each person is on his highest possible indifference curve, given the indifference curve of the other person. The line connecting such points is known as the contract curve.

A Pareto efficient allocation can be described as an allocation where:

1. There is no way to make all the people involved better off; or
2. there is no way to make some individual better off without making someone else worse off; or
3. all of the gains from trade have been exhausted; or
4. there are no mutually advantageous trades to be made, and so on.

Indeed we have mentioned the concept of Pareto efficiency several times already in the context of a single market: we spoke of the Pareto efficient level of output in a single market as being that amount of output where the marginal willingness to buy equaled the marginal willingness to sell. At any level of output where these two numbers differed, there would be a way to make both sides of the market better off by carrying out a trade. In this chapter we will examine more deeply the idea of Pareto efficiency involving many goods and many traders.

Note the following simple geometry of Pareto efficient allocations: the indifference curves of the two agents must be tangent at any Pareto efficient allocation in the interior of the box. It is easy to see why. If the two indifference curves are not tangent at an allocation in the interior of the box, then they must cross. But if they cross, then there must be some mutually advantageous trade—so that point cannot be Pareto efficient. (It is possible to have Pareto efficient allocations on the sides of the box—where one consumer has zero consumption of some good—in which the indifference curves are not tangent. These boundary cases are not important for the current discussion.)

From the tangency condition it is easy to see that there are a lot of Pareto efficient allocations in the Edgeworth box. In fact, given any indifference curve for person A, for example, there is an easy way to find a Pareto efficient allocation. Simply move along A's indifference curve until you find the point that is the best point for B. This will be a Pareto efficient point, and thus both indifference curves must be tangent at this point.

The set of *all* Pareto efficient points in the Edgeworth box is known as the **Pareto set**, or the **contract curve**. The latter name comes from the idea that all “final contracts” for trade must lie on the Pareto set—otherwise they wouldn't be final because there would be some improvement that could be made!

In a typical case the contract curve will stretch from A's origin to B's origin across the Edgeworth box, as shown in Figure 31.2. If we start at A's origin, A has none of either good and B holds everything. This is Pareto efficient since the only way A can be made better off is to take something away from B. As we move up the contract curve A is getting more and more well-off until we finally get to B's origin.

The Pareto set describes all the possible outcomes of mutually advantageous trade from starting anywhere in the box. If we are given the starting point—the initial endowments for each consumer—we can look at the subset of the Pareto set that each consumer prefers to his initial endowment. This is simply the subset of the Pareto set that lies in the lens-shaped region depicted in Figure 31.1. The allocations in this lens-shaped region are the possible outcomes of mutual trade starting from the particular initial endowment depicted in that diagram. But the Pareto set itself doesn't depend on the initial endowment, except insofar as the endowment determines the total amounts of both goods that are available and thereby determines the dimensions of the box.

31.4 Market Trade


The equilibrium of the trading process described above—the set of Pareto efficient allocations—is very important, but it still leaves a lot of ambiguity about where the agents end up. The reason is that the trading process we have described is very general. Essentially we have only assumed that the two parties will move to *some* allocation where they are both made better off.

If we have a *particular* trading process, we will have a more precise description of equilibrium. Let's try to describe a trading process that mimics the outcome of a competitive market.

Suppose that we have a third party who is willing to act as an “auctioneer” for the two agents A and B. The auctioneer chooses a price for good 1 and a price for good 2 and presents these prices to the agents A and B. Each agent then sees how much his or her endowment is worth at the prices (p_1, p_2) and decides how much of each good he or she would want to buy at those prices.

One warning is in order here. If there are really only two people involved in the transaction, then it doesn't make much sense for them to behave in a competitive manner. Instead they would probably attempt to bargain over the terms of trade. One way around this difficulty is to think of the Edgeworth box as depicting the average demands in an economy with only two *types* of consumers, but with many consumers of each type. Another way to deal with this is to point out that the behavior is implausible in the two-person case, but it makes perfect sense in the many-person case, which is what we are really concerned with.

Either way, we know how to analyze the consumer-choice problem in this framework—it is just the standard consumer-choice problem we described in Chapter 5. In Figure 31.3 we illustrate the two demanded bundles of the two agents. (Note that the situation depicted in Figure 31.3 is not an equilibrium configuration since the demand by one agent is not equal to the supply of the other agent.)



Gross demands and net demands. Gross demands are the amounts the person wants to consume; net demands are the amounts the person wants to purchase.

As in Chapter 9 there are two relevant concepts of “demand” in this framework. The **gross demand** of agent A for good 1, say, is the total amount of good 1 that he wants at the going prices. The **net demand** of agent A for good 1 is the difference between this total demand and the initial endowment of good 1 that agent A holds. In the context of general equilibrium analysis, net demands are sometimes called **excess demands**. We will denote the excess demand of agent A for good 1 by e_A^1 . By definition, if A’s gross demand is x_A^1 , and his endowment is ω_A^1 , we have

$$e_A^1 = x_A^1 - \omega_A^1.$$

The concept of excess demand is probably more natural, but the concept of gross demand is generally more useful. We will typically use the word “demand” to mean gross demand and specifically say “net demand” or “excess demand” if that is what we mean.

For arbitrary prices (p_1, p_2) there is no guarantee that supply will equal demand—in either sense of demand. In terms of net demand, this means that the amount that A wants to buy (or sell) will not necessarily equal the amount that B wants to sell (or buy). In terms of gross demand, this means that the total amount that the two agents want hold of the goods is not equal to the total amount of that goods available. Indeed, this is true in the example depicted in Figure 31.3. In this example the agents will not

be able to complete their desired transactions: the markets will not clear.

We say that in this case the market is in **disequilibrium**. In such a situation, it is natural to suppose that the auctioneer will change the prices of the goods. If there is excess demand for one of the goods, the auctioneer will raise the price of that good, and if there is excess supply for one of the goods, the auctioneer will lower its price.

Suppose that this adjustment process continues until the demand for each of the goods equals the supply. What will the final configuration look like?

The answer is given in Figure 31.4. Here the amount that A wants to buy of good 1 just equals the amount that B wants to sell of good 1, and similarly for good 2. Said another way, the total amount that each person wants to buy of each good at the current prices is equal to the total amount available. We say that the market is in **equilibrium**. More precisely, this is called a **market equilibrium**, a **competitive equilibrium**, or a **Walrasian equilibrium**.² Each of these terms refers to the same thing: a set of prices such that each consumer is choosing his or her most-preferred affordable bundle, and all consumers' choices are compatible in the sense that demand equals supply in every market.

We know that if each agent is choosing the best bundle that he can afford, then his marginal rate of substitution between the two goods must be equal to the ratio of the prices. But if all consumers are facing the same prices, then all consumers will have to have the *same* marginal rate of substitution between each of the two goods. In terms of Figure 31.4, an equilibrium has the property that each agent's indifference curve is tangent to his budget line. But since each agent's budget line has the slope $-p_1/p_2$, this means that the two agents' indifference curves must be tangent to each other.

31.5 The Algebra of Equilibrium


If we let $x_A^1(p_1, p_2)$ be agent A's demand function for good 1 and $x_B^1(p_1, p_2)$ be agent B's demand function for good 1, and define the analogous expressions for good 2, we can describe this equilibrium as a set of prices (p_1^*, p_2^*) such that

$$x_A^1(p_1^*, p_2^*) + x_B^1(p_1^*, p_2^*) = \omega_A^1 + \omega_B^1$$

$$x_A^2(p_1^*, p_2^*) + x_B^2(p_1^*, p_2^*) = \omega_A^2 + \omega_B^2.$$

These equations say that in equilibrium the total demand for each good should be equal to the total supply.

² Leon Walras (1834–1910) was a French economist at Lausanne who was an early investigator of general equilibrium theory.



Equilibrium in the Edgeworth box. In equilibrium, each person is choosing the most-preferred bundle in his budget set, and the choices exhaust the available supply.

Another way to describe the equilibrium is to rearrange these two equations to get

$$[x_A^1(p_1^*, p_2^*) - \omega_A^1] + [x_B^1(p_1^*, p_2^*) - \omega_B^1] = 0$$

$$[x_A^2(p_1^*, p_2^*) - \omega_A^2] + [x_B^2(p_1^*, p_2^*) - \omega_B^2] = 0.$$

These equations say that the sum of *net demands* of each agent for each good should be zero. Or, in other words, the net amount that A chooses to demand (or supply) must be equal to the net amount that B chooses to supply (or demand).

Yet another formulation of these equilibrium equations comes from the concept of the **aggregate excess demand function**. Let us denote the net demand function for good 1 by agent A by

$$e_A^1(p_1, p_2) = x_A^1(p_1, p_2) - \omega_A^1$$

and define $e_B^1(p_1, p_2)$ in a similar manner.

The function $e_A^1(p_1, p_2)$ measures agent A's **net demand** or his **excess demand**—the difference between what she wants to consume of good 1 and what she initially has of good 1. Now let us add together agent A's net demand for good 1 and agent B's net demand for good 1. We get

$$\begin{aligned} z_1(p_1, p_2) &= e_A^1(p_1, p_2) + e_B^1(p_1, p_2) \\ &= x_A^1(p_1, p_2) + x_B^1(p_1, p_2) - \omega_A^1 - \omega_B^1, \end{aligned}$$

which we call the **aggregate excess demand** for good 1. There is a similar aggregate excess demand for good 2, which we denote by $z_2(p_1, p_2)$.

Then we can describe an equilibrium (p_1^*, p_2^*) by saying that the aggregate excess demand for each good is zero:

$$\begin{aligned} z_1(p_1^*, p_2^*) &= 0 \\ z_2(p_1^*, p_2^*) &= 0. \end{aligned}$$

Actually, this definition is stronger than necessary. It turns out that if the aggregate excess demand for good 1 is zero, then the aggregate excess demand for good 2 must necessarily be zero. In order to prove this, it is convenient to first establish a property of the aggregate excess demand function known as **Walras' law**.

31.6 Walras' Law

Using the notation established above, Walras' law states that

$$p_1 z_1(p_1, p_2) + p_2 z_2(p_1, p_2) \equiv 0.$$

That is, *the value of aggregate excess demand is identically zero*. To say that the value of aggregate demand is identically zero means that it is zero for *all* possible choices of prices, not just equilibrium prices.

The proof of this follows from adding up the two agents' budget constraints. Consider first agent A. Since her demand for each good satisfies her budget constraint, we have

$$p_1 x_A^1(p_1, p_2) + p_2 x_A^2(p_1, p_2) \equiv p_1 \omega_A^1 + p_2 \omega_A^2$$

or

$$\begin{aligned} p_1 [x_A^1(p_1, p_2) - \omega_A^1] + p_2 [x_A^2(p_1, p_2) - \omega_A^2] &\equiv 0 \\ p_1 e_A^1(p_1, p_2) + p_2 e_A^2(p_1, p_2) &\equiv 0. \end{aligned}$$

This equation says that the *value of agent A's net demand is zero*. That is, the value of how much A wants to buy of good 1 plus the value of how much she wants to buy of good 2 must equal zero. (Of course the amount that she wants to buy of *one* of the goods must be negative—that is, she intends to sell some of one of the goods to buy more of the other.)

We have a similar equation for agent B:

$$\begin{aligned} p_1 [x_B^1(p_1, p_2) - \omega_B^1] + p_2 [x_B^2(p_1, p_2) - \omega_B^2] &\equiv 0 \\ p_1 e_B^1(p_1, p_2) + p_2 e_B^2(p_1, p_2) &\equiv 0. \end{aligned}$$

Adding the equations for agent A and agent B together and using the definition of aggregate excess demand, $z_1(p_1, p_2)$ and $z_2(p_1, p_2)$, we have

$$\begin{aligned} p_1 [e_A^1(p_1, p_2) + e_B^1(p_1, p_2)] + p_2 [e_A^2(p_1, p_2) + e_B^2(p_1, p_2)] &\equiv 0 \\ p_1 z_1(p_1, p_2) + p_2 z_2(p_1, p_2) &\equiv 0. \end{aligned}$$

Now we can see where Walras' law comes from: since the value of each agent's excess demand equals zero, the value of the sum of the agents' excess demands must equal zero.

We can now demonstrate that if demand equals supply in one market, demand must also equal supply in the other market. Note that Walras' law must hold for all prices, since each agent must satisfy his or her budget constraint for all prices. Since Walras' law holds for all prices, in particular, it holds for a set of prices where the excess demand for good 1 is zero:

$$z_1(p_1^*, p_2^*) = 0.$$

According to Walras' law it must also be true that

$$p_1^* z_1(p_1^*, p_2^*) + p_2^* z_2(p_1^*, p_2^*) = 0.$$

It easily follows from these two equations that if $p_2 > 0$, then we must have

$$z_2(p_1^*, p_2^*) = 0.$$

Thus, as asserted above, if we find a set of prices (p_1^*, p_2^*) where the demand for good 1 equals the supply of good 1, we are guaranteed that the demand for good 2 must equal the supply of good 2. Alternatively, if we find a set of prices where the demand for good 2 equals the supply of good 2, we are guaranteed that market 1 will be in equilibrium.

In general, if there are markets for k goods, then we only need to find a set of prices where $k - 1$ of the markets are in equilibrium. Walras' law then implies that the market for good k will automatically have demand equal to supply.

31.7 Relative Prices

As we've seen above, Walras' law implies that there are only $k - 1$ independent equations in a k -good general equilibrium model: if demand equals supply in $k - 1$ markets, demand must equal supply in the final market. But if there are k goods, there will be k prices to be determined. How can you solve for k prices with only $k - 1$ equations?

The answer is that there are really only $k - 1$ *independent* prices. We saw in Chapter 2 that if we multiplied all prices and income by a positive number t , then the budget set wouldn't change, and thus the demanded bundle wouldn't change either. In the general equilibrium model, each consumer's income is just the value of his or her endowment at the market prices. If we multiply all prices by $t > 0$, we will automatically multiply each consumer's income by t . Thus, if we find some equilibrium set of prices (p_1^*, p_2^*) , then (tp_1^*, tp_2^*) are equilibrium prices as well, for any $t > 0$.

This means that we are free to choose one of the prices and set it equal to a constant. In particular it is often convenient to set one of the prices equal to 1 so that all of the other prices can be interpreted as being measured relative to it. As we saw in Chapter 2, such a price is called a **numeraire** price. If we choose the first price as the numeraire price, then it is just like multiplying all prices by the constant $t = 1/p_1$.

The requirement that demand equal supply in every market can only be expected to determine the equilibrium relative prices, since multiplying all prices by a positive number will not change anybody's demand and supply behavior.

EXAMPLE: An Algebraic Example of Equilibrium

The Cobb-Douglas utility function described in Chapter 6 has the form $u_A(x_A^1, x_A^2) = (x_A^1)^a (x_A^2)^{1-a}$ for person A, and a similar form for person B. We saw there that this utility function gave rise to the following demand functions:

$$\begin{aligned} x_A^1(p_1, p_2, m_A) &= a \frac{m_A}{p_1} \\ x_A^2(p_1, p_2, m_A) &= (1-a) \frac{m_A}{p_2} \\ x_B^1(p_1, p_2, m_B) &= b \frac{m_B}{p_1} \\ x_B^2(p_1, p_2, m_B) &= (1-b) \frac{m_B}{p_2}, \end{aligned}$$

where a and b are the parameters of the two consumers' utility functions.

We know that in equilibrium, the money income of each individual is given by the value of his or her endowment:

$$\begin{aligned} m_A &= p_1 \omega_A^1 + p_2 \omega_A^2 \\ m_B &= p_1 \omega_B^1 + p_2 \omega_B^2. \end{aligned}$$

Thus the aggregate excess demands for the two goods are

$$\begin{aligned} z_1(p_1, p_2) &= a \frac{m_A}{p_1} + b \frac{m_B}{p_1} - \omega_A^1 - \omega_B^1 \\ &= a \frac{p_1 \omega_A^1 + p_2 \omega_A^2}{p_1} + b \frac{p_1 \omega_B^1 + p_2 \omega_B^2}{p_1} - \omega_A^1 - \omega_B^1 \end{aligned}$$

and

$$\begin{aligned} z_2(p_1, p_2) &= (1-a) \frac{m_A}{p_2} + (1-b) \frac{m_B}{p_2} - \omega_A^2 - \omega_B^2 \\ &= (1-a) \frac{p_1 \omega_A^1 + p_2 \omega_A^2}{p_2} + (1-b) \frac{p_1 \omega_B^1 + p_2 \omega_B^2}{p_2} - \omega_A^2 - \omega_B^2. \end{aligned}$$

You should verify that these aggregate demand functions satisfy Walras' law.

Let us choose p_2 as the numeraire price, so that these equations become

$$z_1(p_1, 1) = a \frac{p_1 \omega_A^1 + \omega_A^2}{p_1} + b \frac{p_1 \omega_B^1 + \omega_B^2}{p_1} - \omega_A^1 - \omega_B^1$$

$$z_2(p_1, 1) = (1 - a)(p_1 \omega_A^1 + \omega_A^2) + (1 - b)(p_1 \omega_B^1 + \omega_B^2) - \omega_A^2 - \omega_B^2.$$

All we've done here is set $p_2 = 1$.

We now have an equation for the excess demand for good 1, $z_1(p_1, 1)$, and an equation for the excess demand for good 2, $z_2(p_1, 1)$, with each equation expressed as a function of the relative price of good 1, p_1 . In order to find the *equilibrium* price, we set either of these equations equal to zero and solve for p_1 . According to Walras' law, we should get the same equilibrium price, no matter which equation we solve.

The equilibrium price turns out to be

$$p_1^* = \frac{a\omega_A^2 + b\omega_B^2}{(1 - a)\omega_A^1 + (1 - b)\omega_B^1}.$$

(Skeptics may want to insert this value of p_1 into the demand equals supply equations to verify that the equations are satisfied.)

31.8 The Existence of Equilibrium

In the example given above, we had specific equations for each consumer's demand function and we could explicitly solve for the equilibrium prices. But in general, we don't have explicit algebraic formulas for each consumer's demands. We might well ask how do we know that there is *any* set of prices such that demand equals supply in every market? This is known as the question of the **existence of a competitive equilibrium**.

The existence of a competitive equilibrium is important insofar as it serves as a "consistency check" for the various models that we have examined in previous chapters. What use would it be to build up elaborate theories of the workings of a competitive equilibrium if such an equilibrium commonly did not exist?

Early economists noted that in a market with k goods there were $k - 1$ relative prices to be determined, and there were $k - 1$ equilibrium equations stating that demand should equal supply in each market. Since the number of equations equaled the number of unknowns, they asserted that there would be a solution where all of the equations were satisfied.

Economists soon discovered that such arguments were fallacious. Merely counting the number of equations and unknowns is not sufficient to prove that an equilibrium solution will exist. However, there are mathematical

tools that can be used to establish the existence of a competitive equilibrium. The crucial assumption turns out to be that the aggregate excess demand function is a **continuous function**. This means, roughly speaking, that small changes in prices should result in only small changes in aggregate demand: a small change in prices should not result in a big jump in the quantity demanded.

Under what conditions will the aggregate demand functions be continuous? Essentially there are two kinds of conditions that will guarantee continuity. One is that each individual's demand function be continuous—that small changes in prices will lead to only small changes in demand. This turns out to require that each consumer have convex preferences, which we discussed in Chapter 3. The other condition is more general. Even if consumers themselves have discontinuous demand behavior, as long as all consumers are small relative to the size of the market, the aggregate demand function will be continuous.

This latter condition is quite nice. After all, the assumption of competitive behavior only makes sense when there are a lot of consumers who are small relative to the size of the market. This is exactly the condition that we need in order to get the aggregate demand functions to be continuous. And continuity is just the ticket to ensure that a competitive equilibrium exists. Thus the very assumptions that make the postulated behavior reasonable will ensure that the equilibrium theory will have content.

31.9 Equilibrium and Efficiency

We have now analyzed market trade in a pure exchange model. This gives us a specific model of trade that we can compare to the general model of trade that we discussed in the beginning of this chapter. One question that might arise about the use of a competitive market is whether this mechanism can really exhaust all of the gains from trade. After we have traded to a competitive equilibrium where demand equals supply in every market, will there be any more trades that people will desire to carry out? This is just another way to ask whether the market equilibrium is Pareto efficient: will the agents desire to make any more trades after they have traded at the competitive prices?

We can see the answer by inspecting Figure 31.4: it turns out that the market equilibrium allocation *is* Pareto efficient. The proof is this: an allocation in the Edgeworth box is Pareto efficient if the set of bundles that A prefers doesn't intersect the set of bundles that B prefers. But at the market equilibrium, the set of bundles preferred by A must lie above her budget set, and the same thing holds for B, where "above" means "above from B's point of view." Thus the two sets of preferred allocations can't intersect. This means that there are no allocations that both agents prefer to the equilibrium allocation, so the equilibrium is Pareto efficient.

31.10 The Algebra of Efficiency

We can also show this algebraically. Suppose that we have a market equilibrium that is *not* Pareto efficient. We will show that this assumption leads to a logical contradiction.

To say that the market equilibrium is not Pareto efficient means that there is some other feasible allocation $(y_A^1, y_A^2, y_B^1, y_B^2)$ such that

$$y_A^1 + y_B^1 = \omega_A^1 + \omega_B^1 \quad (31.1)$$

$$y_A^2 + y_B^2 = \omega_A^2 + \omega_B^2 \quad (31.2)$$

and

$$(y_A^1, y_A^2) \succ_A (x_A^1, x_A^2) \quad (31.3)$$

$$(y_B^1, y_B^2) \succ_B (x_B^1, x_B^2). \quad (31.4)$$

The first two equations say that the y -allocation is feasible, and the next two equations say that it is preferred by each agent to the x -allocation. (The symbols $\succ_A$ and $\succ_B$ refer to the preferences of agents A and B .)

But by hypothesis, we have a market equilibrium where each agent is purchasing the best bundle he or she can afford. If (y_A^1, y_A^2) is better than the bundle that A is choosing, then it must cost more than A can afford, and similarly for B :

$$p_1 y_A^1 + p_2 y_A^2 > p_1 \omega_A^1 + p_2 \omega_A^2$$

$$p_1 y_B^1 + p_2 y_B^2 > p_1 \omega_B^1 + p_2 \omega_B^2.$$

Now add these two equations together to get

$$p_1(y_A^1 + y_B^1) + p_2(y_A^2 + y_B^2) > p_1(\omega_A^1 + \omega_B^1) + p_2(\omega_A^2 + \omega_B^2).$$

Substitute from equations (31.1) and (31.2) to get

$$p_1(\omega_A^1 + \omega_B^1) + p_2(\omega_A^2 + \omega_B^2) > p_1(\omega_A^1 + \omega_B^1) + p_2(\omega_A^2 + \omega_B^2),$$

which is clearly a contradiction, since the left-hand side and the right-hand side are the same.

We derived this contradiction by assuming that the market equilibrium was not Pareto efficient. Therefore, this assumption must be wrong. It follows that all market equilibria are Pareto efficient: a result known as the **First Theorem of Welfare Economics**.

The First Welfare Theorem guarantees that a competitive market will exhaust all of the gains from trade: an equilibrium allocation achieved by a

set of competitive markets will necessarily be Pareto efficient. Such an allocation may not have any other desirable properties, but it will necessarily be efficient.

In particular, the First Welfare Theorem says nothing about the distribution of economic benefits. The market equilibrium might not be a “just” allocation—if person A owned everything to begin with, then she would own everything after trade. That would be efficient, but it would probably not be very fair. But, after all, efficiency does count for something, and it is reassuring to know that a simple market mechanism like the one we have described is capable of achieving an efficient allocation.

EXAMPLE: Monopoly in the Edgeworth Box

In order to understand the First Welfare Theorem better, it is useful to consider another resource allocation mechanism that does not lead to efficient outcomes. A nice example of this occurs when one consumer attempts to behave as a monopolist. Suppose now that there is no auctioneer and that instead, agent A is going to quote prices to agent B, and agent B will decide how much he wants to trade at the quoted prices. Suppose further that A knows B’s “demand curve” and will attempt to choose the set of prices that makes A as well-off as possible, given the demand behavior of B.


In order to examine the equilibrium in this process, it is appropriate to recall the definition of a consumer’s **price offer curve**. The price offer curve, which we discussed in Chapter 6, represents all of the optimal choices of the consumer at different prices. B’s offer curve represents the bundles that he will purchase at different prices; that is, it describes B’s demand behavior. If we draw a budget line for B, then the point where that budget line intersects his offer curve represents B’s optimal consumption.

Thus, if agent A wants to choose the prices to offer to B that make A as well-off as possible, she should find that point on B’s offer curve where A has the highest utility. Such a choice is depicted in Figure 31.5.

This optimal choice will be characterized by a tangency condition as usual: A’s indifference curve will be tangent to B’s offer curve. If B’s offer curve cut A’s indifference curve, there would be some point on B’s offer curve that A preferred—so we couldn’t be at the optimal point for A.

Once we have identified this point—denoted by X in Figure 31.5—we just draw a budget line to that point from the endowment. At the prices that generate this budget line, B will choose the bundle X , and A will be as well-off as possible.

Is this allocation Pareto efficient? In general the answer is no. To see this simply note that A’s indifference curve will not be tangent to the budget line at X , and therefore A’s indifference curve will not be tangent to B’s indifference curve. A’s indifference curve is tangent to B’s *offer curve*,




Monopoly in the Edgeworth box. A chooses the point on B's offer curve that gives him the highest utility.

but it cannot then be tangent to B's indifference curve. The monopoly allocation is Pareto inefficient.

In fact, it is Pareto inefficient in exactly the same way as described in the discussion of monopoly in Chapter 24. At the margin A would like to sell more at the equilibrium prices, but she can only do so by lowering the price at which she sells—and this will lower her income received from all her inframarginal sales.

We saw in Chapter 25 that a perfectly discriminating monopolist would end up producing an efficient level of output. Recall that a discriminating monopolist was one who was able to sell each unit of a good to the person who was willing to pay the most for that unit. What does a perfectly discriminating monopolist look like in the Edgeworth box?

The answer is depicted in Figure 31.6. Let us start at the initial endowment, W , and imagine A selling each unit of good 1 to B at a different price—the price at which B is just indifferent between buying or not buying that unit of the good. Thus, after A sells the first unit, B will remain on the same indifference curve through W . Then A sells the second unit of good 1 to B for the maximum price he is willing to pay. This means that the allocation moves further to the left, but remains on B's indifference curve through W . Agent A continues to sell units to B in this manner, thereby moving up B's indifference curve to find her—A's—most preferred point, denoted by an X in Figure 31.6.



A perfectly discriminating monopolist. Person A chooses the point X on person B's indifference curve through the endowment that gives him the highest utility. Such a point must be Pareto efficient.


It is easy to see that such a point must be Pareto efficient. Agent A will be as well-off as possible given B's indifference curve. At such a point, A has managed to extract all of B's consumer's surplus: B is no better off than he was at his endowment.

These two examples provide useful benchmarks with which to think about the First Welfare Theorem. The ordinary monopolist gives an example of a resource allocation mechanism that results in inefficient equilibria, and the discriminating monopolist gives another example of a mechanism that results in efficient equilibria.

31.11 Efficiency and Equilibrium

The First Welfare Theorem says that the equilibrium in a set of competitive markets is Pareto efficient. What about the other way around? Given a Pareto efficient allocation, can we find prices such that it is a market equilibrium? It turns out that the answer is yes, under certain conditions. The argument is illustrated in Figure 31.7.

Let us pick a Pareto efficient allocation. Then we know that the set of allocations that A prefers to her current assignment is disjoint from the set that B prefers. This implies of course that the two indifference curves are




tangent at the Pareto efficient allocation. So let us draw in the straight line that is their common tangent, as in Figure 31.7.

Suppose that the straight line represents the agents' budget sets. Then if each agent chooses the best bundle on his or her budget set, the resulting equilibrium will be the original Pareto efficient allocation.

Thus the fact that the original allocation is efficient automatically determines the equilibrium prices. The endowments can be any bundles that give rise to the appropriate budget set—that is, bundles that lie somewhere on the constructed budget line.

Can the construction of such a budget line always be carried out? Unfortunately, the answer is no. Figure 31.8 gives an example. Here the illustrated point X is Pareto efficient, but there are no prices at which A and B will want to consume at point X . The most obvious candidate is drawn in the diagram, but the optimal demands of agents A and B don't coincide for that budget. Agent A wants to demand the bundle Y , but agent B wants the bundle X —demand does not equal supply at these prices.

The difference between Figure 31.7 and Figure 31.8 is that the preferences in Figure 31.7 are convex while the ones in Figure 31.8 are not. If the preferences of both agents are convex, then the common tangent will not intersect either indifference curve more than once, and everything will work out fine. This observation gives us the **Second Theorem of Welfare**



Economics: if all agents have convex preferences, then there will always be a set of prices such that each Pareto efficient allocation is a market equilibrium for an appropriate assignment of endowments.

The proof is essentially the geometric argument we gave above. At a Pareto efficient allocation, the bundles preferred by agent A and by agent B must be disjoint. Thus if both agents have convex preferences we can draw a straight line between the two sets of preferred bundles that separates one from the other. The slope of this line gives us the relative prices, and any endowment that puts the two agents on this line will lead to the final market equilibrium being the original Pareto efficient allocation.

31.12 Implications of the First Welfare Theorem

The two theorems of welfare economics are among the most fundamental results in economics. We have demonstrated the theorems only in the simple Edgeworth box case, but they are true for much more complex models with arbitrary numbers of consumers and goods. The welfare theorems have profound implications for the design of ways to allocate resources.

Let us consider the First Welfare Theorem. This says that any competitive equilibrium is Pareto efficient. There are hardly any explicit assump-

tions in this theorem—it follows almost entirely from the definitions. But there are some implicit assumptions. One major assumption is that agents only care about their own consumption of goods, and not about what other agents consume. If one agent does care about another agent's consumption, we say that there is a **consumption externality**. We shall see that when consumption externalities are present, a competitive equilibrium need not be Pareto efficient.

To take a simple example, suppose that agent A cares about agent B's consumption of cigars. Then there is no particular reason why each agent choosing his or her own consumption bundle at the market prices will result in a Pareto efficient allocation. After each person has purchased the best bundle he or she can afford, there may still be ways to make both of them better off—such as A paying B to smoke fewer cigars. We will discuss externalities in more detail in Chapter 33.

Another important implicit assumption in the First Welfare Theorem is that agents actually behave competitively. If there really were only two agents, as in the Edgeworth box example, then it is unlikely that they would each take price as given. Instead, the agents would probably recognize their market power and would attempt to use their market power to improve their own positions. The concept of competitive equilibrium only makes sense when there are enough agents to ensure that each behaves competitively.

Finally, the First Welfare Theorem is only of interest if a competitive equilibrium actually exists. As we have argued above, this will be the case if the consumers are sufficiently small relative to the size of the market.

Given these provisos, the First Welfare Theorem is a pretty strong result: a private market, with each agent seeking to maximize his or her own utility, will result in an allocation that achieves Pareto efficiency.

The importance of the First Welfare Theorem is that it gives a general mechanism—the competitive market—that we can use to ensure Pareto efficient outcomes. If there are only two agents involved, this doesn't matter very much; it is easy for two people to get together and examine the possibilities for mutual trades. But if there are thousands, or even millions, of people involved there must be some kind of structure imposed on the trading process. The First Welfare Theorem shows that the particular structure of competitive markets has the desirable property of achieving a Pareto efficient allocation.

If we are dealing with a resource problem involving many people, it is important to note that the use of competitive markets economizes on the information that any one agent needs to possess. The only things that a consumer needs to know to make his consumption decisions are the prices of the goods he is considering consuming. Consumers don't need to know anything about how the goods are produced, or who owns what goods, or where the goods come from in a competitive market. If each consumer knows only the prices of the goods, he can determine his demands, and if

the market functions well enough to determine the competitive prices, we are guaranteed an efficient outcome. The fact that competitive markets economize on information in this way is a strong argument in favor of their use as a way to allocate resources.

31.13 Implications of the Second Welfare Theorem

The Second Theorem of Welfare Economics asserts that under certain conditions, every Pareto efficient allocation can be achieved as a competitive equilibrium.

What is the meaning of this result? The Second Welfare Theorem implies that the problems of distribution and efficiency can be separated. Whatever Pareto efficient allocation you want can be supported by the market mechanism. The market mechanism is distributionally neutral; whatever your criteria for a good or a just distribution of welfare, you can use competitive markets to achieve it.

Prices play two roles in the market system: an *allocative* role and a *distributive* role. The allocative role of prices is to indicate relative scarcity; the distributive role is to determine how much of different goods different agents can purchase. The Second Welfare Theorem says that these two roles can be separated: we can redistribute endowments of goods to determine how much wealth agents have, and then use prices to indicate relative scarcity.

Policy discussions often become confused on this point. One often hears arguments for intervening in pricing decisions on grounds of distributional equity. However, such intervention is typically misguided. As we have seen above, a convenient way to achieve efficient allocations is for each agent to face the true social costs of his or her actions and to make choices that reflect those costs. Thus in a perfectly competitive market the marginal decision of whether to consume more or less of some good will depend on the price—which measures how everyone else values this good on the margin. The considerations of efficiency are inherently marginal decisions—each person should face the correct marginal tradeoff in making his or her consumption decisions.

The decision about *how much* different agents should consume is a totally different issue. In a competitive market this is determined by the value of the resources that a person has to sell. From the viewpoint of the pure theory, there is no reason why the state can't transfer purchasing power—endowments—among consumers in any way that is seen fit.

In fact the state doesn't need to transfer the physical endowments themselves. All that is necessary is to transfer the purchasing power of the endowment. The state could tax one consumer on the basis of the value of his endowment and transfer this money to another. As long as the taxes are based on the value of the consumer's *endowment* of goods there will

be no loss of efficiency. It is only when taxes depend on the *choices* that a consumer makes that inefficiencies result, since in this case, the taxes will affect the consumer's marginal choices.

It is true that a tax on endowments will generally change people's behavior. But, according to the First Welfare Theorem, trade from any initial endowments will result in a Pareto efficient allocation. Thus no matter how one redistributes endowments, the equilibrium allocation as determined by market forces will still be Pareto efficient.

However, there are practical matters involved. It would be easy to have a lump-sum tax on consumers. We could tax all consumers with blue eyes, and redistribute the proceeds to consumers with brown eyes. As long as eye color can't be changed, there would be no loss in efficiency. Or we could tax consumers with high IQs and redistribute the funds to consumers with low IQs. Again, as long as IQ can be measured, there is no efficiency loss in this kind of tax.

But there's the problem. How do we measure people's endowment of goods? For most people, the bulk of their endowment consists of their own labor power. People's endowments of labor consist of the labor that they *could* consider selling, not the amount of labor that they actually end up selling. Taxing labor that people decide to sell to the market is a **distortionary tax**. If the sale of labor is taxed, the labor supply decision of consumers will be distorted—they will likely supply less labor than they would have supplied in the absence of a tax. Taxing the potential value of labor—the endowment of labor—is not distortionary. The potential value of labor is, by definition, something that is not changed by taxation. Taxing the value of the endowment sounds easy until we realize that it involves identifying and taxing something that *might* be sold, rather than taxing something that is sold.

We could *imagine* a mechanism for levying this kind of tax. Suppose that we considered a society where each consumer was required to give the money earned in 10 hours of his labor time to the state each week. This kind of tax would be independent of how much the person actually worked—it would only depend on the endowment of labor, not on how much was actually sold. Such a tax is basically transferring some part of each consumer's endowment of labor time to the state. The state could then use these funds to provide various goods, or it could simply transfer these funds to other agents.

According to the Second Welfare Theorem, this kind of lump-sum taxation would be nondistortionary. Essentially any Pareto efficient allocation could be achieved by such lump-sum redistribution.

However, no one is advocating such a radical restructuring of the tax system. Most people's labor supply decisions are relatively insensitive to variations in the wage rate, so the efficiency loss from taxing labor may not be too large anyway. But the message of the Second Welfare Theorem is important. Prices should be used to reflect scarcity. Lump-sum transfers of

wealth should be used to adjust for distributional goals. To a large degree, these two policy decisions can be separated.

People's concern about the distribution of welfare can lead them to advocate various forms of manipulation of prices. It has been argued, for example, that senior citizens should have access to less expensive telephone service, or that small users of electricity should pay lower rates than large users. These are basically attempts to redistribute income through the price system by offering some people lower prices than others.

When you think about it this is a terribly inefficient way to redistribute income. If you want to redistribute income, why don't you simply redistribute income? If you give a person an extra dollar to spend, then he can choose to consume more of any of the goods that he wants to consume—not necessarily just the good being subsidized.

Summary

1. General equilibrium refers to the study of how the economy can adjust to have demand equal supply in all markets at the same time.
2. The Edgeworth box is a graphical tool to examine such a general equilibrium with 2 consumers and 2 goods.
3. A Pareto efficient allocation is one in which there is no feasible reallocation of the goods that would make all consumers at least as well-off and at least one consumer strictly better off.
4. Walras' law states that the value of aggregate excess demand is zero for all prices.
5. A general equilibrium allocation is one in which each agent chooses a most preferred bundle of goods from the set of goods that he or she can afford.
6. Only relative prices are determined in a general equilibrium system.
7. If the demand for each good varies continuously as prices vary, then there will always be some set of prices where demand equals supply in every market; that is, a competitive equilibrium.
8. The First Theorem of Welfare Economics states that a competitive equilibrium is Pareto efficient.
9. The Second Theorem of Welfare Economics states that as long as preferences are convex, then every Pareto efficient allocation can be supported as a competitive equilibrium.

REVIEW QUESTIONS

1. Is it possible to have a Pareto efficient allocation where someone is worse off than he is at an allocation that is not Pareto efficient?
2. Is it possible to have a Pareto efficient allocation where everyone is worse off than they are at an allocation that is not Pareto efficient?
3. True or false? If we know the contract curve, then we know the outcome of any trading.
4. Can some individual be made better off if we are at a Pareto efficient allocation?
5. If the value of excess demand in 8 out of 10 markets is equal to zero, what must be true about the remaining two markets?

APPENDIX

Let us examine the calculus conditions describing Pareto efficient allocations. By definition, a Pareto efficient allocation makes each agent as well-off as possible, given the utility of the other agent. So let us pick $\bar{u}$ as the utility level for agent B, say, and see how we can make agent A as well-off as possible.

The maximization problem is

$$\max_{x_A^1, x_A^2, x_B^1, x_B^2} u_A(x_A^1, x_A^2)$$

$$\text{such that } u_B(x_B^1, x_B^2) = \bar{u}$$

$$x_A^1 + x_B^1 = \omega^1$$

$$x_A^2 + x_B^2 = \omega^2.$$

Here $\omega^1 = \omega_A^1 + \omega_B^1$ is the total amount of good 1 available and $\omega^2 = \omega_A^2 + \omega_B^2$ is the total amount of good 2 available. This maximization problem asks us to find the allocation $(x_A^1, x_A^2, x_B^1, x_B^2)$ that makes person A's utility as large as possible, given a fixed level for person B's utility, and given that the total amount of each good used is equal to the amount available.

We can write the Lagrangian for this problem as

$$L = u_A(x_A^1, x_A^2) - \lambda(u_B(x_B^1, x_B^2) - \bar{u}) \\ - \mu_1(x_A^1 + x_B^1 - \omega^1) - \mu_2(x_A^2 + x_B^2 - \omega^2).$$

Here λ is the Lagrange multiplier on the utility constraint, and the μ 's are the Lagrange multipliers on the resource constraints. When we differentiate with

respect to each of the goods, we have four first-order conditions that must hold at the optimal solution:

$$\begin{aligned}\frac{\partial L}{\partial x_A^1} &= \frac{\partial u_A}{\partial x_A^1} - \mu_1 = 0 \\ \frac{\partial L}{\partial x_A^2} &= \frac{\partial u_A}{\partial x_A^2} - \mu_2 = 0 \\ \frac{\partial L}{\partial x_B^1} &= -\lambda \frac{\partial u_B}{\partial x_B^1} - \mu_1 = 0 \\ \frac{\partial L}{\partial x_B^2} &= -\lambda \frac{\partial u_B}{\partial x_B^2} - \mu_2 = 0.\end{aligned}$$

If we divide the first equation by the second, and the third equation by the fourth, we have

$$MRS_A = \frac{\partial u_A / \partial x_A^1}{\partial u_A / \partial x_A^2} = \frac{\mu_1}{\mu_2} \quad (31.5)$$

$$MRS_B = \frac{\partial u_B / \partial x_B^1}{\partial u_B / \partial x_B^2} = \frac{\mu_1}{\mu_2}. \quad (31.6)$$

The interpretation of these conditions is given in the text: at a Pareto efficient allocation, the marginal rates of substitution between the two goods must be the same. Otherwise, there would be some trade that would make each consumer better off.

Let us recall the conditions that must hold for optimal choice by consumers. If consumer A is maximizing utility subject to her budget constraint and consumer B is maximizing utility subject to his budget constraint, and both consumers face the same prices for goods 1 and 2, we must have

$$\frac{\partial u_A / \partial x_A^1}{\partial u_A / \partial x_A^2} = \frac{p_1}{p_2} \quad (31.7)$$

$$\frac{\partial u_B / \partial x_B^1}{\partial u_B / \partial x_B^2} = \frac{p_1}{p_2}. \quad (31.8)$$

Note the similarity with the efficiency conditions. The Lagrange multipliers in the efficiency conditions, μ_1 and μ_2 , are just like the prices p_1 and p_2 in the consumer choice conditions. In fact the Lagrange multipliers in this kind of problem are sometimes known as **shadow prices** or **efficiency prices**.

Every Pareto efficient allocation has to satisfy conditions like those in equations (31.5) and (31.6). Every competitive equilibrium has to satisfy conditions like those in equations (31.7) and (31.8). The conditions describing Pareto efficiency and the conditions describing individual maximization in a market environment are virtually the same.