

Economic Models: >> Trade-offs and Trade

Section 1: Models in Economics: Some Important Examples

A **model** is a simplified representation of a real situation that is used to better understand real-life situations.

A **model** is any simplified representation of reality that is used to better understand real-life situations. But how do we create a simplified representation of an economic situation?

One possibility—an economist’s equivalent of a wind tunnel—is to find or create a real but simplified economy. For example, economists interested in the economic role of money have studied the system of exchange that developed in World War II prison camps, in which cigarettes became a universally accepted form of payment even among prisoners who didn’t smoke.

Another possibility is to simulate the workings of the economy on a computer. For example, when changes in tax law are proposed, government officials use *tax models*—large computer programs—to assess how the proposed changes would affect different types of people.

The **other things equal assumption** means that all other relevant factors remain unchanged.

The importance of models is that they allow economists to focus on the effects of only one change at a time. That is, they allow us to hold everything else constant and study how one change affects the overall economic outcome. So the **other things equal assumption**, which means that all other relevant factors remain unchanged, is an important assumption when building economic models.

But you can't always find or create a small-scale version of the whole economy, and a computer program is only as good as the data it uses. (Programmers have a saying: garbage in, garbage out.) For many purposes, the most effective form of economic modeling is the construction of "thought experiments": simplified, hypothetical versions of real-life situations.

In Chapter 1 we illustrated the concept of equilibrium with the example of how customers at a supermarket would rearrange themselves when a new cash register opens. Though we didn't say it, this was an example of a simple model—an imaginary supermarket, in which many details were ignored (what are the customers buying? never mind), that could be used to answer a "what if" question: what if another cash register were opened?

As the cash register story showed, it is often possible to describe and analyze a useful economic model in plain English. However, because much of economics involves changes in quantities—in the price of a product, the number of units produced, or the number of workers employed in its production—economists often find that using some mathematics helps clarify an issue. In particular, a numerical example, a simple equation, or—especially—a graph can be key to understanding an economic concept.



Whatever form it takes, a good economic model can be a tremendous aid to understanding. The best way to make this point is to consider some simple but important economic models and what they tell us. First, we will look at the *production possibility frontier*, a model that helps economists think about the trade-offs every economy faces. Then we will turn to *comparative advantage*, a model that

clarifies the principle of gains from trade—trade both between individuals and between countries. Finally, we'll examine the *circular-flow model*, which helps economists analyze the monetary transactions taking place in the economy as a whole.

Note: in discussing these models, we make considerable use of graphs to represent mathematical relationships. Such graphs will play an important role throughout this book. If you are already familiar with the use of graphs, the material that follows should not present any problem. If you are not, this would be a good time to turn to the appendix of this chapter, which provides a brief introduction to the use of graphs in economics.

Trade-offs: The Production Possibility Frontier

The hit movie *Cast Away*, starring Tom Hanks, was an update of the classic story of Robinson Crusoe, the hero of Daniel Defoe's eighteenth-century novel. Hanks played the sole survivor of a plane crash, stranded on a remote island. As in the original story of Robinson Crusoe, the character played by Hanks had limited resources: the natural resources of the island, a few items he managed to salvage from the plane, and, of course, his own time and effort. With only these resources, he had to make a life. In effect, he became a one-man economy.

The first principle of economics we introduced in Chapter 1 was that resources are scarce and that, as a result, any economy—whether it contains one person or millions of people—faces trade-offs. For example, if a castaway devotes resources to catching fish, he cannot use those same resources to gather coconuts.

To think about the trade-offs that face any economy, economists often use the model known as the **production possibility frontier**. The idea behind this model is to improve our understanding of trade-offs by considering a simplified economy that produces only two goods. This simplification enables us to show the trade-off graphically.

The **production possibility frontier** illustrates the trade-offs facing an economy that produces only two goods. It shows the maximum quantity of one good that can be produced for any given production of the other.

Figure 2-1 shows a hypothetical production possibility frontier for Tom, a castaway alone on an island, who must make a trade-off between production of fish and production of coconuts. The frontier—the curve in the diagram—shows the maximum number of fish Tom can catch during a week *given* the quantity of coconuts he gathers, and vice versa. That is, it answers questions of the form, “What is the maximum number of fish Tom can catch if he also gathers 20 (or 25, or 30) coconuts?” (We’ll explain the bowed-out shape of the curve in Figure 2-1 shortly, after we’ve seen how to interpret the production possibility frontier.)

There is a crucial distinction between points *inside or on* the curve (the shaded area) and outside the curve. If a production point lies inside the frontier—like the point labeled C, at which Tom catches 20 fish and gathers 20 coconuts—it is feasible. After all, the frontier tells us that if Tom catches 20 fish, he could also gather a maximum of 25 coconuts, so he could certainly gather 20 coconuts. On the other hand, a production point that lies outside the frontier—such as the hypothetical production point shown in the figure as point D, where Tom catches 40 fish and gathers 30 coconuts—isn’t feasible. (In this case, Tom could catch 40 fish and gather no coconuts *or* he could gather 30 coconuts and catch no fish, but he can’t do both.)

In Figure 2-1 the production possibility frontier intersects the horizontal axis at 40 fish. This means that if Tom devoted all his resources to catching fish, he would catch 40 fish per week but would have no resources left over to gather coconuts. The production possibility frontier intersects the vertical axis at 30 coconuts; this means that if Tom devoted all his resources to gathering coconuts, he could gather 30 coconuts per week but would have no resources left over to catch fish.

The figure also shows less extreme trade-offs. For example, if Tom decides to catch 20 fish, he is able to gather 25 coconuts; this production choice is illustrated by point A. If Tom decides to catch 30 fish, he can gather at most only 20 coconuts, as shown by point B.

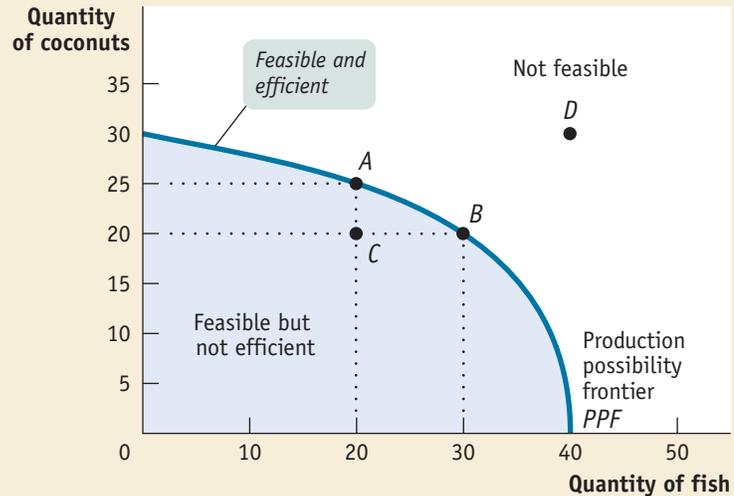
Thinking in terms of a production possibility frontier simplifies the complexities of reality. The real-world economy produces millions of different goods. Even a cast-away on an island would produce more than two different items (for example, he would need clothing and housing as well as food). But in this model we imagine an economy that produces only two goods.

If we simplify reality, however, the production possibility frontier helps us understand some aspects of the real economy better than we could without the model.

Figure 2-1

The Production Possibility Frontier

The production possibility frontier illustrates the trade-offs facing an economy that produces two goods. It shows the maximum quantity of one good that can be produced given the quantity of the other good produced. Here, the maximum number of coconuts that Tom can gather depends on the number of fish he catches, and vice versa. His feasible production is shown by the area *inside or on* the curve. Production at point *C* is feasible but not efficient. Points *A* and *B* are efficient and feasible, but point *D* is not feasible.



First of all, the production possibility frontier is a good way to illustrate the general economic concept of *efficiency*. Recall from Chapter 1 that an economy is efficient if there are no missed opportunities—there is no way to make someone better off without making others worse off. A key element of efficiency is that there are no missed opportunities in production—there is no way to produce more of one good without producing less of other goods.

As long as Tom is on the production possibility frontier, his production is efficient. At point *A*, the 25 coconuts he gathers are the maximum number he can get *given* that he has chosen to catch 20 fish; at point *B*, the 20 coconuts he gathers are the maximum he can get *given* his choice to catch 30 fish; and so on.

But suppose that for some reason Tom was at point *C*, producing 20 fish and 20 coconuts. Then this one-person economy would definitely be *inefficient*: it could be producing more of both goods.

The production possibility frontier is also useful as a reminder of the fundamental point that the true cost of any good is not just the amount of money it costs to buy, but everything else in addition to money that must be given up in order to get that good—the *opportunity cost*. If Tom were to catch 30 fish instead of 20, he would be able to gather only 20 coconuts instead of 25. So the opportunity cost of those 10 extra fish is the 5 coconuts not gathered. And if 10 extra fish have an opportunity cost of 5 coconuts, each 1 fish has an opportunity cost of $\frac{5}{10} = 0.5$ coconuts.

We can now explain the bowed-out shape of the production possibility frontier in Figure 2-1: it reflects an assumption about how opportunity costs change as the mix of output changes. Figure 2-2 shows the same production possibility frontier as Figure 2-1. The arrows in Figure 2-2 illustrate the fact that with this bowed-out production possibility frontier, Tom faces *increasing opportunity cost*: the more fish he catches, the more coconuts he has to give up to catch an additional fish, and vice versa. For example, to go from producing zero fish to producing 20 fish, he has to give up 5 coconuts.

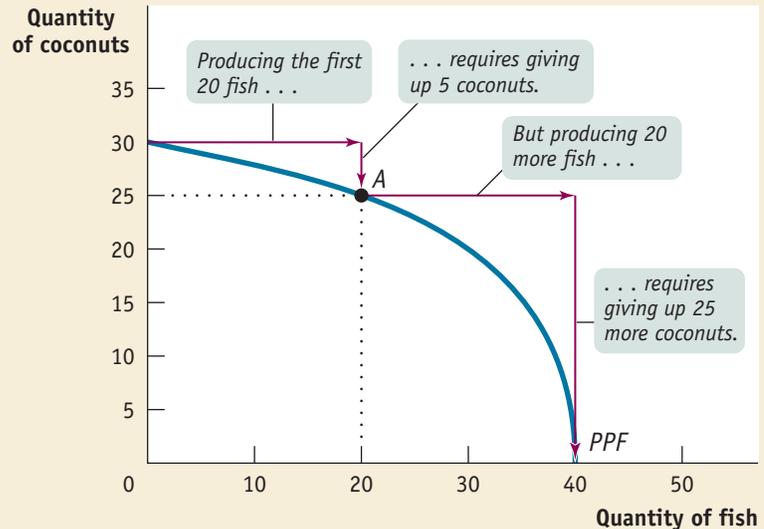
That is, the opportunity cost of those 20 fish is 5 coconuts. But to increase his fish production to 40—that is, to produce an additional 20 fish—he must give up 25 more coconuts, a much higher opportunity cost.

Economists believe that opportunity costs are usually increasing. The reason is that when only a small amount of a good is produced, the economy can use resources that are especially well suited for that production. For example, if an econ-

Figure 2-2

Increasing Opportunity Cost

The bowed-out shape of the production possibility frontier reflects increasing opportunity cost. In this example, to produce the first 20 fish, Tom must give up 5 coconuts. But to produce an additional 20 fish, he must give up 25 more coconuts.



omy grows only a small amount of corn, that corn can be grown in places where the soil and climate are perfect for corn-growing but less suitable for growing anything else, like wheat. So growing that corn involves giving up only a small amount of potential wheat production. If the economy grows a lot of corn, however, land that isn't so great for corn and would have been well suited for wheat must be pressed into service, so the additional corn production will involve sacrificing considerably more wheat production.

Finally, the production possibility frontier helps us understand what it means to talk about *economic growth*. We introduced the concept of economic growth in the Introduction, defining it as *the growing ability of the economy to produce goods and services*. As we saw, economic growth is one of the fundamental features of the real economy. But are we really justified in saying that the economy has grown? After all, although the U.S. economy produces more of many things than it did a century ago, it produces less of other things—for example, horse-drawn carriages. Production of many goods, in other words, is actually down. So how can we say for sure that the economy as a whole has grown?

The answer, illustrated in Figure 2-3, is that economic growth means an *expansion of the economy's production possibilities*: the economy *can* produce more of everything. For example, if Tom's production is initially at point A (20 fish and 25 coconuts), economic growth means that he could move to point E (25 fish and 30 coconuts). E lies outside the original frontier; so in the production possibility frontier model, growth is shown as an outward shift of the frontier.

What the economy actually produces depends on the choices people make. After his production possibilities expand, Tom might not actually choose to produce both more fish and more coconuts—he might choose to increase production of only one good, or he might even choose to produce less of one good. But even if, for some reason, he chooses to produce either fewer coconuts or fewer fish than before, we

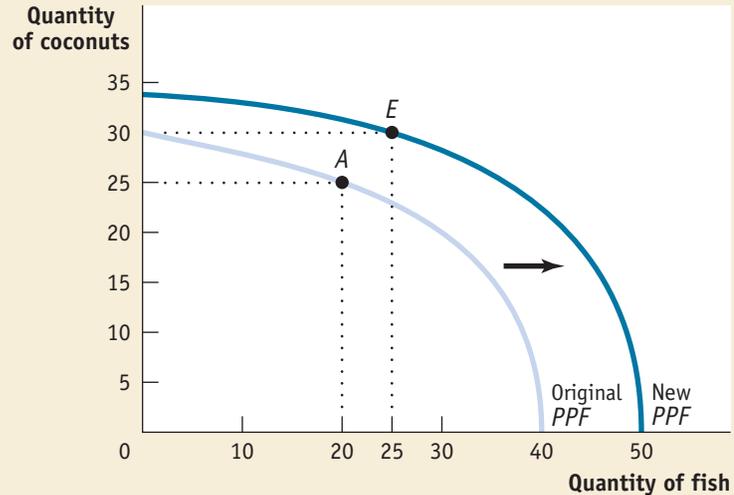
would still say that his economy has grown—because he *could* have produced more of everything.

The production possibility frontier is a very simplified model of an economy. Yet it teaches us important lessons about real-life economies. It gives us our first clear sense of a key element of economic efficiency, it illustrates the concept of opportunity cost, and it makes clear what economic growth is all about.

Figure 2-3

Economic Growth

Economic growth results in an *outward shift* of the production possibility frontier because production possibilities are expanded. The economy can now produce more of everything. For example, if production is initially at point A (20 fish and 25 coconuts), it can move to point E (25 fish and 30 coconuts).



Comparative Advantage and Gains from Trade

Among the nine principles of economics described in Chapter 1 was that of *gains from trade*—the mutual gains that individuals can achieve by specializing in doing different things and trading with one another. Our second illustration of an economic model is a particularly useful model of gains from trade—trade based on *comparative advantage*.

Let's stick with Tom stranded on his island, but now let's suppose that a second castaway, who just happens to be named Hank, is washed ashore. Can they benefit from trading with each other?

It's obvious that there will be potential gains from trade if the two castaways do different things particularly well. For example, if Tom is a skilled fisherman and Hank is very good at climbing trees, clearly it makes sense for Tom to catch fish and Hank to gather coconuts—and for the two men to trade the products of their efforts.

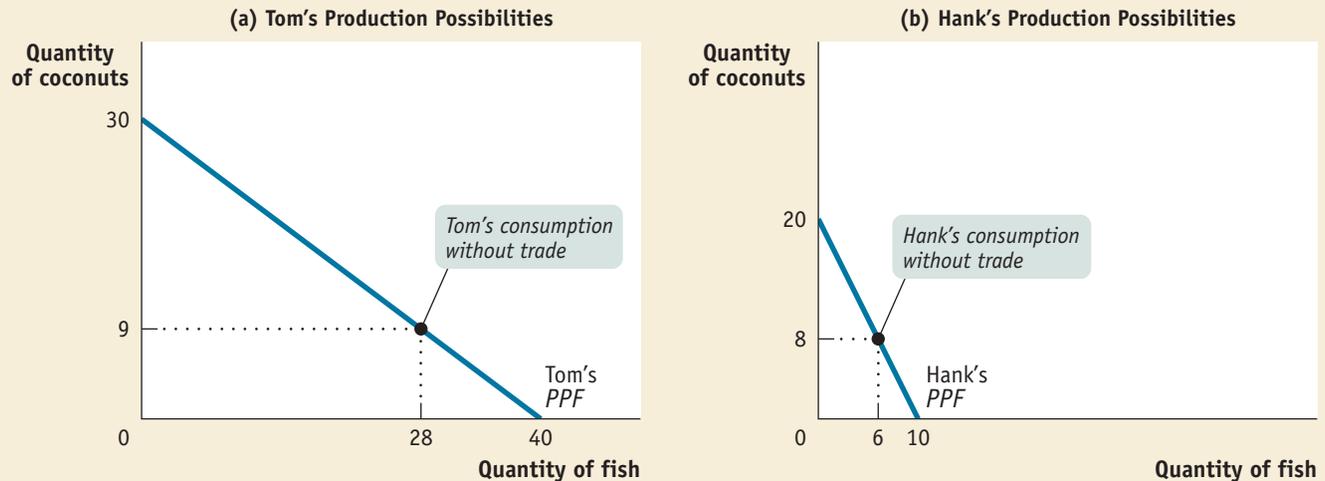
But one of the most important insights in all of economics is that there are gains from trade even if one of the trading parties isn't especially good at anything. Suppose, for example, that Hank is less well suited to primitive life than Tom; he's not nearly as good at catching fish and compared to Tom even his coconut-gathering leaves something to be desired. Nonetheless, what we'll see is that both Tom and Hank can live better by trading with each other than either could alone.

For the purposes of this example, let's slightly redraw Tom's production possibilities represented by the production possibility frontier in panel (a) of Figure 2-4. According to this diagram, Tom could catch at most 40 fish, but only if he gathered no coconuts, and could gather 30 coconuts, but only if he caught no fish, as before.

In Figure 2-4, we have replaced the curved production possibility frontier of Figure 2-1 with a straight line. Why do this, when we've already seen that economists regard a bowed-out production possibility frontier as normal? The answer is that it simplifies our discussion—and as we have explained, modeling is all about simplification. The principle of comparative advantage doesn't depend on the assumption of straight-line production possibility frontiers, but it is easier to explain with that assumption.

The straight-line production possibility frontier in panel (a) of Figure 2-4 has a constant *slope* of $-3/4$. (The appendix to this chapter explains how to calculate the slope of a line.) That is, for every 4 additional fish that Tom chooses to catch, he gathers 3 fewer coconuts. So Tom's opportunity cost of a fish is $3/4$ of a coconut regardless of how many

Figure 2-4 Production Possibilities for Two Castaways



Here, each of the two castaways has a constant opportunity cost of fish and a linear production possibility frontier: In Tom's case, each fish always has an opportunity cost of $3/4$

of a coconut. In Hank's case, each fish always has an opportunity cost of 2 coconuts.

or how few fish he catches. In contrast, a production possibility frontier is curved when the opportunity cost of a good changes according to how much of the good has already been produced. For example, you can see from Figure 2-2 that if Tom starts at the point of having caught zero fish and gathers 30 coconuts, his opportunity cost of catching 20 fish is 5 coconuts. But once he has already caught 20 fish, the opportunity cost of an additional 20 fish increases to 25 coconuts.

Panel (b) of Figure 2-4 shows Hank's production possibilities. Like Tom's, Hank's production possibility frontier is a straight line, implying a constant opportunity cost of fish in terms of coconuts. His production possibility frontier has a constant slope of -2 . Hank is less productive all around: at most he can produce 10 fish or 20 coconuts. But he is particularly bad at fishing; whereas Tom sacrifices $\frac{3}{4}$ of a coconut per fish caught, for Hank the opportunity cost of a fish is 2 whole coconuts. Table 2-1 summarizes the two castaways' opportunity costs for fish and coconuts.

TABLE 2-1**Tom and Hank's Opportunity Costs of Fish and Coconuts**

	Tom's Opportunity Cost	Hank's Opportunity Cost
One fish	$\frac{3}{4}$ coconut	2 coconuts
One coconut	$\frac{4}{3}$ fish	$\frac{1}{2}$ fish

Now Tom and Hank could go their separate ways, each living on his own side of the island, catching his own fish and gathering his own coconuts. Let's suppose that they start out that way and make the consumption choices shown in Figure 2-4: in the absence of trade, Tom consumes 28 fish and 9 coconuts per week, while Hank consumes 6 fish and 8 coconuts.

But is this the best they can do? No, it isn't. Given that the two castaways have different opportunity costs, they can strike a deal that makes both of them better off.

Table 2-2 shows how such a deal works: Tom specializes in the production of fish, catching 40 per week, and gives 10 to Hank. Meanwhile, Hank specializes in the production of coconuts, gathering 20 per week, and gives 10 to Tom. The result is shown in Figure 2-5. Tom now consumes more of both goods than before: instead of 28 fish and 9 coconuts, he consumes 30 fish and 10 coconuts. And Hank also consumes more, going from 6 fish and 8 coconuts to 10 fish and 10 coconuts. As Table 2-2 also shows, both Tom and Hank experience gains from trade: Tom's consumption of fish increases by two, and his consumption of coconuts increases by one. Hank's consumption of fish increases by four, and his consumption of coconuts by two.

So both castaways are better off when they each specialize in what they are good at and trade. It's a good idea for Tom to catch the fish for both of them, because his opportunity cost of a fish in terms of coconuts not gathered is only $\frac{3}{4}$ of a coconut, versus 2 coconuts for Hank. Correspondingly, it's a good idea for Hank to gather coconuts for the both of them.

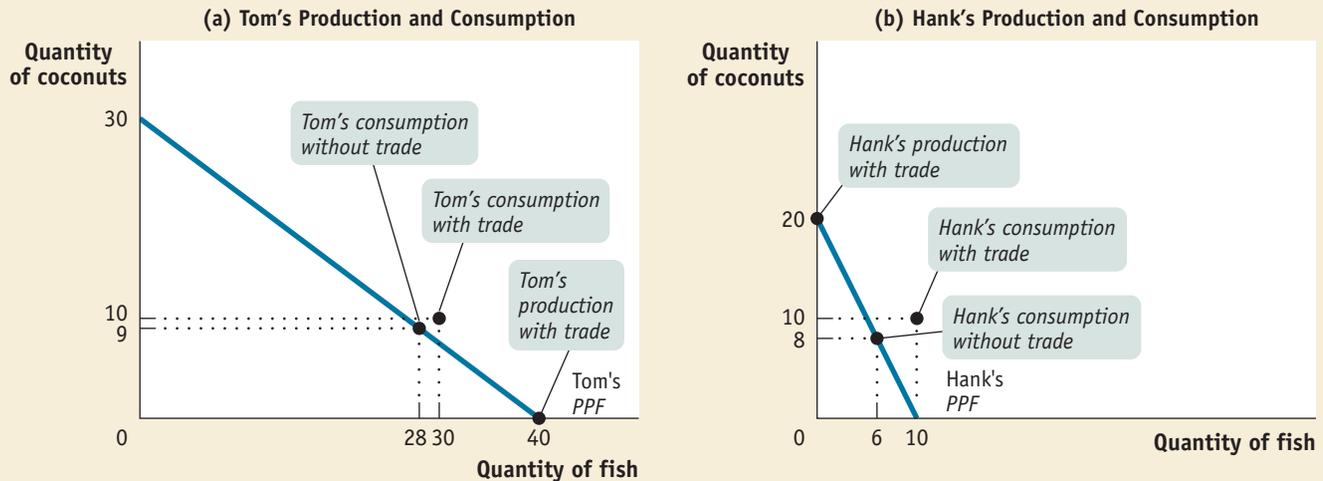
TABLE 2-2

How the Castaways Gain from Trade

		Without Trade		With Trade		Gains from Trade
		Production	Consumption	Production	Consumption	
Tom	Fish	28	28	40	30	+2
	Coconuts	9	9	0	10	+1
Hank	Fish	6	6	0	10	+4
	Coconuts	8	8	20	10	+2

Or we could put it the other way around: Because Tom is so good at catching fish, his opportunity cost of gathering coconuts is high: $\frac{4}{3}$ fish not caught for every coconut gathered. Because Hank is a pretty poor fisherman, his opportunity cost of gathering coconuts is much less, only $\frac{1}{2}$ a fish per coconut.

Figure 2-5 Comparative Advantage and the Gains From Trade



By specializing and trading, the two castaways can produce and consume more of both goods. Tom specializes in catching fish, his comparative advantage, and Hank—who has an *absolute* disadvantage in both goods but a

comparative advantage in coconuts—specializes in gathering coconuts. The result is that each castaway can consume more of both goods than either could without trade.

An individual has a **comparative advantage** in producing a good or service if the opportunity cost of producing the good is lower for that individual than for other people.

An individual has an **absolute advantage** in an activity if he or she can do it better than other people. Having an absolute advantage is not the same thing as having a comparative advantage.

What we would say in this case is that Tom has a **comparative advantage** in catching fish and Hank has a comparative advantage in gathering coconuts. An individual has a comparative advantage in producing something if the opportunity cost of that production is less for that individual than for other people. In other words, Hank has a comparative advantage over Tom in producing a particular good or service if Hank's opportunity cost of producing that good or service is less than Tom's.

The story of Tom and Hank clearly simplifies reality. Yet it teaches us some very important lessons that apply to the real economy, too.

First, the model provides a clear illustration of the gains from trade: by agreeing to specialize and provide goods to each other, Tom and Hank can produce more and therefore both be better off than if they tried to be self-sufficient.

Second, the model demonstrates a very important point that is often overlooked in real-world arguments: as long as people have different opportunity costs, *everyone has a comparative advantage in something, and everyone has a comparative disadvantage in something*.

Notice that in our example Tom is actually better than Hank at producing both goods: Tom can catch more fish in a week, and he can also gather more coconuts. That is, Tom has an **absolute advantage** in both activities: he can produce more output with a given amount of input (in this case, his time) than Hank. You might therefore be tempted to think that Tom has nothing to gain from trading with the less competent Hank.

But we've just seen that Tom can indeed benefit from a deal with Hank because *comparative*, not *absolute*, advantage is the basis for mutual gain. It doesn't matter that it takes Hank more time to gather a coconut; what matters is that for him the opportunity cost of that coconut is lower in terms of fish. So Hank, despite his absolute disadvantage, even in coconuts, has a comparative advantage in coconut-gathering. Meanwhile Tom, who can use his time better by catching fish, has a comparative *disadvantage* in coconut-gathering.

PITFALLS

**MISUNDERSTANDING
COMPARATIVE ADVANTAGE**

Students do it, pundits do it, and politicians do it all the time: they confuse *comparative* advantage with *absolute* advantage. For example, back in the 1980s, when the U.S. economy seemed to be lagging behind that of Japan, one often heard commentators warn that if we didn't improve our productivity, we would soon have no comparative advantage in anything.

What those commentators meant was that we would have no *absolute* advantage in anything—that there might come a time when the Japanese were better at everything than we were. (It didn't turn out that way, but that's another story.) And they had the idea that in that case we would no longer be able to benefit from trade with Japan.

But just as Hank is able to benefit from trade with Tom (and vice versa) despite the fact that Tom is better at everything, nations can still gain from trade even if they are less productive in all industries than the countries they trade with.

If comparative advantage were relevant only to castaways, it might not be that interesting. In fact, however, the idea of comparative advantage applies to many activities in the economy. Perhaps its most important application is to trade—not between individuals, but between countries. So let's look briefly at how the model of comparative advantage helps in understanding both the causes and the effects of international trade.

Comparative Advantage and International Trade

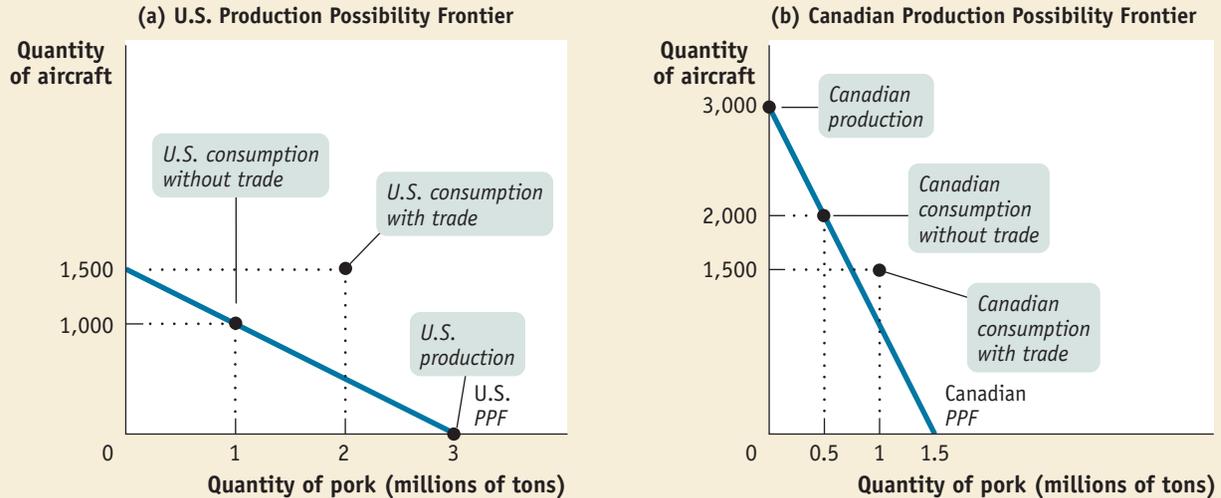
Look at the label on a manufactured good sold in the United States, and there's a good chance you will find that it was produced in some other country—in China, or Japan, or even in Canada, eh? On the other side, many U.S. industries sell a large fraction of their output overseas (this is particularly true of agriculture, high technology, and entertainment).

Should all this international exchange of goods and services be celebrated, or is it cause for concern? Politicians and the public often question the desirability of international trade, arguing that the nation should produce goods for itself rather than buying them from foreigners. Industries around the world demand protection from foreign competition: Japanese farmers want to keep out American rice, American steelworkers want to keep out European steel. And these demands are often supported by public opinion.

Economists, however, have a very positive view of international trade. Why? Because they view it in terms of comparative advantage.

Figure 2-6 shows, with a simple example, how international trade can be interpreted in terms of comparative advantage. Although the example as constructed is hypothetical, it is based on an actual pattern of international trade: American exports of pork to Canada and Canadian exports of aircraft to the United States. Panels (a) and (b) of Figure 2-6 illustrate hypothetical production possibility frontiers for the

Figure 2-6 Comparative Advantage and International Trade



In this hypothetical example, Canada and the United States produce only two goods: pork and aircraft. Aircraft are measured on the vertical axis and tons of pork on the horizontal axis. Panel (a) shows the U.S. production possibility frontier. It is relatively flat, implying that the United States has a comparative

advantage in pork production. Panel (b) shows the Canadian production possibility frontier. It is relatively steep, implying that Canada has a comparative advantage in aircraft production. Just like two individuals, both countries gain from specialization and trade. [>web...](#)

United States and Canada, with tons of pork measured on the horizontal axis and aircraft measured on the vertical axis. The U.S. production possibility frontier is flatter than the Canadian frontier, implying that the United States has a comparative advantage in pork and Canada has a comparative advantage in aircraft.

Although the consumption points in Figure 2-6 are hypothetical, they illustrate a general principle: just like the example of Tom and Hank, the United States and Canada can both achieve mutual gains from trade. If the United States concentrates on producing pork and ships some of its output to Canada, while Canada concentrates on aircraft and ships some of its output to the United States, both countries can consume more than if they insisted on being self-sufficient.

Moreover, these mutual gains don't depend on each country being better at producing one kind of good. Even if one country has, say, higher output per person-hour in both industries—that is, even if one country has an absolute advantage in both industries—there are still mutual gains from trade.

But how does trade actually take place in market interactions? This brings us to our final model, the circular-flow diagram, which helps economists analyze the transactions that take place in a market economy.

Transactions: The Circular-Flow Diagram

The little economy created by Tom and Hank on their island lacks many features of the economy modern Americans live in. For one thing, though millions of Americans are self-employed, most workers are employed by someone else, usually a company with hundreds or thousands of employees. Also, Tom and Hank engage only in the simplest of economic transactions, **barter**, in which an individual directly trades a good or service he or she has for a good or service he or she wants. In the modern economy, simple barter is rare: usually people trade goods or services for money—pieces of colored paper with no inherent value—and then trade those pieces of col-

Trade takes the form of **barter** when people directly exchange goods or services that they have for goods or services that they want.

The **circular-flow diagram** is a model that represents the transactions in an economy by flows around a circle.

A **household** is a person or a group of people that share their income.

A **firm** is an organization that produces goods and services for sale.

Firms sell goods and services that they produce to households in **markets for goods and services**.

Firms buy the resources they need to produce—**factors of production**—in **factor markets**.

ored paper for the goods or services they want. That is, they sell goods or services and buy other goods or services.

And they both sell and buy a lot of different things. The U.S. economy is a vastly complex entity, with more than a hundred million workers employed by hundreds of thousands of companies, producing millions of different goods and services. Yet you can learn some very important things about the economy by considering the simple model shown in Figure 2-7, the **circular-flow diagram**. This diagram represents the transactions that take place in an economy by two kinds of flows around a circle: flows of physical things such as goods, labor, or raw materials in one direction, and flows of money that pay for these physical things in the opposite direction. In this case the physical flows are shown in yellow, the money flows in green.

The simplest circular-flow diagram models an economy that contains only two kinds of “inhabitants”: **households** and **firms**. A household consists of either an individual or a group of people (usually, but not necessarily, a family) that share their income. A firm is an organization (usually, but not necessarily, a corporation) that produces goods and services for sale—and that employs members of households.

As you can see in Figure 2-7, there are two kinds of markets in this model economy. On one side (here the left side) there are **markets for goods and services** in which households buy the goods and services they want from firms. This produces a flow of goods and services to households and a return flow of money to firms.

On the other side, there are **factor markets**. A **factor of production** is a resource used to produce goods and services. Economists usually use the term *factor of production* to refer to a resource that is not used up in production. For example, workers use sewing machines to convert cloth into shirts; the workers and the sewing machines are factors of production, but the cloth is not. Broadly speaking, the main factors of production are labor, land, capital, and human capital. Labor is the work of human beings; land is a resource supplied by nature; capital refers to “created” resources such

as machines and buildings; and human capital refers to the educational achievements and skills of the labor force, which enhance its productivity. Of course, each of these is really a category rather than a single factor: land in North Dakota is quite different from land in Florida.

The factor market most of us know best is the *labor market*, in which workers are paid for their time. Besides labor, we can think of households as owning and selling the other factors of production to firms. For example, when a corporation pays dividends to its stockholders, who are members of households, it is in effect paying them for the use of the machines and buildings that ultimately belong to those investors.

In what sense is Figure 2-7 a model? That is, in what sense is it a *simplified* representation of reality? The answer is that this picture ignores a number of real-world complications. A few examples:

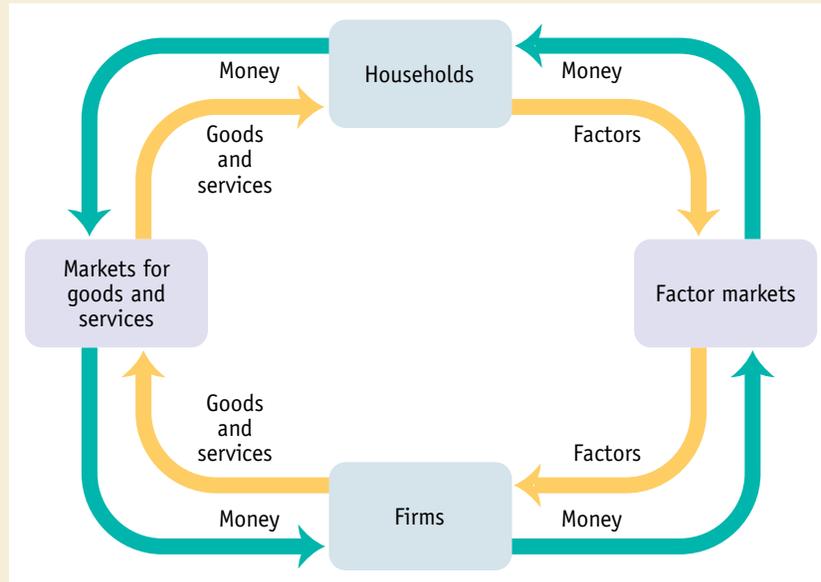
- In the real world, the distinction between firms and households isn't always that clear-cut. Consider a small, family-run business—a farm, a shop, a small hotel. Is this a firm or a household? A more complete picture would include a separate box for family businesses.
- Many of the sales firms make are not to households but to other firms; for example, steel companies sell mainly to other companies such as auto manufacturers, not to households. A more complete picture would include these flows of goods and money within the business sector.
- The figure doesn't show the government, which in the real world diverts quite a lot of money out of the circular flow in the form of taxes but also injects a lot of money back into the flow in the form of spending.

Figure 2-7, in other words, is by no means a complete picture either of all the types of “inhabitants” of the real economy or of all the flows of money and physical items that take place among these inhabitants.

Figure 2-7

The Circular-Flow Diagram

This model represents the flows of money and goods and services in the economy. In the markets for goods and services, households purchase goods and services from firms, generating a flow of money to the firms and a flow of goods and services to the households. The money flows back to households as firms purchase factors of production from the households in factor markets.



Despite its simplicity, the circular-flow diagram, like any good economic model, is a very useful aid to thinking about the economy.

For example, a circular-flow diagram can help us understand how the economy manages to provide jobs for a growing population. To illustrate, consider the huge expansion in the U.S. labor force—the number of people who want to work—between the early 1960s and the late 1980s. This increase was partly caused by the 15-year baby boom that followed World War II; the first baby boomers began looking for jobs

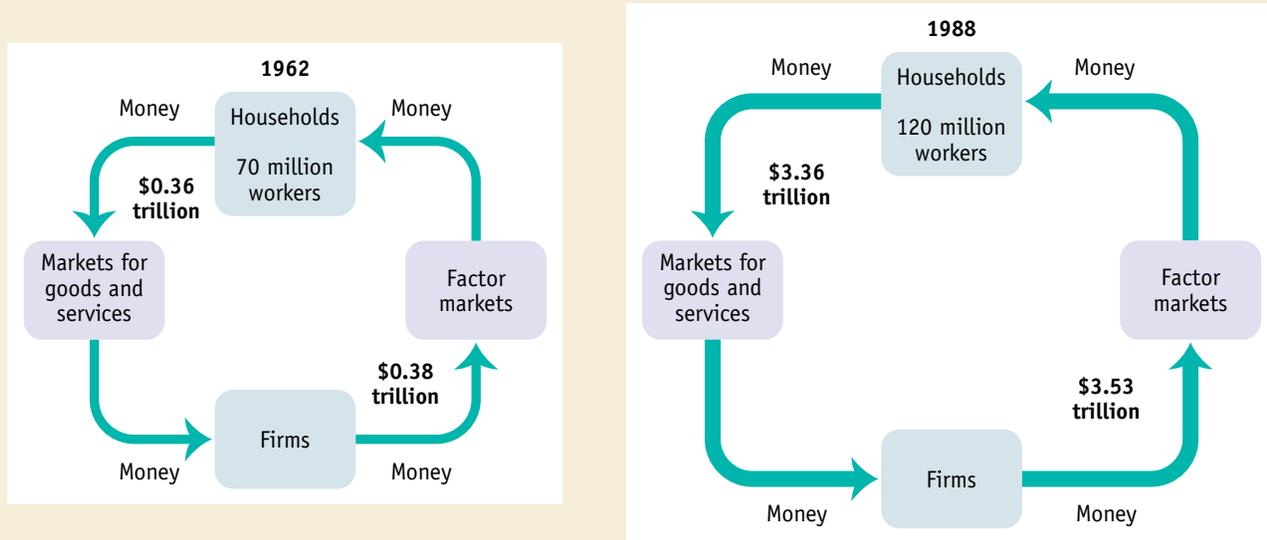
in the early 1960s and the last of them went to work in the late 1980s. In addition, social changes led a much higher fraction of women to seek paid work outside the home. As a result, between 1962 and 1988 the number of Americans employed or seeking jobs increased by 71 percent.

That's a lot of new job seekers. But luckily, the number of jobs also expanded during the same period, by almost exactly the same percentage.

Or was it luck? The circular-flow diagram helps us understand why the number of jobs available grew along with the expansion of the labor force. Figure 2-8 compares the money flows around the circle for the U.S. economy in 1962 and 1988. Both the money paid to households and the money spent by households increased enormously over the period—and that was no accident. As more people went to work—that is, as more labor was sold in the factor markets—households had more income to spend. They used that increased income to buy more goods and services in the market for goods and services. And in order to produce these goods and services, firms had to hire more workers!

So, despite being an extremely simple model of the economy, the circular-flow diagram helps us to understand some important facts about the real U.S. economy. The number of jobs isn't fixed, the model tells us, because it depends on how much households spend; and the amount households spend depends on how many people are working. It is, in other words, no accident that the economy somehow creates enough jobs even when the working population grows rapidly.

Figure 2-8 Growth in the U.S. Economy from 1962 to 1988



These two circular-flow diagrams—one corresponding to 1962, the other corresponding to 1988—help us understand how the U.S. economy was able to produce enough jobs for its rapidly growing labor force. A roughly twofold increase in the number of workers from 1962 to 1988 was accompanied by a ninefold

increase in money flows between households and firms. As the labor force grew, money going to households increased and their spending on goods and services increased. This led firms to hire more workers to meet the increased desire for goods and services and generated more jobs for households.