chapter 7

>> Making Decisions

What you will learn in this chapter:

- How economists model decision making by individuals and firms
- The importance of **implicit** as well as **explicit costs** in decision making
- The difference between **accounting profit** and **economic profit**, and why economic profit is the correct basis for decisions
- The difference between "either–or" and "how much" decisions
- The **principle of marginal analysis**
- What **sunk costs** are and why they should be ignored
- How to make decisions in cases where time is a factor

A TALE OF TWO INVASIONS

On June 6, 1944, Allied soldiers stormed the beaches of Normandy, beginning the liberation of France from German rule. Long before the assault, however, Allied generals had to make a crucial decision: *where* would the soldiers land?

They had to make what we call an "either–or" decision. *Either* the invasion force could cross the English Channel at its narrowest point, Calais—which was what the Germans expected—*or* it could try to surprise the Germans by landing farther west, in Normandy. Since men and landing craft were in limited supply, the Allies could not do both. In fact, they chose to rely on surprise. The German defences in Normandy were too weak to stop the landings, and the Allies went on to liberate France and win the war.

Thirty years earlier, at the beginning of World War I, German generals had to make a different kind of decision. They, too, planned to invade France, in this case via land, and had decided to mount that invasion through Belgium. The decision they had to make was not an either–or but a "how much" decision: *how much* of their army should be allocated to the invasion force, and how much should be used to defend Germany's border with France? The original plan, devised by General Alfred von Schlieffen, allocated most of the German army to the invasion force; on his deathbed, Schlieffen is supposed to have pleaded, "Keep the right wing [the invasion force] strong!" But his successor, General Helmuth von Moltke, weakened the plan: he reallocated some of the divisions that were supposed to race through Belgium to the defence. The weakened invasion force wasn't strong enough: the defending French army stopped it 30 miles from Paris. Most military historians believe that by allocating too few men to the attack, von Moltke cost Germany the war.


MGM/The Kobal Collection

Decision: Attack here? Or there?

So Allied generals made the right decision in 1944; German generals made the wrong decision in 1914. The important point for this chapter is that in both cases the generals had to apply the same logic that applies to economic decisions, like production decisions by businesses and consumption decisions by households.

In this chapter we will survey the principles involved in making economic decisions. These principles will help us understand how any individual—whether a consumer or a producer—makes an economic decision. We begin by taking a deeper look at the significance of opportunity cost for economic decisions and the role it plays in "either–or" decisions. Next we turn to the problem of making "how much" decisions and the usefulness of *marginal analysis*. We then examine what kind of costs should be ignored in making a decision—costs which economists call *sunk costs*. We end by considering the concept of *present value* and its importance for making decisions when costs and benefits arrive at different times.

Opportunity Cost And Decisions

In Chapter 1 we introduced some core principles underlying economic decisions. We've just seen two of those principles at work in our tale of two invasions. The first is that *resources are scarce*—the invading Allies had a limited number of landing craft, and the invading Germans had a limited number of divisions. Because resources are scarce, the true cost of anything is its *opportunity cost*—that is, the real cost of something is what you must give up to get it. When it comes to making decisions, it is crucial to think in terms of opportunity cost, because the opportunity cost of an action is often considerably more than the simple monetary cost.

Explicit versus Implicit Costs

Suppose that, upon graduation from university, you have two options: to go to school for an additional year to get an advanced degree or to take a job immediately. You would like to take the extra year in school but are concerned about the cost.

But what exactly is the cost of that additional year of school? Here is where it is important to remember the concept of opportunity cost: the cost of that year spent getting an advanced degree is what you forgo by not taking a job for that year.

This cost, like any cost, can be broken into two parts: the *explicit cost* of the year's schooling and the *implicit cost*.

An **explicit cost** is a cost that requires an outlay of money. For example, the explicit cost of the additional year of schooling includes tuition. An **implicit cost,** on the other hand, does not involve an outlay of money; instead, it is measured by the value, in dollar terms, of all the benefits that are forgone. For example, the implicit cost of the year spent in school includes the income you would have earned if you had taken that job instead.

An **explicit cost** is a cost that involves actually laying out money. An **implicit cost** does not require an outlay of money; it is measured by the value, in dollar terms, of the benefits that are forgone.

A common mistake, both in economic analysis and in real business situations, is to ignore implicit costs and focus exclusively on explicit costs. But often the implicit cost of an activity is quite substantial—indeed, sometimes it is much larger than the explicit cost.

Table 7-1 gives a breakdown of hypothetical explicit and implicit costs associated with spending an additional year in school instead of taking a job. The explicit cost consists of tuition, books, supplies, and a home computer for doing assignments—all of which require you to spend money. The implicit cost is the salary you would have earned if you had taken a job instead. As you can see, the forgone salary is $35,000 and the explicit cost is $9,500, making the implicit cost more than three times as much as the explicit cost. So ignoring the implicit cost of an action can lead to a seriously misguided decision.

TABLE 7-1

Opportunity Cost of an Additional Year of School

Explicit cost		Implicit cost	
Tuition	$7,000	Forgone salary	$35,000
Books and supplies	1,000		
Home computer	1,500		
Total explicit cost	**9,500**	**Total implicit cost**	**35,000**
Total opportunity cost = Total explicit cost + Total implicit cost = $44,500			

There is another, slightly different way of looking at the implicit cost in this example that can deepen our understanding of opportunity cost. The forgone salary is the cost of using your own resources—your time—in going to school rather than working. The use of your *time* for more schooling, despite the fact that you don't have to spend any money, is nonetheless costly to you. This illustrates an important aspect of opportunity cost: in considering the cost of an activity, you should include the cost of using any of your own resources for that activity. You can calculate the cost of using your own resources by determining what they would have earned in their next best use.

FOR INQUIRING MINDS

FAMOUS COLLEGE DROPOUTS

What do Bill Gates, Tiger Woods, and Sarah Michelle Gellar (a.k.a. Buffy the Vampire Slayer) have in common? None of them have a college degree.

Nobody doubts that all three are easily smart enough to have gotten their diplomas. However, they all made the rational decision that the implicit cost of getting a degree would have been too high—by their late teens, each had a very promising career that would have had to be put on hold to get a college degree. Gellar would have had to postpone her acting career; Woods would have had to put off winning one major tournament after another and becoming the world's best golfer; Gates would have had to delay developing the most successful and most lucrative software ever sold, Microsoft's computer operating system.

In fact, extremely successful people—especially those in careers like acting or athletics, where starting early in life is especially crucial—are often college dropouts. It's a simple matter of economics: the opportunity cost of their time at that stage in their lives is just too high to postpone their careers for a college degree.

Accounting Profit versus Economic Profit

As the example of going to school suggests, taking account of implicit as well as explicit costs can be very important for individuals making decisions. The same is true of businesses.

Consider the case of Kathy's Copy Shop, a small business operating in a local shopping centre. Kathy makes copies for customers, who pay for her services. Out of that revenue, she has to pay her expenses: the cost of supplies and the rent for her store space. We suppose that Kathy owns the copy machines themselves. This year Kathy has $100,000 in revenues and $60,000 in expenses. Is her business profitable?

At first it might seem that the answer is obviously yes: she receives $100,000 from her customers and has expenses of only $60,000. Doesn't this mean that she has a profit of $40,000? Not according to her accountant, who reduces the number by $5,000, for the yearly *depreciation* (reduction in value) of the copy machines.

Depreciation occurs because machines wear out over time. The yearly depreciation amount reflects what an accountant estimates to be the reduction in the value of the machines due to wear and tear that year. This leaves $35,000, which is the business's **accounting profit.** Basically, the accounting profit of a company is its revenue minus its explicit costs and depreciation. The accounting profit is the number that Kathy has to report on her income tax forms and that she would be obliged to report to anyone thinking of investing in her business.

The **accounting profit** of a business is the business's revenue minus the explicit cost and depreciation.

Accounting profit is a very useful number, but suppose that Kathy wants to decide whether to keep her business going or to do something else. To make this decision, she will need to calculate her **economic profit**—the revenue she receives minus her opportunity cost, which may include implicit as well as explicit costs. In general, when economists use the simple term *profit*, they are referring to economic profit. (We will adopt this simplification in later chapters of this book.)

The **economic profit** of a business is the business's revenue minus the opportunity cost of its resources. It is usually less than the accounting profit.

Why does Kathy's economic profit differ from her accounting profit? Because she may have implicit costs over and above the explicit cost her accountant has calculated. Businesses can face implicit costs for two reasons. First, a business's **capital**—its equipment, buildings, tools, inventory, and financial assets—could have been put to use in some other way. If the business owns its capital, it does not pay any money for its use, but it pays an implicit cost because it does not use the capital in some other way. Second, the owner devotes time and energy to the business that could have been used elsewhere—a particularly important factor in small businesses, whose owners tend to put in many long hours.

The **capital** of a business is the value of its assets—equipment, buildings, tools, inventory, and financial assets.

If Kathy had rented her copy machines from the manufacturer, their rent would have been an explicit cost. But because Kathy owns her own machines, she does not pay rent on them and her accountant deducts an estimate of their depreciation in the profit statement. However, this does not account for the opportunity cost of the machines—what Kathy forgoes by owning them. Suppose that instead of using the machines for her own business, the best alternative Kathy has is to sell them for $50,000 and put the money into a bank account where it would earn yearly interest of $3,000. This $3,000 is an implicit cost of running the business.

It is generally known as the **implicit cost of capital,** the opportunity cost of the capital used by a business; it reflects the income that could have been realized if the capital had been used in its next best alternative way. It is just as much a true cost as if Kathy had rented the machines instead of owning them.

The **implicit cost of capital** is the opportunity cost of the capital used by a business—the income the owner could have realized from that capital if it had been used in its next best alternative way.

Finally, Kathy should take into account the opportunity cost of her own time. Suppose that instead of running her own shop, she could earn $34,000 as an office manager. That $34,000 is also an implicit cost of her business.

Table 7-2 summarizes the accounting for Kathy's Copy Shop, taking both explicit and implicit costs into account. It turns out, unfortunately, that although the business makes an accounting profit of $35,000, its economic profit is actually negative.

TABLE 7-2

Profits at Kathy's Copy Shop

Revenue	$100,000
Explicit cost	− 60,000
Depreciation	− 5,000
Accounting profit	**35,000**
Implicit cost of business	
Income Kathy could have earned on capital in the next best way	− 3,000
Income Kathy could have earned as manager	− 34,000
Economic profit	**−2,000**

"I've done the numbers, and I will marry you."

This means that Kathy would be better off financially if she closed the business and devoted her time and capital to something else.

In real life, discrepancies between accounting profits and economic profits are extremely common. As the following Economics in Action explains, this is a message that has found a receptive audience among real-world businesses.

economics in action

Urban Sprawl and the Loss of Farmland in Canada

It seems irrational that some of the most fertile agricultural land in Canada is buried under concrete and tarmac—a victim of urban sprawl. Why does this occur and what, if anything, can be done about it?

The root cause is simple enough. Historically, people congregated in fertile areas, and villages and cities grew up in the middle of the most productive land, especially if there were also water routes nearby that facilitated trading. Given these beginnings, city growth almost inevitably absorbs some of the best farmland.

The mechanics of urban growth work like this. As land prices increase on the edge of an urban area, so the implicit cost of farming increases. It doesn't matter whether the farmer owns the land or not. Because land is a form of capital used to run the business, keeping one's land as a farm instead of selling it to a developer constitutes an implicit cost of capital. Higher land prices increase the implicit cost of capital, which raises the cost of farming—even if the farmer owns the land. This puts intense pressure on farmers to generate incomes that are substantial enough to justify keeping the land in agriculture. Eventually, these pressures become too great and the land is sold for development.

How great are these pressures? Well, farmland in Canada can sell for anything from $100 to $100,000 an acre, depending on its quality and location. Around small urban centres, one would expect to pay about $7000 an acre for good farmland in 2004, depending on the province. But around large metropolitan areas, prices are much higher. For example, a farm within 30 miles of the greater Toronto area, where urban development is foreseeable within the next 5 to 10 years, would command prices of about $100,000 an acre. It's nearly impossible to operate a farm with such a high implicit cost of capital. That's why developers succeed in buying up land where development is foreseen many years before the development takes place. They buy it and hold it as an investment—a so-called "land bank".

Before we get too alarmed about urban sprawl, we should note that urban areas do need space to grow. We should be happy we've got it. Moreover, we should bear in mind that much of the decrease in the amount of land devoted to farming over the last 100 years has nothing to do with the growth of urban centres. Rather, it is due to the replacement of the horse with the tractor as the primary farm vehicle, which reduced the amount of land needed to produce hay and led to the abandonment of much pastureland.

Nevertheless, all provinces feel the need to control urban sprawl in various ways. Most attempt to do this through zoning regulations and urban plans created by the municipality or local service district. The big drawback with this approach is that farmers comprise only about 3% of the rural population, so rural zoning laws may not offer much effective protection against urban development. Partly as a result of this problem, British Columbia set up an Agricultural Land Reserve in the 1970s. In essence, this took zoning decisions out of local hands and put them under provincial jurisdiction, and it has been very effective in preventing urban sprawl around southwestern BC and the Okanagan Valley.

But there are other options for controlling urban sprawl. For example, New Brunswick has a farmland identification program under which provincial property taxes can be deferred indefinitely while the land remains farmland; but should the land be abandoned or developed, the last 15 years' worth of property taxes (plus accrued interest) becomes due immediately. Another method, more common in the U.S. than Canada, is to sell the development rights to a trust. Any developer must then not only buy the land from the farmer but also must buy the development rights from the trust. This insulates the farmer against increases in the implicit cost of capital due to higher land prices caused by impending development. Moreover, farmers benefit from the money they receive from selling the development rights to the land trust, but can meanwhile continue to use the land for agriculture.

So, the main point is two-pronged: first, high implicit costs of capital put enormous pressure on farmers to sell their land to urban developers; and second, attempts to contain this pressure use zoning, farmland identification programs, and land trusts. ■

>>QUICK REVIEW

- All costs are opportunity costs. They can be divided into *explicit costs* and *implicit costs*.
- Companies report their *accounting profit*, which is not necessarily equal to their *economic profit*.
- Due to the *implicit cost of capital*, the opportunity cost of a company's *capital*, and the opportunity cost of the owner's time, economic profit is often substantially less than accounting profit.

>>CHECK YOUR UNDERSTANDING 7-1

Karma and Don run a furniture-refinishing business from their home. Which of the following represent an explicit cost of the business and which represent an implicit cost?

a. Supplies such as paint stripper, varnish, polish, sandpaper, and so on
b. Basement space that has been converted into a workroom
c. Wages paid to a part-time helper
d. A van that they inherited and use only for transporting furniture
e. The job at a larger furniture restorer that Karma gave up in order to run the business

Solutions appear at back of book.

Making "How Much" Decisions: The Role Of Marginal Analysis

As the story of the two wars at the beginning of this chapter demonstrated, there are two types of decisions: "either–or" decisions and "how much" decisions. To help you get a better sense of that distinction, Table 7-3 offers some examples of each kind of decision.

TABLE 7-3

"How Much" versus "Either–Or" Decisions

"How much" decisions	"Either–or" decisions
How many days before you do your laundry?	Tide or Cheer?
How many miles do you go before an oil change in your car?	Buy a car or not?
How many jalapenos on your nachos?	An order of nachos or a sandwich?
How many workers should you hire in your company?	Run your own business or work for someone else?
How much should a patient take of a drug that generates side effects?	Prescribe drug A or drug B for your patients?
How many troops do you allocate to your invasion force?	Invade at Calais or in Normandy?

Although many decisions in economics are "either–or," many others are "how much." Not many people will stop driving if the price of gasoline goes up, but many people will drive less. How much less? A rise in wheat prices won't necessarily persuade a lot of people to take up farming for the first time, but it will persuade farmers who were already growing wheat to plant more. How much more?

To understand "how much" decisions, we use an approach known as *marginal analysis*. Marginal analysis involves comparing the benefit of doing a little bit more of some activity with the cost of doing a little bit more of that activity. The benefit of doing a little bit more of something is what economists call its *marginal benefit*, and the cost of doing a little bit more of something is what they call its *marginal cost*.

Why is this called "marginal" analysis? A margin is an edge; what you do in marginal analysis is push out the edge a bit, and see whether that is a good move.

We will begin our study of marginal analysis by focusing on marginal cost, and we'll do that by considering a hypothetical company called Felix's Lawn-mowing Service, operated by Felix himself with his tractor-mower.

Marginal Cost

Felix is a very hardworking individual; if he works continuously, he can mow 7 lawns in a day. It takes him an hour to mow each lawn. The opportunity cost of an hour of Felix's time is $10.00 because he could make that much in his next best job.

His one and only mower, however, presents a problem when Felix works this hard. Running his mower for longer and longer periods on a given day takes an increasing toll on the engine and ultimately necessitates more—and more costly—maintenance and repairs.

The second column of Table 7-4 shows how the total daily cost of Felix's business depends on the quantity of lawns he mows in a day. For simplicity, we assume that Felix's only costs are the opportunity cost of his time and the cost of upkeep for his mower.

TABLE 7-4

Felix's Marginal Cost of Mowing Lawns

Quantity of lawns mowed	Felix's total cost	Felix's marginal cost of lawn mowed
0	$0	
		$10.50
1	10.50	
		11.25
2	21.75	
		13.25
3	35.00	
		15.50
4	50.50	
		18.00
5	68.50	
		20.75
6	89.25	
		23.75
7	$113.00	

At only 1 lawn per day, Felix's daily cost is $10.50: $10.00 for an hour of his time plus $0.50 for some oil. At 2 lawns per day, his daily cost is $21.75: $20 for 2 hours of his time and $1.75 for mower repair and maintenance. At 3 lawns per day, the daily cost has risen to $35.00: $30.00 for 3 hours of his time and $5.00 for mower repair and maintenance.

The third column of Table 7-4 contains the cost incurred by Felix for each *additional* lawn he mows, calculated from information in the second column. The 1st lawn he mows costs him $10.50; this number appears in the third column between the lines representing 0 lawn and 1 lawn because $10.50 is Felix's cost of going from 0 to 1 lawn mowed. The next lawn, going from 1 to 2, costs him an additional $11.25. So $11.25 appears in the third column between the lines representing the 1st and 2nd lawn, and so on.

The **marginal cost** of an activity is the additional cost incurred by doing one more unit of that activity.

There is **increasing marginal cost** from an activity when each additional unit of the activity costs more than the previous unit.

The increase in Felix's cost when he mows one more lawn is his **marginal cost** of lawn-mowing. In general, the marginal cost of any activity is the additional cost incurred by doing one more unit of that activity.

The marginal costs shown in Table 7-4 have a clear pattern: Felix's marginal cost is greater the more lawns he has already mowed. That is, each time he mows a lawn, the additional cost of doing yet another lawn goes up. Felix's lawn-mowing business has what economists call **increasing marginal cost:** each additional lawn costs

Figure 7-1

The Marginal Cost Curve

The height of each bar is equal to the marginal cost of mowing the corresponding lawn. For example, the 1st lawn mowed has a marginal cost of $10.50, equal to the height of the bar extending from 0 to 1 lawn. The bars ascend in height, reflecting increasing marginal cost: each additional lawn is more costly to mow than the previous one. As a result, the marginal cost curve (drawn by plotting points in the top center of each bar) is upward sloping.

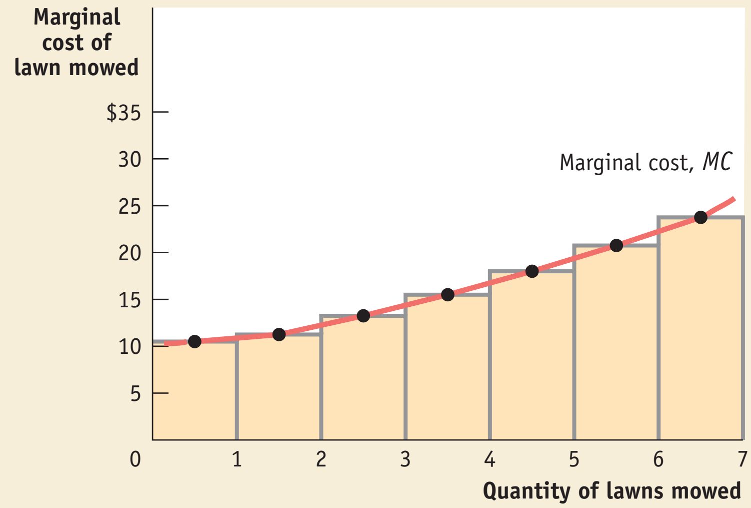

more to mow than the previous one. Or, to put it slightly differently, with increasing marginal cost, the marginal cost of an activity rises as the quantity already done rises.

Figure 7-1 is a graphical representation of the third column in Table 7-4. The horizontal axis measures the quantity of lawns mowed, and the vertical axis measures the marginal cost of a mowed lawn. The height of each shaded bar represents the marginal cost incurred by mowing a given lawn. For example, the bar stretching from 4 to 5 lawns is at a height of $18.00, equal to the cost of mowing the 5th lawn. Notice that the bars form a series of ascending steps, a reflection of the increasing marginal cost of lawn mowing. The **marginal cost curve,** the red curve in Figure 7-1, shows the relationship between marginal cost and the quantity of the activity already done. We draw it by plotting a point in the center at the top of each bar and connecting the points.

The **marginal cost curve** shows how the cost of undertaking one more unit of an activity depends on the quantity of that activity that has already been done.

The marginal cost curve is upward sloping, due to increasing marginal cost. Not all activities have increasing marginal cost; for example, it is possible for marginal cost to be the same regardless of the number of lawns already mowed. Economists call this case *constant* marginal cost. It is also possible for some activities to have a marginal cost that initially falls as we do more of the activity and then eventually rises. These sorts of activities involve gains from specialization: as more output is produced, more workers are hired, allowing each one to specialize in the task that he or she performs best. The gains from specialization yield a lower marginal cost of production.

Now that we have established the concept of marginal cost, we move to the parallel concept of marginal benefit.

Marginal Benefit

Felix's business is in a town where some of the residents are very busy but others are not so busy. For people who are very busy, the opportunity cost of an hour of their time spent mowing the lawn is very high. So they are willing to pay Felix a fairly high sum to do it for them. People with lots of free time, however, have a lower opportunity cost of an hour of their time spent mowing the lawn. So they are willing to pay

PITFALLS

INCREASING TOTAL COST VERSUS INCREASING MARGINAL COST

The concept of *increasing marginal cost* plays an important role in economic analysis, but students sometimes get confused about what it means. That's because it is easy to wrongly conclude that whenever total cost is increasing, marginal cost must also be increasing. But the following example shows that this conclusion is misguided.

Suppose that we change the numbers of our example: the marginal cost of mowing the 6th lawn is now $20, and the marginal cost of mowing the 7th lawn is now $15. In both instances total cost increases as Felix does an additional lawn: it increases by $20 for the 6th lawn and by $15 for the 7th lawn. But in this example marginal cost is *decreasing:* the marginal cost of the 7th lawn is less than the marginal cost of the 6th lawn. So we have a case of increasing total cost and decreasing marginal cost. What this shows us is that, in fact, totals and marginals can sometimes move in opposite directions.

Felix only a relatively small sum. And between these two extremes lie other residents who are moderately busy and so are willing to pay a moderate price to have their lawns mowed.

We'll assume that on any given day, Felix has one potential customer who will pay him $35 to mow her lawn, another who will pay $30, a third who will pay $26, a fourth who will pay $23, and so on. Table 7-5 lists what he can receive from each of his seven potential customers per day, in descending order according to price. So if Felix goes from 0 to 1 lawn mowed, he can earn $35; if he goes from 1 to 2 lawns mowed, he can earn an additional $30; and so on. The third column of Table 7-5 shows us the **marginal benefit** to Felix of each additional lawn mowed. In general, marginal benefit is the additional benefit derived from undertaking one more unit of an activity. Because it arises from doing one more lawn, each marginal benefit value appears between the lines associated with successive quantities of lawns.

The **marginal benefit** from an activity is the additional benefit derived from undertaking one more unit of that activity.

TABLE 7-5

Felix's Marginal Benefit of Mowing Lawns

Quantity of lawns mowed	Felix's total benefit	Felix's marginal benefit of lawn mowed
0	$0	
		$35.00
1	35.00	
		30.00
2	65.00	
		26.00
3	91.00	
		23.00
4	114.00	
		21.00
5	135.00	
		19.00
6	154.00	
		18.00
7	$172.00	

It's clear from Table 7-5 that the more lawns Felix has already mowed, the smaller his marginal benefit from mowing one more. So Felix's lawn-mowing business has what economists call **decreasing marginal benefit:** each additional lawn mowed produces less benefit than the previous lawn. Or, to put it slightly differently, with decreasing marginal benefit, the marginal benefit of an activity falls as the quantity already done rises.

There is **decreasing marginal benefit** from an activity when each additional unit of the activity produces less benefit than the previous unit.

The **marginal benefit curve** shows how the benefit from undertaking one more unit of an activity depends on the quantity of that activity that has already been done.

Just as marginal cost could be represented with a marginal cost curve, marginal benefit can be represented with a **marginal benefit curve,** shown in blue in Figure 7-2. The height of each bar shows the marginal benefit of each additional lawn mowed; the curve through the middle of each bar's top shows how the benefit of each additional unit of the activity depends on the number of units that have already been undertaken.

Felix's marginal benefit curve is downward sloping, because he faces decreasing marginal benefit from lawn-mowing. Not all activities have decreasing marginal benefit; in fact, there are many activities for which marginal benefit is constant—that is, it is the same regardless of the number of units already undertaken. In later chapters where we study firms, we will see that the shape of a firm's marginal benefit curve from producing output has important implications for how it behaves within its industry. We'll also see in Chapters 10 and 11 why economists assume that declining marginal benefit is the norm when considering choices made by consumers. Like increasing marginal cost, decreasing marginal benefit is so common that for now we can take it as the norm.

Now we are ready to see how the concepts of marginal benefit and marginal cost can be brought together to answer the question of "how much" of an activity an individual should undertake.

Figure 7-2

The Marginal Benefit Curve

The height of each bar is equal to the marginal benefit of mowing the corresponding lawn. For example, the 1st lawn mowed has a marginal benefit of $35, equal to the height of the bar extending from 0 to 1 lawn. The bars descend in height, reflecting decreasing marginal benefit: each additional lawn produces a smaller benefit than the previous one. As a result, the marginal benefit curve (drawn by plotting points in the top center of each bar) is downward sloping. >web...

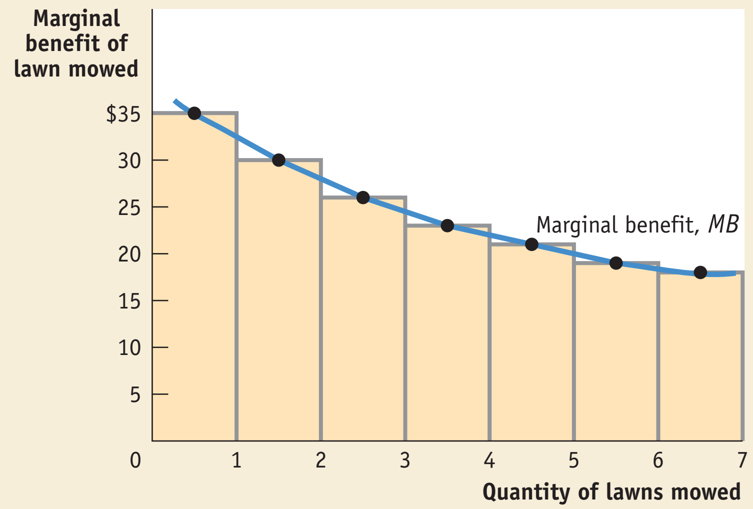

Marginal Analysis

Table 7-6 shows the marginal cost and marginal benefit numbers from Tables 7-4 and 7-5. It also adds an additional column: the net gain to Felix from one more lawn mowed, equal to the difference between the marginal benefit and the marginal cost.

We can use Table 7-6 to determine how many lawns Felix should mow. To see this, imagine for a moment that Felix planned to mow only 3 lawns today. We can immediately see that this is too small a quantity. If Felix mows an additional lawn, increasing the quantity of lawns mowed from 3 to 4, he realizes a marginal benefit of $23.00 and incurs a marginal cost of only $15.50—so his net gain would be $23 – $15.50 = $7.50. But even 4 lawns is still too few: if Felix increases the quantity from 4 to 5, his marginal benefit is $21.00 and his marginal cost is only $18.00, for a net gain of $21.00 – $18.00 = $3.00 (as indicated by the highlighting in the table).

But if Felix goes ahead and mows 7 lawns, that is too many. We can see this by looking at the net gain from mowing that 7th lawn: Felix's marginal benefit is $18.00, but his marginal cost is $23.75. So mowing that 7th lawn would produce a net gain

TABLE 7-6

Felix's Net Gain from Mowing Lawns

Quantity of lawns mowed	Felix's marginal benefit of lawn mowed	Felix's marginal cost of lawn mowed	Felix's net gain of lawn mowed
0			
	$35.00	$10.50	$24.50
1			
	30.00	11.25	18.75
2			
	26.00	13.25	12.75
3			
	23.00	15.50	7.50
4			
	21.00	18.00	3.00
5			
	19.00	20.75	−1.75
6			
	18.00	23.75	−5.75
7			

of $18.00 – $23.75 = –$5.75; that is, a net *loss* for his business. And even 6 lawns is too many: by increasing the quantity of lawns mowed from 5 to 6, Felix incurs a marginal cost of $20.75 compared with a marginal benefit of only $19.00. He is best off at mowing 5 lawns, the largest quantity of lawns at which marginal benefit is at least as great as marginal cost.

The upshot is that Felix should mow 5 lawns—no more and no less. If he mows fewer than 5 lawns, his marginal benefit from one more is greater than his marginal cost; he would be passing up a net gain by not mowing more lawns. If he mows more than 5 lawns, his marginal benefit from the last lawn mowed is less than his marginal cost, resulting in a loss for that lawn. So 5 lawns is the quantity that generates Felix's maximum possible total net gain; it is what economists call the **optimal quantity** of lawns mowed.

The **optimal quantity** of an activity is the quantity that generates the maximum possible total net gain.

Figure 7-3 shows graphically how the optimal quantity can be determined. Felix's marginal benefit and marginal cost curves are both shown. If Felix mows fewer than 5 lawns, the marginal benefit curve is *above* the marginal cost curve, so he can make himself better off by mowing more lawns; if he mows more than 5 lawns, the marginal benefit curve is *below* the marginal cost curve, so he would be better off mowing fewer lawns.

The table in Figure 7-3 confirms our result. The second column repeats information from Table 7-6, showing marginal benefit minus marginal cost—or the net gain—for each lawn. The third column shows total net gain according to the quantity of lawns mowed. The total net gain after doing a given lawn is simply the sum of numbers in the second column up to and including that lawn. For example, the net gain

Figure 7-3 The Optimal Quantity

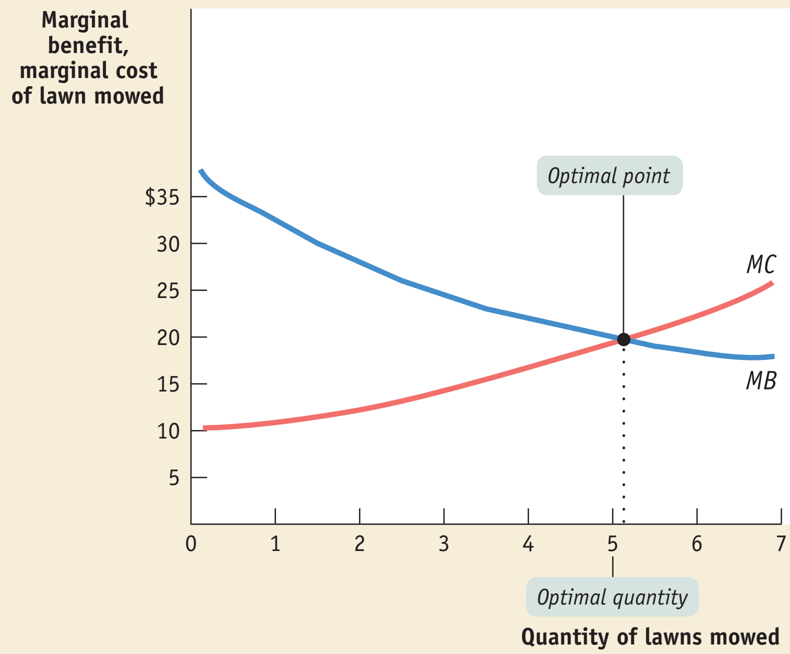

Quantity of lawns mowed	Felix's net gain of lawn mowed	Felix's total net gain
0		$0
	$24.50	
1		24.50
	18.75	
2		43.25
	12.75	
3		56.00
	7.50	
4		63.50
	3.00	
5		66.50
	−1.75	
6		64.75
	−5.75	
7		59.00

The optimal quantity of an activity is the quantity that generates the highest possible total net gain. It is the quantity at which marginal benefit is equal to marginal cost. Equivalently, it is the quantity at which the marginal benefit curve and the marginal cost curve intersect. Here they intersect at approximately 5 lawns. The table beside the graph confirms that 5 is indeed the optimal quantity: the total net gain is maximized at 5 lawns, generating $66.50 in total net gain for Felix.

is $24.50 for the first lawn and $18.75 for the second. So the total net gain after doing the first lawn is $24.50, and the total net gain after doing the second lawn is $24.50 + $18.75 = $43.25. Our conclusion that 5 is the optimal quantity is confirmed by the fact that the greatest total net gain, $66.50, occurs when the 5th lawn is mowed.

The example of Felix's lawn-mowing business shows how you go about finding the optimal quantity: increase the quantity as long as the marginal benefit from one more unit is greater than the marginal cost, but stop before the marginal benefit becomes less than the marginal cost.

In many cases, however, it is possible to state this rule more simply. When a "how much" decision involves relatively large quantities, the rule simplifies to this: the optimal quantity is the quantity at which marginal benefit is equal to marginal cost.

To see why this is so, consider the example of a farmer who finds that her optimal quantity of wheat produced is 5,000 bushels. Typically, she will find that in going from 4,999 to 5,000 bushels, her marginal benefit is only very slightly greater than her marginal cost—that is, the difference between marginal benefit and marginal cost is close to zero. Similarly, in going from 5,000 to 5,001 bushels, her marginal cost is only very slightly greater than her marginal benefit—again, the difference between marginal cost and marginal benefit is very close to zero. So a simple rule for her in choosing the optimal quantity of wheat is to produce the quantity at which the difference between marginal benefit and marginal cost is approximately zero—that is, the quantity at which marginal benefit equals marginal cost.

PITFALLS

MUDDLED AT THE MARGIN

The idea of setting marginal benefit equal to marginal cost sometimes confuses people. Aren't we trying to maximize the *difference* between benefits and costs? And don't we wipe out our gains by setting benefits and costs equal to each other? But what we are doing is setting *marginal,* not *total,* benefit and cost equal to each other.

Once again, the point is to maximize the total net gain from an activity. If the marginal benefit from the activity is greater than the marginal cost, doing a bit more will increase that gain. If the marginal benefit is less than the marginal cost, doing a bit less will increase the total net gain. So only when the *marginal* benefit and cost are equal is the difference between *total* benefit and cost at a maximum.

Economists call this rule the **principle of marginal analysis.** It says that the optimal quantity of an activity is the quantity at which marginal benefit equals marginal cost. Graphically, the optimal quantity is the quantity of an activity at which the marginal benefit curve *intersects* the marginal cost curve. In fact, this graphical method works quite well even when the numbers involved aren't that large. For example, in Figure 7-3 the marginal benefit and marginal cost curves cross each other at about 5 lawns mowed—that is, marginal benefit equals marginal cost at about 5 lawns mowed, which we have already seen is Felix's optimal quantity.

The **principle of marginal analysis** says that the optimal quantity of an activity is the quantity at which marginal benefit is equal to marginal cost.

A Principle with Many Uses

The principle of marginal analysis can be applied to just about any "how much" decision—including those decisions where the benefits and costs are not necessarily expressed in dollars and cents. Here are a few examples:

- The number of traffic deaths can be reduced by spending more on highways, requiring better protection in cars, and so on. But these measures are expensive. So we can talk about the marginal cost to society of eliminating one more traffic fatality. And we can then ask whether the marginal benefit of that life saved is large enough to warrant doing this. (If you think no price is too high to save a life, see the following Economics in Action.)
- Many useful drugs have side effects that depend on the dosage. So we can talk about the marginal cost, in terms of these side effects, of increasing the dosage of a drug. The drug also has a marginal benefit in helping fight the disease. So the optimal quantity of the drug is the quantity that makes the best of this trade-off.
- Studying for an exam has costs because you could have done something else with the time, such as studying for another exam or sleeping. So we can talk about the marginal cost of devoting another hour to studying for your chemistry final. The optimal quantity of studying is the level at which the marginal benefit in terms of a higher grade is just equal to the marginal cost.

economics in action

The Cost of a Life

What's the marginal benefit to society of saving a human life? You might be tempted to answer that human life is infinitely precious. But in the real world, resources are scarce, so we must decide how much to spend on saving lives since we cannot spend infinite amounts. After all, we could surely reduce highway deaths by dropping the speed limit on major highways to 60 kilometres per hour, but the cost of such a lower speed limit—in time and money—is more than anyone is willing to pay.

Generally, people are reluctant to talk in a straightforward way about comparing the marginal cost of a life saved with the marginal benefit—it sounds too callous. Sometimes, however, the question becomes unavoidable.

For example, the cost of saving a life became an object of intense discussion in the United Kingdom in 1999, after a horrifying train crash near London's Paddington Station killed 31 people. There were accusations that the British government was spending too little on rail safety. However, the government estimated that improving rail safety would cost an additional $4.5 million per life saved. But if that amount was worth spending—that is, if the estimated marginal benefit of saving a life exceeded $4.5 million—then the implication was that the British government was spending way too little on *traffic safety*. The estimated marginal cost per life saved through highway improvements was only $1.5 million, making it a much better deal than saving lives through greater rail safety. ■

>>QUICK REVIEW

- A "how much" decision is made by using marginal analysis.
- The *marginal cost* of an activity is represented graphically by the *marginal cost curve*. An upward-sloping marginal cost curve reflects *increasing marginal cost*.
- The *marginal benefit* of an activity is represented by the *marginal benefit curve*. A downward-sloping marginal benefit curve reflects *decreasing marginal benefit*.
- The *optimal quantity* of an activity is found by applying the *principle of marginal analysis*. It says that the optimal quantity of an activity is the quantity at which marginal benefit is equal to marginal cost. Equivalently, it is the quantity at which the marginal cost curve intersects the marginal benefit curve.

>>CHECK YOUR UNDERSTANDING 7-2

1. In each of the "how much" decisions listed in Table 7-3, describe the nature of the marginal cost and of the marginal benefit.
2. Suppose that Felix's marginal cost, instead of increasing, is the same for every lawn he mows.
   a. Assume Felix's marginal cost is $18.50. Using Table 7-6, find the optimal quantity of mowed lawns. What is his total net gain?
   b. How high would marginal cost have to be such that Felix's optimal quantity of lawns mowed is 0? Can you specify a marginal cost for which the optimal quantity is 3?

Solutions appear at back of book.

Sunk Costs

When making decisions, knowing what to ignore is important. Although we have devoted much attention in this chapter to costs that are important to take into account when making a decision, some costs should be ignored when doing so. In this section we will focus on the kinds of costs that people should ignore—what economists call *sunk costs*—and why they should be ignored.

To gain some intuition, consider the following scenario. You own a car that is a few years old, and you have just replaced the brake pads at a cost of $250. But then you find out that the entire brake system is defective and must be replaced—including the newly installed brake pads. This will cost you an additional $1,500. Alternatively, you could sell the car and buy another of comparable quality, but with no brake defects, by spending an additional $1,600. What should you do: fix your old car, or sell it and buy another?

Roy Morsch/Corbis

This vet left law school to pursue his dream career. The cost for a year of law school was lost—a sunk cost. But he and his patients are now happy.

Some might say that you should take the latter option. After all, this line of reasoning goes, if you repair your car you will end up having spent $1,750: $1,500 for the brake system and $250 for the brake pads you replaced. If you were instead to sell your old car and buy another, you would spend only $1,600.

But this reasoning, although it sounds plausible, is wrong. It is wrong because it ignores the fact that you have *already* spent the amount of $250 on brake pads, and that $250 is *non-recoverable*. That is, having been spent already, the $250 cannot be recouped. Therefore, it should be ignored and should have no effect on your decision whether to repair your car and keep it or not. From an economist's viewpoint, the real cost at this time of repairing and keeping your car is $1,500 and not $1,750. Therefore, the correct decision is to repair your car and keep it rather than spend $1,600 on a new car.

In this example, the $250 that has already been spent and cannot be recovered is what economists call a **sunk cost.** Sunk costs should be ignored in making decisions about future actions because they have no influence on their costs and benefits. It's like the old saying, "There's no use crying over spilt milk": once something is gone and can't be recovered, it is irrelevant in making decisions about what to do in the future.

A **sunk cost** is a cost that has already been incurred and is non-recoverable. A sunk cost should be ignored in decisions about future actions.

It is often psychologically hard to ignore sunk costs. And if, in fact, the costs haven't yet been incurred, then they should be taken into consideration. That is, if you had known at the beginning that it would cost $1,750 to repair your car, then the right choice *at that time* would have been to buy a new car for $1,600. But once the $250 had already been paid for brake pads, it is no longer something that should be included in your decision making about your next actions. It may be hard to accept that "bygones are bygones," but it is the right thing to do.

economics in action

The Next Generation

In 2000 and early 2001, several European countries held "spectrum auctions", auctions in which telephone companies bid for portions of a country's airwave space. The telephone companies planned to use this airwave space to offer new mobile phone services to consumers. Companies believed they could earn large profits by providing these new services, so-called third-generation, or 3G, mobile phone services, which included features such as video calling and mobile Internet access. Eager to capture what they expected to be large future profits, telephone companies paid billions of dollars for portions of the European airwave space.

But some technology experts were worried. They believed that the companies had exaggerated expectations of future profits and, as a result, had paid too much for the airwave space. These experts feared that once the companies realized that the airwave space was worth much less than what they had paid, the companies would be unwilling to put up the additional money needed for physical infrastructure, such as the towers used to transmit the signals that are necessary to the 3G services.

It turned out that the technology experts were right about the exaggerated expectations: within a few months of the spectrum auctions, telephone companies realized that they had paid far more for the portions of airwave space than they were really worth.

But was the experts' second conjecture correct: would the overpayment for the airwaves really prevent the future investment needed to provide 3G services? The answer at this point is no. Several companies, including Vodaphone, the British company that owns a substantial part of the American company Verizon, have pushed ahead in building the required infrastructure. As of 2004, 3G was available in over 30 countries worldwide.

Technology experts were wrong about the effect of overpayment because they didn't understand the concept of sunk costs. That is, they didn't understand that once made, those payments for airwave space couldn't be recovered; therefore, they wouldn't affect the telephone companies' willingness to spend additional money to complete the project. After the companies came to the painful—and quite embarrassing—realization of their overpayment, it didn't change the fact

>> QUICK REVIEW

- *Sunk costs*, costs that have already been incurred and that cannot be recovered, should be ignored in decisions regarding future actions. Because they have already been incurred and are unrecoverable, they have no influence on future costs and benefits.

that it was still profitable to build the infrastructure needed to provide the new services. In the end, they appear to have made the right economic calculation—and in the process admitted to themselves that there's no use crying over a lost billion or two. ■

>> CHECK YOUR UNDERSTANDING 7-3

1. You have decided to go into the ice-cream business and have bought a used ice-cream truck for $8,000. Now you are reconsidering. What is your sunk cost in the following scenarios?
   a. The truck cannot be resold.
   b. It can be resold, but only at a 50% discount.
2. You have gone through two years of medical school but are suddenly wondering whether you wouldn't be happier as a musician. Which of the following statements are potentially valid arguments and which are not?
   a. "I can't give up now, after all the time and money I've put in."
   b. "If I had thought about it from the beginning, I never would have gone to med school, so I should give it up now."
   c. "I wasted two years, but never mind—let's start from here."
   d. "My parents would kill me if I stopped now." (*Hint*: we're discussing *your* decision-making ability, not your parents'.)

Solutions appear at back of book.

The Concept of Present Value

In many cases, individuals must make decisions whose consequences extend some ways into the future. For example, when you decide to attend university, you are committing yourself to years of study, which you expect will pay off for the rest of your life. So the decision to attend university is the decision to embark on a long-term project.

As we have already seen, the basic rule in deciding whether or not to undertake a project is that you should compare the benefits of that project with its costs, implicit as well as explicit. But sometimes there can be a problem in making these comparisons: the benefits and costs of a project may not arrive at the same time.

Sometimes the costs of a project come at an earlier date than the benefits. For example, going to university involves large immediate costs: tuition, income forgone because you are in school, and so on. The benefits, such as a higher salary in your future career, come later, often much later.

In other cases, the benefits of a project come at an earlier date than the costs. If you take out a loan to pay for a vacation cruise, the satisfaction of the vacation will come immediately, but the burden of making payments will come later.

But why is time an issue?

Borrowing, Lending, and Interest

In general, having a dollar today is worth more than having a dollar a year from now. To see why, let's consider two examples.

First, suppose that you get a new job that comes with a $1,000 bonus, which will be paid at the end of the first year. But you would like to spend the extra money now—say, on new clothes for work. Can you do that?

The answer is yes—you can borrow money today, and use the bonus to repay the debt a year from now. But if that is your plan, you cannot borrow the full $1,000 today. You must borrow less than that, because a year from now you will have to repay the amount borrowed *plus interest*.

Now consider a different scenario. Suppose that you are paid a bonus of $1,000 today, and you decide that you don't want to spend the money right now. What do you do with it? You put it in the bank; in effect, you are lending the $1,000 to the bank, which in turn lends it out to its customers who wish to borrow. At the end of a year, you will get more than $1,000 back—you will have the $1,000 plus the interest earned.

What all of this means is that $1,000 today is worth more than $1,000 a year from now. The reason is that if you want to have the money today, you must borrow it and pay interest. That is, you must pay a price for using the money today. And, correspondingly, if you forgo using the money today and lend it to someone else, you earn interest on the money. That is, you earn something by letting someone else use your money. When someone borrows money for a year, the **interest rate** is the price, calculated as a percentage of the amount borrowed, charged by the lender.

When someone borrows money for a year, the **interest rate** is the price, calculated as a percentage of the amount borrowed, charged by the lender.

Because of the interest paid on borrowing, you can't evaluate a project just by adding up all the costs and benefits when those costs and benefits arrive at different times. You must take time into account when evaluating the project because a $1 benefit that comes today is worth more than a $1 benefit that comes a year from now; and a $1 cost that comes today is more burdensome to you than a $1 cost that comes next year. Fortunately, there is a simple way to adjust for these complications.

What we will now see is that the interest rate can be used to convert future benefits and costs into what economists call their *present values*. By using present values in evaluating a project, you can evaluate a project *as if* all its costs and benefits were occurring today rather than at different times. This allows people to "factor out" the complications created by time. We'll start by defining exactly what the concept of present value is.

Defining Present Value

The key to the concept of present value is to understand that you can use the interest rate to compare the value of a dollar realized today with the value of a dollar realized later. Why the interest rate? Because the interest rate correctly measures the cost of delaying a dollar of benefit and, correspondingly, the benefit of delaying a dollar of cost. Let's illustrate this with some examples.

Suppose, first, that you are evaluating whether or not to take a job in which your employer promises to pay you a bonus at the end of the first year. What is the value to you today of $1 of bonus money to be paid to you one year in the future? A slightly different way of asking the same question: what would you be willing to accept today in place of receiving $1 one year in the future?

The way to answer this question is to observe that you need *less* than $1 today in order to be assured of having $1 one year from now. Why? Because any money that you have today can be lent out at interest, turning it into a greater sum at the end of the year.

The symbol r is used to represent the rate of interest, expressed as a fraction—that is, if the interest rate is 10%, then $r = 0.10$. If you lend out \$X, at the end of a year you will receive your \$X back, plus the interest on your \$X, which is $\$X \times r$. Thus, at the end of the year you will receive $\$X + \$X \times r$, which is $\$X \times (1 + r)$. What we want to know is how much you would have to lend out today to have $1 a year from now. If the amount you lend out is \$X, it must be true that

(7-1) $$\$X \times (1 + r) = \$1$$

Rearranging, we can solve for \$X, the amount you need today in order to generate $1 one year from now.

(7-2) $$\$X = \$1/(1 + r)$$

This means that you would be willing to accept \$X today for every $1 to be paid one year from now. The reason is that by lending out \$X today, you can be assured of having $1 one year from now. If we plug into the equation the value of the yearly interest rate—say it is 10%, which means that $r = 0.10$—then we can solve for \$X: \$X

The **present value** of \$1 realized one year from now is equal to \$1/(1 + *r*): the amount of money you must lend out today in order to have \$1 in one year. It is the value to you today of \$1 realized one year from now.

is equal to \$1/1.10, which is approximately \$0.91. So you would be willing to accept \$0.91 today in exchange for every \$1 to be paid to you one year from now. Economists have a special name for \$X—it's called the **present value** of \$1.

To see that this technique works for future costs as well as future benefits, consider the following example. Suppose you enter into an agreement that obliges you to pay \$1 one year from now—say to pay off your student loan when you graduate in a year. How much money would you need today to ensure that you have \$1 in a year? The answer is \$X, the present value of \$1, which in our example is \$0.91. The reason \$0.91 is the right answer is that if you lend it out for one year at an interest rate of 10%, you will receive \$1 in return at the end.

What these two examples show us is that the present value concept provides a way to calculate the value today of \$1 that is realized in the future—regardless of whether that \$1 is realized as a benefit (the bonus) or a cost (the student loan payback). This means that to evaluate a project today that has benefits and/or costs to be realized in the future, we just use the relevant interest rate to convert those future dollars into their present values. In that way, we have "factored out" the complication that time creates for decision making.

In the next section we will work out an example of using the present value concept to evaluate a project. But before we do that, it is worthwhile to note that the present value method can be used for projects in which the \$1 is realized more than a year later—say two, three, or even more years.

Suppose you are considering a project that will pay you \$1 *two* years from today. What is the value to you today of \$1 received two years into the future? We can find the answer to that question by expanding our formula for present value.

Let's call \$V the amount of money you need to lend today at an interest rate of r in order to have \$1 in two years. So if you lend \$V today, you will receive $\$V \times (1 + r)$ in one year. And if you *re-lend* that sum for yet another year, you will receive $\$V \times (1 + r) \times (1 + r) = \$V \times (1 + r)^2$ at the end of the second year. At the end of two years, \$V will be worth $\$V \times (1 + r)^2$; if $r = 0.10$, then this becomes $\$V \times (1.10)^2 = \$V \times (1.21)$.

Now we are ready to answer the question of what \$1 realized two years in the future is worth today. In order for the amount lent today, \$V, to be worth \$1 two years from now, it must satisfy this formula:

$$\textbf{(7-3)} \quad \$V \times (1 + r)^2 = \$1$$

Rearranging, we can solve for \$V:

$$\textbf{(7-4)} \quad \$V = \$1/(1 + r)^2$$

Given $r = 0.10$, this means that $\$V = \$1/1.21 = \$0.83$. So when the interest rate is 10%, \$1 realized two years from today is worth \$0.83 today because by lending out \$0.83 today you can be assured of having \$1 in two years. And that means that the present value of \$1 realized two years into the future is \$0.83.

From this example we can see how the present value concept can be expanded to a number of years even greater than two. If we ask what value of \$1 realized N number of years into the future is, the answer is given by a generalization of the present value formula: it is equal to $\$1/(1 + r)^N$.

Using Present Value

Suppose you have to choose one of three projects to undertake. Project A has an immediate payoff to you of \$100, while project B requires that you put up \$10 of your own money today in order to receive \$115 a year from now. Project C gives you an immediate payoff of $119 but requires that you pay $20 a year from now. We'll assume that the annual interest rate is 10%—that is, $r = 0.10$.

The problem in evaluating these three projects is that they have costs and benefits that are realized at different times. That is, of course, where the concept of present value becomes extremely helpful: by using present value to convert any dollars realized in the future into today's value, you factor out the issue of time. This allows you to calculate the **net present value** of a project—the present value of current and future benefits minus the present value of current and future costs. And the best project is the one with the highest net present value.

The **net present value** of a project is the present value of today's and future benefits minus the present value of today's and future costs.

Table 7-7 shows how this is done for each of the three projects. The second and third columns show how many dollars are realized and when they are realized; costs are indicated by a minus sign. The fourth column shows the equations used to convert the flows of dollars into their present value, and the fifth column shows the actual amounts of the total net present value for each of the three projects.

TABLE 7-7

The Net Present Value of Three Projects

Project	Dollars realized today	Dollars realized one year from today	Present value formula	Net present value given $r = 0.10$
A	$100	—	$100	$100.00
B	–$10	$115	$-\$10 + \$115/(1 + r)$	$94.55
C	$119	–$20	$\$119 - \$20/(1 + r)$	$100.82

For instance, to calculate the net present value of project B, we need to calculate the present value of $115 received in one year. The present value of $1 received in one year would be $\$1/(1 + r)$. So the present value of $115 is 115 times $\$1/(1 + r)$; that is, $\$115/(1 + r)$. The net present value of project B is the present value of today's and future benefits minus the present value of today's and future costs: $-\$10 + \$115/(1 + r)$.

From the fifth column, we can immediately see which is the preferred project—it is project C. That's because it has the highest net present value, $100.82, which is higher than the net present value of project A ($100) and much higher than the net present value of project B ($94.55).

This example shows how important the concept of present value is. If we had failed to use the present value calculations and instead had simply added up the dollars generated by each of the three projects, we could have easily been misled into believing that project B was the best project and project C was the worst one.

economics in action

Should You Take the "Set for Life" or the "Cash Out" Option?

Many lotteries in Canada are organised nationally through the Interprovincial Lotteries Corporation. Both Lotto 649 and Super 7 are national lotteries offering large one-time tax-free winnings. Other lotteries are run provincially and geared to specific regional tastes. However, one common game offers the winner a fixed sum over an extended period of time. For example, Loto-Québec has a game called "La Grande Vie" which pays the lucky winner $100,000 each year for life; in Ontario the game is called "Cash for Life" and pays out $2000 a week for life; and in Atlantic Canada and British Columbia, the game is called "Set for Life" and pays out $1000 a week for 25 years. This adds up to a cool $1.3 million ($\$1000 \times 52 \times 25$); hence the name "Set for Life." In all these cases, the lottery corporation in question offers a "cash-out" option.

For example, the Atlantic Lottery Corporation routinely offers a one-time up-front cash payment of $675,000 instead of the weekly payments. Even though this seems like a stingy amount for a quick payoff, most lucky winners have chosen this "cash-out" option instead of the weekly payments. Are they making good decisions? What would you do?

As economics students, we should calculate the present value of the "set for life" income stream. There are two stages to this. First, we suppose we invest each $1000 payment, and calculate the amount we would have after 25 years, including all the interest earned. In the second stage, we take this amount and discount it back to the present to find its present value. While stage one is quite laborious, these calculations are not that difficult with a calculator or computer. But the result depends crucially on the interest rate used.

It turns out that the present value of the "Set for Life" income stream of $1000 per week for 25 years equals the cash-out option of $675,000 at an interest rate of 6½ percent. At higher interest rates, the cash-out option is worth more; at lower interest rates, the income stream is worth more. In April 2004, the yield on long-term government of Canada bonds was only 4.8 percent; therefore, the best decision at that time was to take the income stream. So, why have most winners chosen the cash-out option?

The most likely explanation is they were just impatient to get all the money as quickly as possible. But in that case, they should have taken the income stream and borrowed the $675,000 at 4.8%. Their weekly payments on this loan would be less than their weekly lottery payments. That's called having your cake and eating it too. See how useful a bit of economics can be? ■

>>QUICK REVIEW

- When costs or benefits arrive at different times, you must take the complication created by time into account. This is done by transforming any dollars realized in the future into their *present value.*
- $1 in benefit realized a year from now is worth $\$1/(1 + r)$ today, where r is the *interest rate*. Similarly, $1 in cost realized a year from now is valued at a cost of $\$1/(1 + r)$ today.
- When comparing several projects in which costs and benefits arrive at different times, you should choose the project that generates the highest *net present value.*

>>CHECK YOUR UNDERSTANDING 7-4

1. Consider the three alternative projects shown in Table 7-7. This time, however, suppose that the interest rate is only 2%.
   a. Calculate the net present values of the three projects. Which one is now preferred?
   b. Explain why the preferred choice is different with a 2% interest rate than with a 10% rate.

Solutions appear at back of book.

• A LOOK AHEAD •

This chapter laid out the basic concepts that we need to understand economic decisions. These concepts, as we will soon see, provide the necessary tools for understanding not only the behaviour behind the supply and demand curves, but also the implications of markets for consumer and producer welfare.

But to get there we need a bit more context—we need to know something more about the kinds of decisions that producers and consumers must make. We start with producers: in the next two chapters we will see how marginal analysis determines how much a profit-maximizing producer chooses to produce.

SUMMARY

1. All economic decisions involve the allocation of scarce resources. Some decisions are "either–or" decisions, in which the question is whether or not to do something. Other decisions are "how much" decisions, in which the question is how many resources to put into some use.
2. The cost of using a resource for a particular activity is the opportunity cost of that resource. Some opportunity costs are **explicit costs;** they involve a direct payment of cash. Other opportunity costs, however, are **implicit costs;** they involve no outlay of money but represent the inflows of cash that are forgone. Both explicit and implicit costs should be taken into account in making decisions. Companies use **capital** and their owners' time. So companies should base decisions on **economic profit,** which takes into account implicit costs such as the opportunity cost of the owners' time and the **implicit**

cost of capital. The **accounting profit,** which companies calculate for the purposes of taxes and public reporting, is often considerably larger than the economic profit because it includes only explicit costs and depreciation, not implicit costs.

3. A "how much" decision is made using marginal analysis, which involves comparing the benefit to the cost of doing an additional unit of an activity. The **marginal cost** of an activity is the additional cost incurred by doing one more unit of the activity, and the **marginal benefit** of an activity is the additional benefit gained by doing one more unit. The **marginal cost curve** is the graphical illustration of marginal cost, and the **marginal benefit curve** is the graphical illustration of marginal benefit.
4. Marginal cost and marginal benefit typically depend on how much of the activity has already been done. In the case of **increasing marginal cost,** each additional unit costs more than the unit before; this is represented by an upward-sloping marginal cost curve. In the case of **decreasing marginal benefit,** each additional unit produces a smaller benefit than the unit before; this is represented by a downward-sloping marginal benefit curve.
5. The **optimal quantity** of an activity is the quantity that generates the maximum possible total net gain. According to the **principle of marginal analysis,** the optimal quantity is the quantity at which marginal benefit is equal to marginal cost. It is the quantity at which the marginal cost curve and the marginal benefit curve intersect.
6. A cost that has already been incurred and that is nonrecoverable is a **sunk cost.** Sunk costs should be ignored in decisions about future actions because they have no effect on future benefits and costs.
7. In order to evaluate a project in which costs or benefits are realized in the future, you must first transform them into their **present values** using the **interest rate,** r. The present value of \$1 realized one year from now is $\$1/(1 + r)$, the amount of money you must lend out today to have \$1 one year from now. Once this transformation is done, you should choose the project with the highest **net present value.**

KEY TERMS

Explicit cost, p.??
Implicit cost, p.??
Accounting profit, p.??
Economic profit, p.??
Capital, p.??
Implicit cost of capital, p.??
Marginal cost, p.??
Increasing marginal cost, p.??
Marginal cost curve, p.??
Marginal benefit, p.??
Decreasing marginal benefit, p.??
Marginal benefit curve, p.??
Optimal quantity, p.??
Principle of marginal analysis, p.??
Sunk cost, p.??
Interest rate, p.??
Present value, p.??
Net present value, p.??

PROBLEMS

1. Scott owns and operates a small business that provides economic consulting services. During the year he spends \$55,000 on travel to clients and other expenses, and the computer that he owns depreciates by \$2,000. If he didn't use the computer, he could sell it and earn yearly interest of \$100 on the money created through this sale. Scott's total revenue for the year is \$100,000. Instead of working as a consultant for the year, he could teach economics at a small local college and make a salary of \$50,000.
   a. What is Scott's accounting profit?
   b. What is Scott's economic profit?
   c. Should Scott continue working as a consultant, or should he teach economics instead?
2. Jackie owns and operates a web-design business. Her computing equipment depreciates by \$5,000 per year. She runs the business out of a room in her home. If she didn't use the room as her business office, she could rent it out for \$2,000 per year. Jackie knows that if she didn't run her own business, she could return to her previous job at a large software company that would pay her a salary of \$60,000 per year. Jackie has no other expenses.
   a. How much total revenue does Jackie need to make in order to break even in the eyes of her accountant? That is, how much total revenue would give Jackie just zero accounting profit?
   b. How much total revenue does Jackie need to make in order for her to want to remain self-employed? That is, how much total revenue would give Jackie just zero economic profit?
3. You own and operate a bike store. Each year, you receive revenue of \$200,000 from your bike sales, while it costs you \$100,000 to obtain the bikes. In addition, you pay \$20,000 for electricity, taxes, and other expenses per year. Instead of running the bike store, you could become an accountant and receive a yearly salary of \$40,000. A large clothing retail chain wants to expand and offers to rent the store from you for \$50,000 per year. How do you explain to your friends that despite making a profit, it is too costly for you to continue running your store?

4. Suppose you have just paid your nonrefundable fees of $1,000 for your meal plan for this academic term. This allows you to eat dinner in the cafeteria every evening.

a. You are offered a part-time job in a restaurant where you can eat for free each evening. Your parents say that you should eat dinner in the cafeteria anyway, since you have already paid for those meals. Are your parents right? Explain why or why not.

b. Now suppose that you are offered a part-time job in a restaurant, but rather than being able to eat there for free, the restaurant only gives you a large discount on your meals there. Each meal at the restaurant will cost you $2, and if you eat there each evening this semester it will add up to $200. Your roommate says that you should eat in the restaurant since it costs less than the $1,000 that you paid for the meal plan. Is your roommate right? Explain why or why not.

5. You have already bought a $10 ticket for the college soccer game in advance. The ticket cannot be resold. You know that going to the soccer game will give you a benefit equal to $20. After you have bought the ticket to the soccer game, you hear that there will be a professional baseball post-season game at the same time. Tickets to the baseball game cost $20, and you know that going to the baseball game will give you a benefit equal to $35. You tell your friends the following: “If I had known about the baseball game before buying the ticket to the soccer game, I would have gone to the baseball game instead. But now that I have the ticket to the soccer game already, it’s better for me to just go to the soccer game.” Are you making the correct decision? Justify your answer by calculating the benefits and costs of your decision.

6. Amy, Bill, and Carla all mow lawns for money. Each of them operates a different lawnmower. The accompanying table shows the total cost to Amy, Bill, and Carla of mowing lawns.

Quantity of lawns mowed	Amy’s total cost	Bill’s total cost	Carla’s total cost
0	$0	$0	$0
1	20	10	2
2	35	20	7
3	45	30	17
4	50	40	32
5	52	50	52
6	53	60	82

a. Calculate Amy’s, Bill’s, and Carla’s marginal costs, and draw each of their marginal cost curves.

b. Who has increasing marginal cost, who has decreasing marginal cost, and who has constant marginal cost?

7. You are the manager of a gym, and you have to decide how many customers to admit each hour. Assume that each customer stays exactly one hour. Customers are costly to admit because they inflict wear and tear on the exercise equipment. Moreover, each additional customer generates more wear and tear than the customer before. As a result, the gym faces increasing marginal cost. The table below shows the marginal costs associated with each number of customers per hour.

Quantity of customers per hour	Marginal cost of customer
0	
	$14.00
1	
	14.50
2	
	15.00
3	
	15.50
4	
	16.00
5	
	16.50
6	
	17.00
7	

a. Suppose that each customer pays $15.25 for a one-hour workout. Use the principle of marginal analysis to find the optimal number of customers that you should admit per hour.

b. You increase the price of a one-hour workout to $16.25. What is the optimal number of customers per hour that you should admit now?

8. Georgia and Lauren are economics students who go to a karate class together. Both have to choose how many classes to go to per week. Each class costs $20. The table below shows Georgia’s and Lauren’s estimates of the marginal benefit that each of them gets from each class per week.

Quantity of classes	Lauren’s marginal benefit of each class	Georgia’s marginal benefit of each class
0		
	$23	$28
1		
	19	22
2		
	14	15
3		
	8	7
4		

a. Use marginal analysis to find Lauren’s optimal number of karate classes per week

b. Will Georgia be willing to go to the same number of classes per week that are optimal for Lauren?

9. Recently, the Atlanta-based Center for Disease Control and Prevention (CDC) recommended against vaccination of the whole population of the United States against the smallpox virus because the vaccination has undesirable, and sometimes fatal, side effects. Suppose the accompanying table gives the available data about the effects of a smallpox vaccination program.

Percent of population vaccinated	Deaths due to smallpox	Deaths due to vaccination side effects
0	200	0
10	180	4
20	160	10
30	140	18
40	120	33
50	100	50
60	80	74

a. Calculate the marginal benefit (in terms of lives saved) and the marginal cost (in terms of lives lost) of each 10% increment of smallpox vaccination. Calculate the net gain of a 10% increment in population vaccinated.

b. Using marginal analysis, decide what percentage of the population should optimally be vaccinated.

10. Patty delivers pizza using her own car, and she is paid according to how many pizzas she delivers. The accompanying table shows Patty's total benefit and total cost when she works a specific number of hours.

Quantity of hours worked	Total benefit	Total cost
0	$0	$0
1	30	10
2	55	21
3	75	34
4	90	50
5	100	70

a. Use marginal analysis to decide how many hours Patty should work. In other words, what is the optimal number of hours Patty should work?

b. Calculate the total net gain to Patty from working 0 hours, 1 hour, 2 hours, and so on. Now suppose Patty chooses to work for one hour. Compare her total net gain from working for one hour with the total net gain from working the optimal number of hours. How much would she lose from working for only 1 hour?

11. Assume De Beers is the sole producer of diamonds. When it wants to sell more diamonds, it must lower its price in order to induce consumers to buy more. Furthermore, each additional diamond that is produced costs more than the previous one due to the difficulty of mining for diamonds. De Beers's total benefit schedule is given in the accompanying table, along with its total cost schedule.

Quantity of diamonds	Total benefit	Total cost
0	$0	$0
1	1,000	50
2	1,900	100
3	2,700	200
4	3,400	400
5	4,000	800
6	4,500	1,500
7	4,900	2,500
8	5,200	3,800

a. Draw the marginal cost curve and the marginal benefit curve and, from your diagram, graphically derive the optimal quantity of diamonds to produce.

b. Calculate the total net gain to De Beers from producing each quantity of diamonds. Which quantity gives De Beers the highest total net gain?

12. You have won the provincial lottery. There are two ways in which you can receive your prize: You can either have $1 million in cash now, or you can have $1.2 million that is paid out as follows: you get $300,000 now, $300,000 in one year's time, $300,000 in two years' time, and $300,000 in three years' time. The interest rate is 20%. How would you prefer to receive your prize?

13. The drug company Pfizer is considering whether to invest in the development of a new cancer drug. Development will require an initial investment of $10 million now; beginning one year from now, the new drug will generate annual profits of $4 million for three years.

a. If the interest rate is 12% should Pfizer invest in the development of the new drug? Why or why not?

b. If the interest rate is 8% should Pfizer invest in the development of the new drug? Why or why not?

>web... To continue your study and review of concepts in this chapter, please visit the Krugman/Wells website for quizzes, animated graph tutorials, web links to helpful resources, and more.

www.worthpublishers.com/krugmanwells