

Stand-Up Economics: The Micro Textbook

Version 5.00

September 2009

Available with and without calculus at

<http://www.standupeconomist.com/books>

Email contact: yoram@standupeconomist.com



This work is licensed under the Creative Commons Attribution-NonCommercial 3.0 Unported License. To view a copy of this license, visit Creative Commons online at <http://creativecommons.org/licenses/by-nc/3.0/>.

Contents

README.TXT	vii
I The optimizing individual	11
1 Introduction	13
2 Decision trees	17
2.1 Decision trees	17
2.2 Example: Monopoly	18
3 Time	23
3.1 Lump sums	24
3.2 Annuities and perpetuities	25
3.3 Capital theory	27
3.4 <i>Algebra</i> : Budget constraints	30
3.5 <i>Algebra</i> : Continuous compounding	31
4 Risk	43
4.1 Reducing risk with diversification	44
4.2 <i>Algebra</i> : Indifference curves	47
5 From one to some	55
5.1 No arbitrage	55
5.2 Rush-hour arbitrage	56
5.3 Financial arbitrage	57
II Strategic interactions	63
6 Cake cutting	65
6.1 Some applications of game theory	66
6.2 Cake-cutting: The problem of fair division	67
6.3 The importance of trade	70

7	Pareto efficiency	75
7.1	Cost-benefit analysis	76
7.2	Pareto	77
7.3	Example: Taxes	78
8	Simultaneous-move games	83
8.1	Strictly dominant strategies	84
8.2	The Prisoners' Dilemma	84
9	Auctions	91
9.1	Kinds of auctions	92
9.2	Bid-shading and truth-revelation	93
9.3	Auction equivalences	94
9.4	Auction miscellany	96
10	From some to many	105
10.1	Monopolies in the long run	106
10.2	Barriers to entry	107
III	Market interactions	109
11	Supply and demand	111
11.1	The story of supply and demand	112
11.2	Shifts in supply and demand	113
11.3	<i>Algebra</i> : The algebra of markets	115
12	Taxes	121
12.1	Per-unit taxes on the sellers	121
12.2	Per-unit taxes on the buyers	123
12.3	Tax equivalence	125
12.4	Tax incidence	125
12.5	<i>Algebra</i> : The algebra of taxes	126
13	Margins	137
13.1	Reinterpreting the supply curve	137
13.2	Reinterpreting the demand curve	139
13.3	Conclusion: Carrots and sticks	141
14	Elasticity	143
14.1	The price elasticity of demand	143
14.2	Beyond the price elasticity of demand	145
14.3	Applications	147
15	Transition: Welfare economics	151
15.1	From theory to reality	153

IV Supplemental chapters	155
16 Inflation	157
16.1 Nominal and real interest rates	157
16.2 Inflation	158
16.3 Mathematics	160
17 Fisheries	167
17.1 An economic perspective	168
17.2 A brief look at government intervention	168
17.3 ITQs to the rescue?	171
18 Sequential-move games	173
18.1 Backward induction	174
19 Iterated dominance and Nash equilibrium	189
19.1 Iterated dominance	189
19.2 Nash equilibrium	192
19.3 Infinitely repeated games	194
19.4 Mixed strategies	196
20 Supply and demand details	203
20.1 Deconstructing supply and demand	203
20.2 <i>Math</i> : The algebra of markets	204
20.3 On the shape of the demand curve	206
20.4 On the shape of the supply curve	207
20.5 Comparing supply and demand	208
Glossary	213
Index	221

README.TXT

Jay Leno says that anybody can become a successful stand-up comic if they give it seven years. . . and the same thing is true about getting an undergraduate degree in economics!

This textbook is a companion to *The Cartoon Introduction to Economics: Volume One, Microeconomics*, by Grady Klein and Yoram Bauman, which you can buy from amazon.com. The first three parts of this textbook cover the same material as the three parts of the cartoon book; the fourth part covers supplemental material not covered in the cartoon book. Both books try to apply three lessons from stand-up comedy to the world of economics:

Make it short Most of the work of stand-up comedy involves boiling down 10 minutes of material that has promise into 2 minutes of material that kills. This book is short, and the *Cartoon Introduction* is even shorter.

Make it funny Look, this book is not a laugh riot—the *Cartoon Introduction* is funnier, and the videos at <http://www.standupeconomist.com> are *much* funnier—but it tries.

Tell a story Stand-up comedy is less about one-liners and more about stories.

The story in this book is about the big question in microeconomics: “Under what circumstances does individual self-interest lead to good outcomes for the group as a whole?” In looking at this question, we’ll follow a path that is not unlike a hiking trip. We start out by putting our boots on and getting our gear together: in Part I we study the optimizing individual. Then we set out on our path and immediately find ourselves hacking through some pretty thick jungle: even simple interactions between just two people (Part II) can be very complicated! As we add even more people (in studying auctions, for example), things get even more complicated, and the jungle gets even thicker. Then a miracle occurs: we add even more people, and a complex situation suddenly becomes simple. After hacking through thick jungle, we find ourselves in a beautiful clearing: competitive markets (Part III) are remarkably easy to analyze and understand.

Part I

The optimizing individual

Chapter 1

Introduction

Two monks walking through a garden stroll onto a small bridge over a goldfish pond and stop to lean their elbows on the railing and look contemplatively down at the fish. One monk turns to the other and says, “I wish I were a fish; they are so happy and content.” The second monk scoffs: “How do you know fish are happy? You’re not a fish!” The reply of the first monk: “Ah, but how do you know what I know, since you are not me?” (Adapted from a story told by the Chinese philosopher Zhuangzi in the 4th century BCE.)

Economics is a social science, and as such tries to explain human behavior. Different disciplines—psychology, sociology, political science, anthropology—take different approaches, each with their own strengths and weaknesses; but all try to shed some light on human behavior and (as the joke suggests) all have to make assumptions about how people work. The basic assumption of economics is that **decisions are made by optimizing individuals**.

Decisions

Economics studies the act and implications of choosing. Without choice, there is nothing to study. As Mancur Olson put it in *The Logic of Collective Action*: “To say a situation is ‘lost’ or hopeless is in one sense equivalent to saying it is perfect, for in both cases efforts at improvement can bring no positive results.”

Individuals

Economics assumes that the power to make choices resides in the hands of individuals. The approach that economists take in studying the behavior of groups of individuals (consumers, businesses, countries, etc.) is to study the incentives and behaviors of each individual in that group.

The key question in economics is whether—and under what circumstances—*individual* decision-making leads to results that are good *for the group as a*

whole. For a pessimistic view on this issue, see the Prisoner’s Dilemma game in Chapter 8. For an optimistic view, see Chapter 15 or consider the words of Adam Smith, who wrote in 1776 that “[man is] led by an invisible hand to promote an end which was no part of his intention. . . . By pursuing his own interest he frequently promotes that of society more effectually than when he really intends to promote it.”

Optimization

Economics assumes that individuals try to do the best they can. Although economics is unwavering in the assumption that individuals are optimizing—i.e., that each has some objective—there *is* flexibility in determining exactly what those objectives are. In particular, economics does not need to assume that individuals are selfish or greedy; their objectives may well involve friends or family, or people they’ve never met, or even plants and animals.

Economics also does not make value judgments about different types of individuals; for example, economists do not say that people who avoid risk are better or worse than people who seek out risk. We simply note that, given identical choices, different people act in different ways. In all cases, we assume that each individual is making the decision that is in their best interest.

An aside about individual firms Economists often treat companies as optimizing individuals with a goal of *profit maximization*. (For our purposes, **profit** is simply money in minus money out.¹) Although the assumption of profit maximization is useful, it does have some problems. One is that some firms—such as food co-operatives—have different goals. A deeper problem is that it is not entirely correct to attribute any goals whatsoever to firms because firms are not optimizing individuals. Rather, a firm is a collection of individuals—workers, managers, stockholders—that is unlikely to function seamlessly as a cohesive optimizing unit because the various individuals have their own objectives. While some branches of economics do analyze the goings-on inside firms, in many cases it is valuable to simplify matters by assuming—as we will throughout this book—that firms act like optimizing individuals and that their objective is to maximize profit.

An aside about individual people Economists say that optimizing individuals pursue *utility maximization*. Although you can think of **utility** as simply a fancy word for happiness, some complex philosophical and mathematical ideas lie under the surface. At the heart of modern economic theory is the idea of *preferences*: given two options, an individual can say that she prefers one over the other (or that she likes them both equally). If we make certain assumptions—e.g., that if the individual prefers A over B and prefers B over C then she prefers A over C—then we can represent individual preferences using a **utility function** that assigns a number (a “happiness level”) to each option.

¹The ideas in Chapter 3 can be used to refine this definition to account for the changing value of money over time. Note also that there are many different definitions of profit, including “accounting profit” and “economic profit”.

Economists in the late 1900s thought that utility might actually be real, something that could be measured using “hedonometers” or “psychogalvanometers”. In contrast, modern economic theory treats utility as simply a handy mathematical technique for representing an individual’s preferences: saying that option A produces a utility of 20 and option B produces a utility of 10 means no more and no less than saying that the individual prefers A to B.²

Problems

Answers are in the next section. The first few words of each question (and each answer) are hyperlinked to allow online users to quickly move back and forth.

- 1.1 [What are some](#) of the advantages and disadvantages of assuming that decisions are made by optimizing individuals? What does it say about, say, programs like Social Security that are intended to help people save for their retirement?

Answers

- 1.1 [One of the strengths](#) of the assumption that decisions are made by optimizing individuals is that it often provides clear predictions of human behavior that we can evaluate in the real world; this allows us to think about under what circumstances our assumption may or may not be correct or useful. One of the disadvantages is that it is not always true that humans act like optimizing individuals, so in cases where “irrational” behavior is important our assumption will lead us in the wrong direction.

Both the strengths and the weaknesses of our assumption are evident when we look at a program like Social Security. If individuals are optimizing, then they will save for their retirement on their own, which would suggest that we don’t need a government program like Social Security anymore than we need a government program to (say) help people save to buy a new car. At the same time, many people would be hesitant to abandon Social Security and go back to a system where people are on their own to save for their retirement; presumably this is because we don’t have confidence that people really are optimizing individuals who will save for their retirement. (PS. Somewhere in the middle ground of this argument is the idea of “privatizing” Social Security, i.e., replacing the government system we have now with a system of “forced savings” where each individual has

²For details, see Jean-Jacques Laffont, *The Economics of Uncertainty and Information* (MIT Press, 1995), or Mas-Colell, Whinston, and Green, *Microeconomic Theory* (Oxford Univ. Press, 1995), which is basically the *Joy of Cooking* of microeconomics.

an account and is required to save a portion of their income in their account. Note, however, that even this type of “forced savings” program is unnecessary if people really are optimizing individuals who will save on their own.)

Chapter 2

Decision trees

2.1 Decision trees

We can imagine that optimizing individuals make decisions by listing all the options and then choosing the best one. We can visualize this process with the help of a **decision tree** (Figure 2.1). The individual starts at the left-most node on the tree, chooses between the various options, and continues moving to the right until he reaches one of the payoff boxes at the end of the tree.

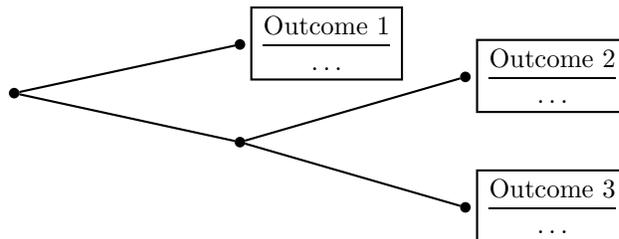


Figure 2.1: A simple decision tree

Comparing the items in the various payoff boxes, we can see that some are in all the payoff boxes and some are not. Items that are in all the payoff boxes are called **sunk costs**. For example, say you pay \$20 to enter an all-you-can-eat restaurant. Once you enter and begin making choices about what to eat, the \$20 you paid to get into the restaurant becomes a sunk cost: no matter what you order, you will have paid the \$20 entrance fee.

The important thing about sunk costs is that they're often not important. *If the same item is in all of the payoff boxes, it's impossible to make a decision solely on the basis of that item.* Sunk costs can provide important background

material for making a decision (see problem 2.2), but the decision-making process depends crucially on items that are not in all the boxes.

Once you think you’ve found the best choice, a good way to check your work is to look at some nearby choices. Such **marginal analysis** basically involves asking “Are you sure you don’t want a little more or a little less?” For example, imagine that an optimizing individual goes to the grocery store, sees that oranges are 25 cents apiece (i.e., that the **marginal cost** of each orange is 25 cents), and decides to buy five. One nearby choice is “Buy four”; since our optimizing individual chose “Buy five” instead, her **marginal benefit** from the fifth orange must be *more* than 25 cents. Another nearby choice is “Buy six”; since our optimizing individual chose “Buy five” instead, her marginal benefit from a sixth orange must be *less* than 25 cents.

As an analogy for marginal analysis, consider the task of finding the highest place on Earth. If you think you’ve found the highest place, marginal analysis helps you verify this by establishing a condition that is *necessary but not sufficient*: in order for some spot to be the highest spot on Earth, it is necessarily true that moving a little bit in any direction cannot take you any higher. Simple though it is, this principle is quite useful for checking your work. (It is not, however, infallible: although all mountain peaks pass the marginal analysis test, not all of them rank as the highest spot on Earth.)

2.2 Example: Monopoly

Recall that one of our assumptions is that firms are profit-maximizing. This assumption takes on extra significance in the case of **monopoly**—i.e., when there is only one seller of a good—because of the lack of competition. We will see in Part III that competition between firms imposes tight constraints on their behavior. In contrast, monopolists have more freedom, and a key component of that freedom is the ability to set whatever price they want for their product.¹ Indeed, we will see in Part II that monopolists will try to charge different people different prices based on their willingness to pay for the product in question.

For now, however, we will focus on a **uniform pricing monopolist**—like the one in Figure 2.2—who must charge all customers the same price. In this situation, the monopolist’s profit is calculated according to

$$\text{Profit} = \text{Total Revenue} - \text{Total Costs} = \text{Price} \cdot \text{Quantity} - \text{Total Costs}$$

From this equation we can see that a profit-maximizing monopolist will try to minimize costs, just like any other firm; every dollar they save in costs is one more dollar of profit. So the idea that monopolists are slow and lazy doesn’t find much support in this model.²

¹For this reason, firms with monopoly power are also known as **price-setting** firms, whereas firms in a competitive market are known as **price-taking** firms.

²More progress can be made in this direction by studying **regulated monopolies** that are guaranteed a certain profit level by the government. Such firms are not always allowed to keep cost savings that they find.

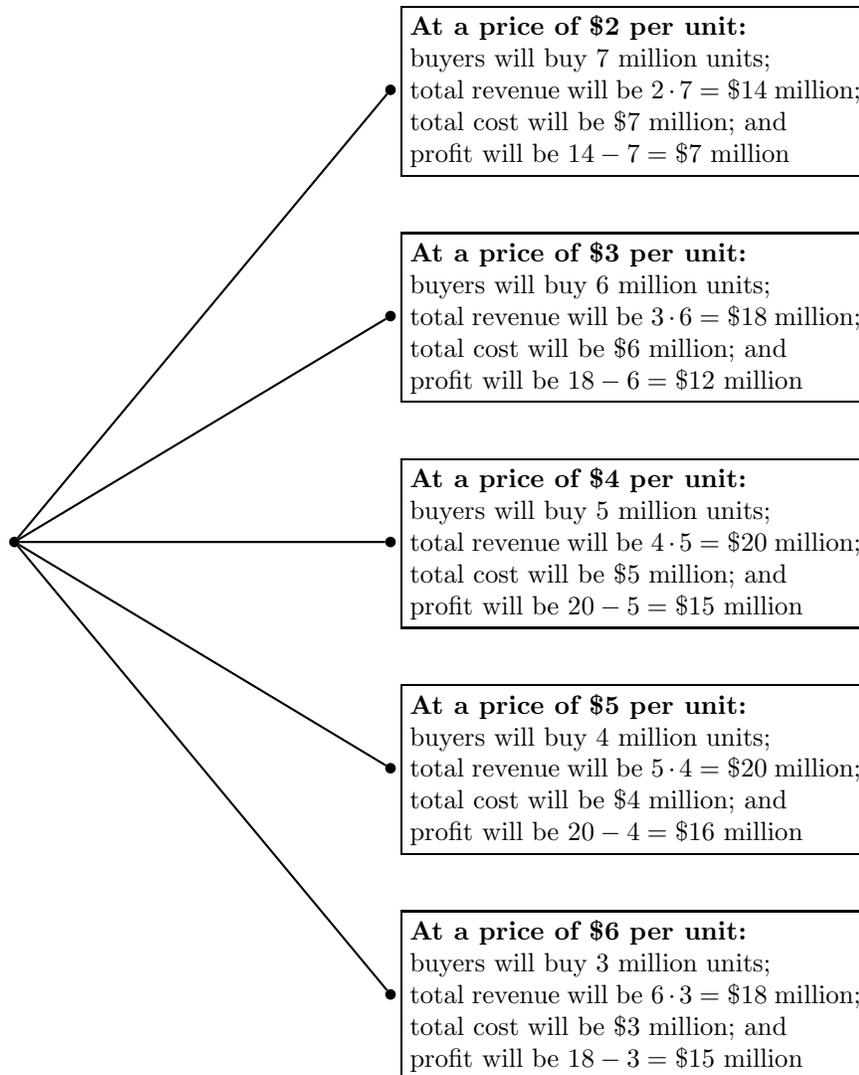


Figure 2.2: A decision tree for a uniform-pricing monopoly. The second line in each payoff box shows the amount that buyers are willing to buy at the given price. The third line shows the monopolist's total revenue if it chooses the given price; note that the highest price (\$6 per unit) yields *less* revenue than lower prices! The fourth line shows the monopolist's costs for producing the amount that buyers want to buy at the given price; in this case there are constant production costs of \$1 per unit. The final line shows the monopolist's profit, identifying a price of \$5 per unit as the profit-maximizing choice.

We can also see the monopolist's **uniform pricing problem**: it would like to sell a large quantity for a high price, but it is forced to choose between selling a smaller quantity for a high price and selling a larger quantity for a low price. In other words, it faces a trade-off between **profit margin** (making a lot of money on each unit by charging a high price) and **volume** (making a lot of money by selling lots of units). The decision tree in Figure 2.2 gives an example showing how a higher price will drive away customers, thereby reducing the quantity that the monopolist can sell. In Part III we will return to the inverse relationship between price and the quantity that customers want to buy, a relationship formally known as the **Law of Demand**.

Problems

Answers are in the next section. The first few words of each question (and each answer) are hyperlinked to allow online users to quickly move back and forth.

- 2.1 [A newspaper column](#) in the summer of 2000 complained about the overwhelming number of hours being devoted to the Olympics by NBC. The columnist argued that NBC had such an extensive programming schedule in order to recoup the millions of dollars it had paid for the rights to televise the games. Do you believe this argument? Why or why not?
- 2.2 [You win \\$1 million](#) in the lottery, and the lottery officials offer you the following bet: You flip a coin; if it comes up heads, you win an additional \$10 million; if it comes up tails, you lose the \$1 million. Will the amount of money you had prior to the lottery affect your decision? (*Hint*: What would you do if you were already a billionaire? What if you were penniless?) What does this say about the importance of sunk costs?
- 2.3 [Alice the axe murderer](#) is on the FBI's Ten Most Wanted list for killing six people. If she is caught, she will be convicted of these murders. The government decides to get tough on crime by passing a new law saying that anybody convicted of murder will get the death penalty. Does this serve as a deterrent for Alice, i.e., does the law give Alice an incentive to stop killing people? Does the law serve as a deterrent for Betty, who is thinking about becoming an axe murderer but hasn't killed anybody yet?
- 2.4 [A drug company](#) comes out with a new pill that prevents baldness. When asked why the drug costs so much, the company spokesman says that they need to recoup the \$1 billion spent on research and development (R&D).
 - (a) Will a profit-maximizing firm pay attention to R&D costs when determining its pricing?
 - (b) If you said "Yes" above: Do you think the company would have charged less for the drug if it had discovered it after spending only \$5 million instead of \$1 billion?

If you said “No” above: Do R&D costs affect the company’s behavior (1) *before* they decide whether or not to invest in the R&D, (2) *after* they invest in the R&D, (3) both before and after, or (4) neither?

2.5 Consider a cap-and-trade system that reduces carbon emissions by creating a limited number of pollution permits that firms can buy and sell. One of the choices in designing such a system is whether to auction permits to the highest bidder or “grandfather” permits by giving them to firms based on their historic emissions.

- (a) Will auctioning versus grandfathering have an impact on how much pollution clean-up a profit-maximizing firm will engage in? Why or why not?
- (b) Will auctioning versus grandfathering have an impact on how much a profit-maximizing firm will charge for its output, e.g., how much a coal plant charges for electricity? Why or why not?

Answers

2.1 You should not believe this argument because the amount spent on the rights to the Olympic Games is a sunk cost. The reason the network showed the Olympics for so many hours was because that was the decision that maximized profit.

2.2 A billionaire would probably take this bet because this is a terrific (although risky) investment. A penniless person would probably not take the bet because the risk of ending up with zero is too great. So sunk costs cannot be entirely ignored in decision-making; rather, the point is that it is not possible to base decisions *solely* on sunk costs.

2.3 The law does not deter Alice from committing additional crimes because she’s already facing the death penalty if she’s caught. The law does deter Betty, because she hasn’t killed anybody yet.

2.4 You should not believe the spokesman’s explanation because the R&D expenditure is a sunk cost. If it spent twice as much or half as much to discover the drug, it should still charge the same price, because that’s the price that maximizes profit. The only time that R&D costs affect the company’s behavior is *before* they’re sunk: when the company is thinking about spending money on R&D, it has to determine whether or not it’s going to be profitable to make that investment given their estimate of how much they’ll be able to charge for the pill. Once they do the R&D, however, it’s a sunk cost and will no longer influence their profit-maximizing decisions.

2.5 [“Consider a cap-and-trade system. . .”]

- (a) No, a profit-maximizing firm will clean up the same amount regardless of whether permits are grandfathered or auctioned; the fact that the firm can sell permits gives it an incentive to clean up even with grandfathered permits. Another way to think about this is to imagine that a firm receiving grandfathered permits sells all of its permits and then buys back what it needs; the revenue from selling the permits becomes a sunk cost that does not affect the firm's decision about how much to clean up.
- (b) Again, a profit-maximizing firm will charge the same amount for its output regardless of whether permits are grandfathered or auctioned. If you imagine that a firm receiving grandfathered permits sells all of its permits and then buys back what it needs, it is clear that it will behave in exactly the same way as a firm that has to purchase all of its permits at auction. (This was a lesson that Europe learned the hard way with their carbon-trading Emissions Trading System that started in 2005–2007; permits were grandfathered and electricity prices still went up! Because of this, a greater percentage of permits will be auctioned in newer versions of the ETS.)

Chapter 3

Time

Motivating question: If you win a “\$20 million” lottery jackpot, you will probably find that you have actually won payments of \$1 million each year for 20 years. Lottery officials generally offer the winner an immediate cash payment as an alternative, but the amount of this **lump sum** payment is usually only half of the “official” prize. So: Would you rather have \$10 million today or \$1 million each year for 20 years? (See Figure 3.1.)

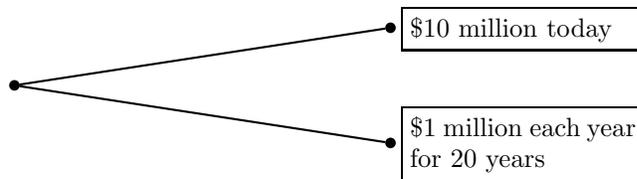


Figure 3.1: A decision tree involving time

This is tricky because simply comparing money today and money tomorrow is like comparing apples and oranges. One reason is **inflation**, a general increase in prices over time. But inflation—which is covered as a supplement in Chapter 16—is not the only reason why money today and money tomorrow are not commensurate, and for clarity *we will assume in this chapter that there is no inflation*.

Despite this assumption, money today is still not commensurate with money tomorrow, for the simple reason that people generally prefer not to wait for things: offered a choice between a TV today and a TV next year, most people prefer the TV today. This *time preference* means that money today is worth more than money tomorrow even in the absence of inflation.

The way to compare money today and money tomorrow is to observe that banks and other financial institutions turn money today into money tomorrow

and vice versa. We can therefore use the relevant **interest rates** to put values from the past or the future into **present value** terms. In other words, we can express everything in terms of today's dollars.

As an example, assume that you can put money in a savings account and earn 10% interest, or that you can borrow money from the bank at 10% interest. Then the present value of receiving \$1 one year from now is about \$0.91: put \$0.91 in the bank today and in a year you'll have about \$1.00; equivalently, borrow \$0.91 today and in a year you'll need \$1.00 to pay back the principal plus the interest.

3.1 Lump sums

Question: If you have \$100 in a savings account at a 5% annual interest rate, how much will be in the account after 30 years?

Answer:

$$\begin{aligned} \text{After 1 year:} & \quad \$100(1.05) = \$100(1.05)^1 = \$105.00 \\ \text{After 2 years:} & \quad \$105(1.05) = \$100(1.05)^2 = \$110.25 \\ \text{After 3 years:} & \quad \$110.25(1.05) = \$100(1.05)^3 = \$115.76. \\ \text{So after 30 years:} & \quad \quad \quad \$100(1.05)^{30} \approx \$432.19. \end{aligned}$$

More generally, the **future value of a lump sum** of \$ x invested for n years at interest rate r is

$$\text{FV} = x(1+r)^n.$$

In addition to being useful in and of itself, the future value formula also sheds light on the topic of present value:

Question: If someone offers you \$100 in 30 years, how much is that worth today if the interest rate is 5%?

Answer: The present value of \$100 in 30 years is that amount of money which, if put in the bank today, would grow to \$100 in 30 years. Using the future value formula, we want to find x such that $x(1.05)^{30} = 100$. Solving this we find that $x = \frac{100}{(1.05)^{30}} \approx \23.14 , meaning that if you put \$23.14 in the bank today at 5% interest, after 30 years you'll have about \$100.

More generally, the **present value of a lump sum payment** of \$ x received at the end of n years at interest rate r —e.g., $r = 0.05$ for a 5% interest rate—is

$$\text{PV} = \frac{x}{(1+r)^n}.$$

Note that the present value formula can also be used to evaluate lump sums received not just in the *future* but also in the *past* (simply use, say, $n = -1$ to calculate the present value of a lump sum received one year ago) and even in the *present*: simply use $n = 0$ to find that you need to put \$ x in the bank today in order to have \$ x in the bank today!

3.2 Annuities and perpetuities

Question: What is the present value of receiving \$100 at the end of each year for the next three years when the interest rate is 5%? (We now have a stream of annual payments, which is called an **annuity**.)

Answer: We need to figure out how much you have to put in the bank today in order to finance \$100 payments at the end of each year for three years. One way is to calculate the present value of each year's payment and then add them all together. The present value of the first \$100 payment, which comes after one year, is $\frac{100}{(1.05)^1} \approx \95.24 , meaning that putting \$95.24 in the bank today will allow you to make the first payment. The present value of the second \$100 payment, which comes after two years, is $\frac{100}{(1.05)^2} \approx \90.70 , meaning that putting \$90.70 in the bank today will allow you to make the second payment. And the present value of the third \$100 is $\frac{100}{(1.05)^3} \approx \86.38 , meaning that putting \$86.38 in the bank today will allow you to make the final payment. So the present value of this annuity is about $\$95.24 + \$90.70 + \$86.38 = \272.32 .

You can also use this method to analyze the lottery question at the beginning of this chapter, but doing twenty separate calculations will get rather tedious. Fortunately, there's a formula that makes things easier. At interest rate r —e.g., $r = 0.05$ for a 5% interest rate—the **present value of an annuity** paying \$ x at the end of each year for the next n years is

$$\text{PV} = x \left[\frac{1 - \frac{1}{(1+r)^n}}{r} \right].$$

The derivation of this formula involves some pretty mathematics, so if you're interested here's an example based on the 3-year \$100 annuity. We want to calculate the present value (PV) of this annuity:

$$\text{PV} = \frac{100}{(1.05)^1} + \frac{100}{(1.05)^2} + \frac{100}{(1.05)^3}.$$

If we multiply both sides of this equation by 1.05, we get

$$1.05 \cdot \text{PV} = 100 + \frac{100}{(1.05)^1} + \frac{100}{(1.05)^2}.$$

Now we subtract the first equation from the second equation:

$$1.05 \cdot \text{PV} - \text{PV} = \left[100 + \frac{100}{(1.05)^1} + \frac{100}{(1.05)^2} \right] - \left[\frac{100}{(1.05)^1} + \frac{100}{(1.05)^2} + \frac{100}{(1.05)^3} \right].$$

The left hand side of this equation simplifies to $0.05 \cdot PV$. The right hand side also simplifies (all the middle terms cancel!), yielding

$$0.05 \cdot PV = 100 - \frac{100}{(1.05)^3}.$$

Dividing both sides by 0.05 and grouping terms gives us a result that parallels the general formula:

$$PV = 100 \left[\frac{1 - \frac{1}{(1.05)^3}}{0.05} \right].$$

Perpetuities

Question: What is the present value of receiving \$100 at the end of each year *forever* at a 5% interest rate? Such a stream of payments is called a **perpetuity**—i.e., a *perpetual annuity*—and it really is forever: you can pass it along in your will!

Answer: As with annuities, there's a nice formula. The **present value of a perpetuity** paying \$ x at the end of each year *forever* at an interest rate of r —e.g., $r = 0.05$ for a 5% interest rate—is

$$PV = \frac{x}{r}.$$

So at a 5% interest rate the present value of receiving \$100 at the end of each year forever is only \$2,000!

One explanation here is *living off the interest*: put \$2,000 in the bank at 5% interest and at the end of a year you'll have \$100 in interest. Take out that interest—leaving the \$2,000 principal—and in another year you'll have another \$100 in interest. Living off the interest from your \$2,000 principal therefore provides you with \$100 each year in perpetuity.

As with annuities, we can also do some nifty mathematics. What we're looking for is

$$PV = \frac{100}{(1.05)^1} + \frac{100}{(1.05)^2} + \frac{100}{(1.05)^3} + \dots$$

To figure this out, we apply the same trick we used with annuities. First multiply through by 1.05 to get

$$1.05 \cdot PV = 100 + \frac{100}{(1.05)^1} + \frac{100}{(1.05)^2} + \frac{100}{(1.05)^3} + \dots$$

Then subtract the first equation above from the second equation—note that almost everything on the right hand side cancels!—to end up with

$$1.05 \cdot PV - PV = 100.$$

This simplifies to $PV = \frac{100}{0.05} = \$2,000$.

3.3 Capital theory

We can apply the material from the previous sections to study investment decisions, also called **capital theory**. This topic has some rather unexpected applications, including natural resource economics.

Motivating question: You inherit a lake with some commercially valuable fish in it. If you are profit-maximizing, how should you manage this resource?

The key economic idea of this section is that **fish are capital**, i.e., fish are an investment, just like a savings account is an investment. To “invest in the fish”, you simply leave them alone, thereby allowing them to reproduce and grow bigger so you can have a bigger harvest next year. This exactly parallels investing in a savings account: you simply leave your money alone, allowing it to gain interest so you have a bigger balance next year.

In managing your lake you need to compare your options, each of which involves some combination of “investing in the fish” (by leaving some or all of the fish in the lake so that they will grow and reproduce) and investing in the bank (by catching and selling some or all the fish and putting the proceeds in the bank). This suggests that you need to compare the interest rate at the Bank of America with the “interest rate” at the “Bank of Fish”. (See Figure 3.2.)



Figure 3.2: The Bank of America and the Bank of Fish

This sounds easy, but there are lots of potential complications: the price of fish could change over time, the effort required to catch fish could vary depending on the number of fish in the lake, etc. To simplify matters, let’s assume that the cost of fishing is zero, that the market price of fish is constant over time at \$1 per pound, and that the interest rate at the Bank of America is 5%.

One interesting result concerns **Maximum Sustainable Yield (MSY)**, defined as the maximum possible catch that you can sustain year after year forever. Although it sounds attractive, MSY is *not* the profit-maximizing approach

to managing this resource. To see why, let's consider the example of Figure 3.3, assuming that the lake starts out with 100 pounds of fish. First calculate the present value of following MSY: each year you start out with 100 pounds of fish and at the end of the year you catch 10 pounds of fish and sell them for \$1 per pound, so the perpetuity formula gives us a present value of $\frac{10}{0.05} = \$200$.

Now consider an alternative policy: as before, you start out with 100 pounds of fish, but now you catch 32 pounds *immediately*, reducing the population to 68. This lower population size corresponds to a sustainable yield of 9 pounds per year (as shown in Figure 3.4), so at the end of every year you can catch 9 pounds of fish. The present value of this harvest policy is $32 + \frac{9}{0.05} = \$212$, which is more than the present value from the MSY policy!

So MSY is *not* the profit-maximizing harvest plan: you can make more money by catching additional fish and investing the proceeds in the Bank of America. Intuitively, the problem with the MSY policy is that population growth is about the same whether you start out with 68 pounds of fish or 100 pounds of fish; *at the margin*, then, the interest rate you're getting at the Bank of Fish is really low. Since the return on those final 32 fish is so small, you're better off catching them and putting the money in the bank.

Question: What happens if the fish are lousy at growing and reproducing?

Answer: Well, you're tempted to kill them all right away, since the interest rate at the Bank of Fish is lousy. This is part of the explanation for the decimation of rockfish, orange roughy, and other slow-growing fish species. (A much bigger part of the explanation—discussed as a supplement in Chapter 17—is that these fish do not live in privately controlled lakes but in oceans and other public areas where anyone who wants to can go fishing. The incentives for “investing in the fish” are much lower in such **open-access** fisheries, which are like a bank account for which everyone has an ATM card.)

But other investments can be even worse than fish at growing and reproducing. Consider oil, for example, or gold, or Microsoft stock: these don't grow at all. So why do people invest in them? It must be *because they think the price is going to go up*. If I'm sitting on an oil well or a gold mine or some Microsoft stock and I don't think the price of oil or gold or Microsoft stock is going to increase faster than the interest rate at the Bank of America, I should withdraw my capital from the Bank of Oil or the Bank of Gold or the Bank of Microsoft and put it in the Bank of America.

Phrases such as “fish are capital” and “oil is capital” emphasize the economic perspective that fish, oil, and many other seemingly unrelated items are all investments. As such, an optimizing individual searching for the best investment strategy needs to examine these different “banks” and compare their rates of return with each other and with the rate of return at an actual bank.

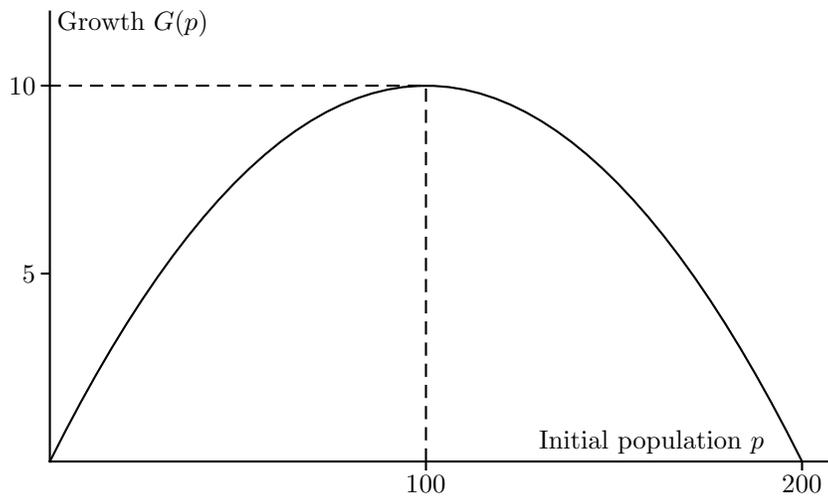


Figure 3.3: With an initial population of 100 pounds of fish, population growth over the course of one year amounts to 10 additional pounds of fish. Harvesting 10 pounds returns the population to 100, at which point the process can begin again. An initial population of 100 pounds of fish therefore produces a **sustainable yield** of 10 pounds per year; the graph shows that this is the **maximum sustainable yield**, i.e., the maximum amount that can be harvested year after year forever.

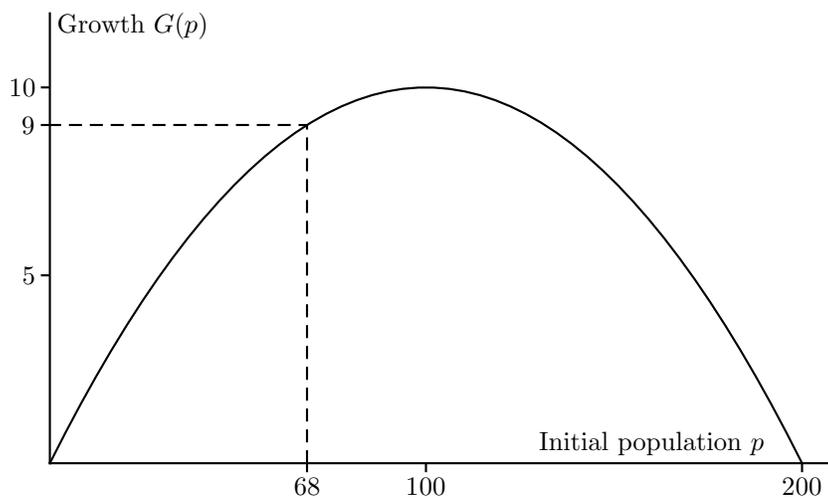


Figure 3.4: With an initial population of 68 pounds of fish, population growth over the course of one year amounts to 9 additional pounds of fish. Harvesting 9 pounds returns the population to 68, at which point the process can begin again. An initial population of 68 pounds of fish therefore produces a sustainable yield of 9 pounds per year.

3.4 Algebra: Budget constraints

The first of two topics here is budget constraints. Consider an individual who lives for two time periods, today and tomorrow (which we will call periods $t = 1$ and $t = 2$, respectively). If he starts out with wealth of $w = 100$ and can borrow or save money in a bank at an interest rate of $100 \cdot s\% = 20\%$, he has many different consumption options. For example, he could spend all his money today and nothing tomorrow (i.e., $(c_1, c_2) = (100, 0)$ where c_1 is consumption today and c_2 is consumption tomorrow), or he could spend nothing today and $100(1.20)$ tomorrow (i.e., $(c_1, c_2) = (0, 120)$), or anything in between. The constraint on his spending (which is called the **budget constraint**) is that the present value of his spending must equal the present value of his wealth:

$$c_1 + \frac{c_2}{1.20} = 100.$$

Figure 3.5 shows the graph of this budget constraint, which we can rewrite as $c_2 = 1.20(100 - c_1)$. Note that the slope of the budget constraint is $-1.20 = -(1 + s)$, i.e., is related to the interest rate.

The budget constraint can also be used to compare different endowments (e.g., lottery payoffs). For example, an endowment of 50 in the first period and 60 in the second period also lies on the budget constraint shown in Figure 3.5. The present value of this endowment is the x -intercept of the budget constraint, i.e., 100; the future value of this endowment is the y -intercept, i.e., 120.

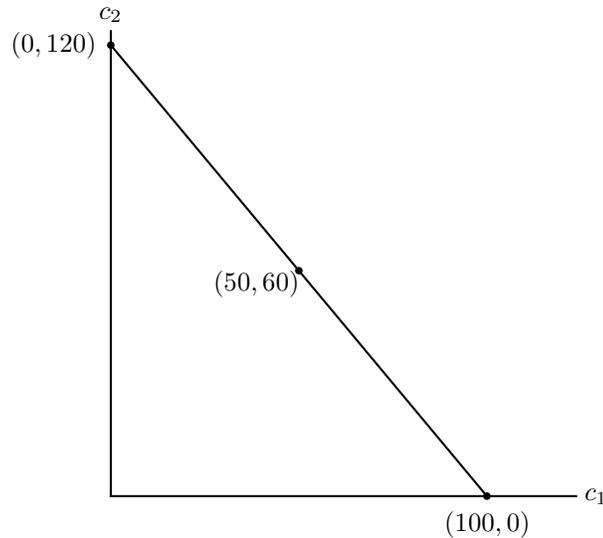


Figure 3.5: A budget constraint corresponding to a present value of 100

3.5 Algebra: Continuous compounding

Motivating question: If you had \$100 in the bank and the interest rate was 10% per year, would you rather have interest calculated every year or every six months? In other words, would you rather get 10% at the end of each year or 5% at the end of every six months?

Answer: Let's compare future values at the end of one year. With an annual calculation (10% every year) you end up with $\$100(1.1) = \110 . With a semi-annual calculation (5% every six months) you end up with $\$100(1.05) = \105 at the end of six months and $\$105(1.05) = \110.25 at the end of one year. So you should choose the semiannual compounding.

The "bonus" from the semi-annual calculation comes from the **compounding** of the interest: at the end of six months you get an interest payment on your \$100 investment of $\$100(.05) = \5 ; at the end of one year you get an *interest payment on that interest payment* of $\$5(.05) = \$.25$, plus an interest payment on your original \$100 investment of $\$100(.05) = \5 .

Next: Let's put \$1 in the bank at an interest rate of 100% per year and see what happens if we make the compounding interval smaller and smaller:

- If we compound once a year we get $\$1(1 + 1) = \2 .
- If we compound every month we get $\$1 \left(1 + \frac{1}{12}\right)^{12} \approx \2.61 .
- If we compound every hour we get $\$1 \left(1 + \frac{1}{365 \cdot 24}\right)^{(365 \cdot 24)} \approx \2.718127 .

Comparing this with the value of $e \approx 2.7182818$ suggests that there might be a connection, and there is: one definition of e is

$$e = \lim_{n \rightarrow \infty} \left(1 + \frac{1}{n}\right)^n.$$

A more general result about e leads to the following definition: **Continuous compounding** is the limit reached by compounding over shorter and shorter intervals. If the interest rate is $100 \cdot r\%$ per year and you invest x with continuous compounding, after t years your investment will be worth $x \cdot e^{rt} = x \cdot \lim_{n \rightarrow \infty} \left(1 + \frac{r}{n}\right)^n$.

Example: What is the present value of receiving \$100 in 3 years at a 10% interest rate, compounded continuously?

Answer: We want an amount x such that $x \cdot e^{(0.1)(3)} = 100$. The solution is $x = 100 \cdot e^{-0.3} \approx \74.08 .

Supplemental material

In this chapter we have assumed that there is no inflation; Chapter 16 expands the discussion to include inflation.

Problems

Answers are in the next section. The first few words of each question (and each answer) are hyperlinked to allow online users to quickly move back and forth.

3.1 Say you have \$100 in the bank today.

- (a) How much will be in the bank after 30 years if the interest rate is 5%? Call this amount y .
- (b) What is the present value of receiving y after 30 years? Call this amount z .
- (c) How does z compare with the \$100 you currently have in the bank? Can you explain why by looking at the relationship between the formulas for the present value and future value of lump sums?

3.2 (The Rule of 72): A rule of thumb is that if you have money in the bank at $r\%$ (e.g., 10%), then your money will double in $\frac{72}{r}$ years, e.g., 7.2 years for a 10% interest rate.

- (a) How many years will it actually take your money to double at 10%? (You can find the answer—plus or minus one year—through trial-and-error; if you know how to use logarithms—no this won't be on the test—you can use them to get a more precise answer.) Compare the true answer with the Rule of 72 approximation.
- (b) Do the same with a 5% interest rate and a 100% interest rate.
- (c) Do your results above suggest something about when the Rule of 72 is a good approximation?

3.3 Investment #1 pays you \$100 at the end of each year for the next 10 years. Investment #2 pays you nothing for the first four years, and then \$200 at the end of each year for the next six years.

- (a) Calculate the present value of each investment if the interest rate is 5%. Which one has a higher present value?
- (b) Which investment has the greater present value at an interest rate of 15%?

- (c) Do higher interest rates favor investment #1 or #2? Can you explain why using intuition and/or math?
 - (d) Can you think of any real-life decisions that have features like these? (Hint: This about deciding whether or not to go to college!)
- 3.4 Intuitively, how much difference do you think there is between an annuity paying \$100 each year for 100 years and a perpetuity paying \$100 each year forever? Can you mathematically confirm your intuition by relating the annuity formula to the perpetuity formula?
- 3.5 Explain the perpetuity formula in terms of “living off the interest”.
- 3.6 Consider a “\$20 million” lottery payoff paying \$1 million at the end of each year for 20 years.
- (a) Calculate the present value of this payoff if the interest rate is 5%.
 - (b) Calculate the present value of the related perpetuity paying \$1 million at the end of each year forever.
 - (c) Assume that the lump sum payoff for your \$20 million lottery is \$10 million, i.e., you can opt to get \$10 million now instead of \$1 million at the end of each year for 20 years. Using trial and error, estimate the interest rate r that makes the lump-sum payoff for the lottery equal to the annuity payoff.
 - (d) Calculate the present value of winning \$1 million at the *beginning* of each year for 20 years. Again, assume the interest rate is 5%. *Hint:* there are easy ways and hard ways to do this!
- 3.7 *Fun.* Compound interest has been called the Eighth Wonder of the World. Here’s why.
- (a) According to the Legend of the Ambalappuzha Paal Paayasam (more [online¹](http://en.wikipedia.org/wiki/Ambalappuzha)), Lord Krishna once appeared in discuss in the court of a great king and challenged the king to a game of chess, with the prize being “just” one grain of rice for the first square of the board, two for the second, four for the third, and so on, doubling on each of the 64 squares of the chessboard. This seemed reasonable enough, so the king agreed and promptly lost the chess match. How many grains of rice was he supposed to deliver for the 64th square?
 - (b) On Day 1, somebody dropped a lily pad in a mountain pond. The number of lily pads (and the percentage of pond they covered) doubled every day. On the 30th day, the pond was completely covered in lily pads. On which day was the pond half-covered?
 - (c) How tall do you think a piece of paper would be if you could fold it in half again and again and again, 40 times? Estimate the thickness of a piece of paper (just guess!) and then calculate this height.

¹<http://en.wikipedia.org/wiki/Ambalappuzha>

Comment: These examples involve interest rates of 100% (i.e., doubling), but you will get similar results with much smaller interest rates as long as your time horizons are long enough. This is because all interest rate problems share a common feature: constant doubling time. Put \$100 in the bank at 100% interest and it will double every year: \$200, \$400, \$800... At 1% interest it will double every 70 years: \$200, \$400, \$800... So 1% growth and 100% growth are different in degree but not in spirit.

- 3.8 Explain (as if to a non-economist) the phrases “fish are capital,” “trees are capital,” and/or “oil is capital,” or otherwise explain the importance of the interest rate at the Bank of America in management decisions regarding natural resources such as fish, trees, and oil.
- 3.9 *Fun.* Here is some information from the National Archives:

In 1803 the United States paid France \$15 million (\$15,000,000) for the Louisiana Territory—828,000 square miles of land west of the Mississippi River. The lands acquired stretched from the Mississippi River to the Rocky Mountains and from the Gulf of Mexico to the Canadian border. Thirteen states were carved from the Louisiana Territory. The Louisiana Purchase nearly doubled the size of the United States, making it one of the largest nations in the world.

At first glance, paying \$15 million for half of the United States seems like quite a bargain! But recall that the Louisiana Purchase was over 200 years ago, and \$15 million then is not the same as \$15 million now. Before we can agree with General Horatio Grant, who told President Jefferson at the time, “Let the Land rejoice, for you have bought Louisiana for a song,” we should calculate the present value of that purchase. So: If President Jefferson had not completed that purchase and had instead put the \$15 million in a bank account, how much would there be after 200 years at an interest rate of: (a) 2%, or (b) 8%? (See problem 16.5 on page 162 for more about the Louisiana Purchase.)

- 3.10 It is sometimes useful to change interest rate time periods, i.e., to convert a monthly interest rate into a yearly interest rate, or vice versa. As with all present value concepts, this is done by considering money in the bank at various points of time.
- (a) To find an annual interest rate that is approximately equal to a monthly interest rate, multiply the monthly interest rate by 12. Use this approximation to estimate an annual interest rate that is equivalent to a monthly rate of 0.5%.
- (b) Assume you have \$100 in the bank at a monthly interest rate of 0.5%. Use the future value formula to determine how much money will actually be in the bank at the end of one year. What annual interest rate is actually equivalent to a monthly rate of 0.5%?

- (c) To find a monthly interest rate that is approximately equal to an annual interest rate, divide the annual interest rate by 12. Use this approximation to estimate a monthly interest rate that is equivalent to an annual rate of 6% (Note that you can use logarithms—no this won't be on the test—to determine the actual monthly interest rate that is equivalent to an annual rate of 6%.)
- 3.11 Imagine that you own a lake and that you're trying to maximize the present value of catching fish from the lake, which currently has 1000 fish in it. The population growth function of the fish is described in Figure 3.6. Assume that fish are worth \$1 per pound and that the interest rate at the bank is 5%.
- (a) The Maximum Sustainable Yield (MSY) policy is to catch 100 fish at the end of this year, 100 fish at the end of the next year, and so on, forever. What is the present value from following MSY?
- (b) An alternative policy is to catch 400 fish *today* (so that 600 remain in the lake), and then catch 84 fish at the end of this year, 84 fish at the end of the next year, and so on, forever. What is the resulting present value? Is it higher or lower than the present value of the maximum sustainable yield policy?
- 3.12 Assume that you've just bought a new carpet. The good news is that the carpet will last forever. The bad news is that you need to steam-clean it at the end of every year, i.e., one year from today, two years from today, etc. What you need to decide is whether to buy a steam-cleaner or just rent one every year. *You can use the bank to save or borrow money at a 5% interest rate.*

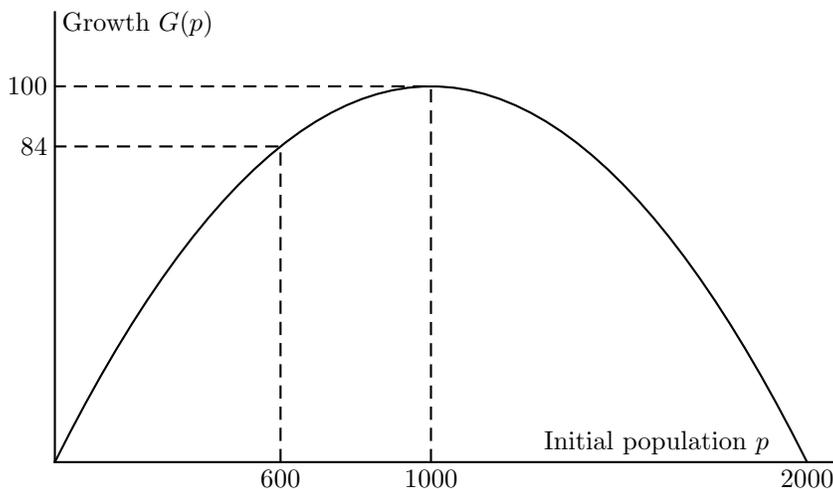


Figure 3.6: A population growth function for fish.

- (a) Will the amount you paid for the carpet affect your decision regarding renting versus buying?
 - (b) One year from today, i.e., when you first need to clean the carpet, you'll be able to buy a steam-cleaner for \$500; like the carpet, the steam-cleaner will last forever. Calculate the present value of this cost.
 - (c) The alternative to buying is renting a steam-cleaner, which will cost you \$20 at the end of every year forever. Calculate the present value of this cost. Is it better to rent or buy?
 - (d) Imagine that your friend Jack calls to tell you that steam-cleaners are on sale (today only!) for \$450: "You'd have to be a dummy to pay \$20 every year forever when you can just pay \$450 today and be done with it!" Write a brief response explaining (as if to a non-economist) why you do or do not agree.
- 3.13 **Geothermal energy** involves "mining" heat by drilling into the earth's crust. Like many clean energy technologies, it has high up-front costs but promises to pay off over time. The made-up numbers in this problem look at the economics of geothermal power. (For the real numbers, see the 2007 MIT report "The Future of Geothermal Energy".)
- (a) Consider spending \$1000 today to build a geothermal plant that will generate \$100 at the end of each year for the next 30 years. Show that the present value of the costs outweigh the present value of the benefits if the interest rate is 13%.
 - (b) In order to make geothermal more attractive, does the interest rate need to go up or down? Briefly explain.
 - (c) What will the present value of benefits be if the plant generates \$100 a year *forever* instead of just for 30 years? (The interest rate is still 13%.)
 - (d) Explain (as if to a non-economist) why the present value of the \$100-a-year stream of benefits will not be infinite even if the plant can operate forever. Do *not* talk about inflation.
- 3.14 **Just about everybody** agrees that the Social Security system faces financial troubles down the road: after the Baby Boomers retire, the money coming into the system through payroll taxes is not expected to be enough to finance the benefits that the system promises to retirees. The Social Security Board of Trustees said in their 2008 report that the system's deficit over the next 75 years has a present value of \$4.3 trillion. Explain what this means in English, i.e., as if to a non-economist.
- 3.15 **Consider a choice** between receiving a lump sum of \$100 today and receiving an annuity of \$20 every year for 10 years. *As always, banks are standing by to accept deposits and/or make loans at the nominal interest rate.*

- (a) One issue that might affect your choice is the interest rate. Compared to a “low” interest rate (say, 3%), does a “high” interest rate (say, 7%) favor the lump sum or the annuity? (Although it will almost certainly help to do a numerical example with these numbers, this question is really about a more general issue: do higher higher interest rates favor “money today” or “money tomorrow”?) Support your answer with a brief, intuitive (i.e., non-mathematical) explanation.
 - (b) Another issue that might affect your choice is your preference for “money today” versus “money tomorrow”; for example, you might *really* want money today so that you can buy a new computer. Does this mean you should choose the \$100 lump sum even if the annuity has a higher present value? Explain *briefly* why or why not.
- 3.16 Consider choosing between an annuity paying \$100 at the end of every year for 250 years and a perpetuity paying \$100 at the end of every year forever. The *difference between these two options* is, well, it’s an infinite number of \$100 payments beginning at the end of year 251. The *difference between the present values of these two options* is, well, at an interest rate of 5% it’s about \$0.01 (or, even more precisely, about \$0.0100857). Explain (as if to a non-economist) what this means and why it makes sense.
- 3.17 For the sake of simplicity, textbooks often assume that interest is calculated and payments are made once a year. In reality, the relevant time frames are often shorter periods, as with monthly car payments. Fortunately, you can use the same formulas and calculations as long as you make appropriate adjustments. For example, these formulas are behind the online “car payment calculators” that allow you to calculate monthly payments given the loan amount (say, \$20,000), the length of the loan (say, 48 months), and the interest rate (say, 9%, given as an Annual Percentage Rate or APR).
- (a) An approximate *monthly* interest rate can be derived from a *yearly* interest rate simply by dividing by 12. Determine the approximate monthly interest rate for a 9% APR and then show that this is an over-estimate of the actual monthly interest rate. (Hint: Put \$100 in the bank at your approximate monthly interest rate and show that you’ll have more than \$109 after a year.) *Explain why this approximation is an over-estimate.*
 - (b) An APR of 9% is equivalent to a monthly interest rate of something closer to 0.72%. Show that this is true by calculating how much you’ll have in the bank after a year if you put in \$100 at a monthly interest rate of 0.72%.
 - (c) Show that the present value of 48 monthly payments of \$494.29 at a 0.72% monthly interest rate is about \$20,000.
 - (d) Estimate the monthly payments (at 0.72% for 48 months) of a \$30,000 loan.

Answers

3.1 [“Say you have...”]

- (a) Plug \$100, 5%, and 30 years into the future value of a lump sum formula to get $y \approx \$432.19$.
- (b) WPlug \$432.19, 5%, and 30 years into the present value of a lump sum formula to get $z \approx \$100$.
- (c) They are equal. The explanation here is that the formulas for present value and future value of lump sums are inverses of each other in that you can rearrange either equation to get the other:

$$PV = \frac{FV}{(1+r)^n} \iff FV = (PV)(1+r)^n.$$

3.2 [“The Rule of 72...”]

- (a) Solving $2x = x(1.10)^t$ using logarithms or trial-and-error yields $t = 7.27$, pretty close to the Rule of 72 estimate of 7.2.
- (b) The Rule of 72 predicts 14.4 years at 5% and .72 years at 100%. The correct answer at 100% is 1 year (obviously, since the interest rate is 100%), so the Rule of 72 is not such a good approximation. The correct answer at 5% comes from solving $2x = x(1.05)^t$ using logarithms. We get $t \approx 14.2$, which is quite accurate.
- (c) They suggest that the Rule of 72 works well for small interest rates, but not for large ones.

3.3 [“Investment #1 pays...”]

- (a) For Investment #1, use the annuity formula to get a present value of about \$772 at a 5% interest rate. For Investment #2, the brute force way to calculate the present value is to calculate the present value of each of the 6 lump sum payments and then add them up to get about \$835. A more elegant way is to note that Investment #2 is equivalent to a ten-year annuity *minus* a four-year annuity. You can therefore use the annuity formula to calculate the present value of the ten-year annuity (\$1,544) and the four-year annuity (\$709). Subtracting one from the other gives a present value for Investment #2 of \$835.
Investment #2 is the better option at a 5% interest rate.
- (b) Following the same process described above, Investments #1 and #2 have present values of \$502 and \$433, respectively. So Investment #1 has a greater present value at a 15% interest rate.
- (c) Higher interest rates favor Investment #1, in which you essentially forfeit money in the distant future in order to get more money in the immediate future. Since higher interest rates make the future less important, they make Investment #1 more attractive.

- (d) There are many real-life decisions with similar features; for example, Investment #1 could be going to work right out of high school, and Investment #2 could be going to college for 4 years first to increase your earnings potential.

3.4 Intuitively, there shouldn't be much difference. Mathematically, the annuity formula approaches the perpetuity formulas as n approaches infinity:

$$\lim_{n \rightarrow \infty} x \left[\frac{1 - \frac{1}{(1+r)^n}}{r} \right] = \frac{x}{r}.$$

In the specific example of comparing a \$100-per-year perpetuity with a \$100-per-year 100-year annuity, let's do the math with an interest rate of 5%. Using the perpetuity formula gives us a present value of \$2000; using the annuity formula with $n = 100$ gives us a present value of \$1984.79. So the difference is only \$15.21. (On a side note, that difference is also the present value of receiving a \$100-per-year perpetuity starting in 100 years: use the lump-sum formula to discount \$2000 back 100 years and you get \$15.21. Can you see why?)

3.5 The perpetuity formula says that the present value of a perpetuity paying x every year is $\frac{x}{r}$. This is like living off the interest because if you put $\frac{x}{r}$ in the bank, every year you will get interest of $r \cdot \frac{x}{r} = x$, so with this principal you can finance annual consumption of x forever.

3.6 ["Consider a "\$20 million" lottery payoff..."]

- Plug \$1 million, 5%, and 20 years into the annuity formula to get about \$12.5 million as the present value of the annuity.
- Plug \$1 million and 5% into the perpetuity formula to get \$20 million as the present value of the perpetuity. Note that the extra payments you get—\$1 million annually beginning in year 21—are only worth about \$7.5 million in present value terms!
- Increasing r will make the lump sum payment more attractive, and decreasing r will make the annual payments more attractive. Trial and error yields $s \approx .075$ as the interest rate that makes the two payoffs equal in present value terms.
- The hard way to do this is to just calculate the present value of each payment and then add them all together. Easy way #1 is to realize that the difference between the end-of-year payments and the beginning-of-year payments is just an extra payment at the beginning of the first year and a lost payment at the end of the 20th year. The present value of \$1 million today is \$1 million, and the present value of \$1 million at the end of 20 years is \$380,000. Their difference is

\$620,000, so adding this to the answer from (a) yields \$13.08 million. Easy way #2 is to see that the answer from (a) is the right answer from the perspective of one year ago, so using the future value of a lump sum formula to push this answer one year into the future gives us $\$12.46(1.05) = \13.08 million.

3.7 [“Compound interest. . .”]

- (a) The king was supposed to deliver $2^{63} \approx 10,000,000,000,000,000,000$ grains of rice for the last square. According to legend, Lord Krishna informed the king that he could pay on the installment plan, and ever since the king made rice freely available to pilgrims in the nearby temple.
- (b) The pond was half-covered on the 29th day.
- (c) The answer is the height of a single sheet of paper multiplied by $2^{40} \approx 1,000,000,000,000$. If 1,000 pages makes an inch, then this gives us 1,000,000,000 inches, or about 83 million feet, or about 16,000 miles.

3.8 To maximize your present value you need to compare the return you’ll get from “investing in the fish” (or the trees, or the oil) with the return you’ll get from investing in the bank. Investing in the bank means catching the fish, cutting down the trees, or selling the oil and putting the proceeds in the bank. Investing in the fish means letting the fish grow and reproduce so there will be more fish next year; investing in the trees means letting the trees grow so there will be more lumber next year; investing in the oil means keeping the oil in the hopes that the price will go up next year.

3.9 For 2%, plug \$15 million, 200 years, and 2% into the future value of a lump sum formula to get a current bank account balance of \$787 million. For 8%, plug \$15 million, 200 years, and 8% into the future value of a lump sum formula to get a current bank account balance of \$72.6 trillion.

3.10 [“It is sometimes useful. . .”]

- (a) The approximate annual interest rate is 6%.
- (b) Use the future value of a lump sum formula to calculate how much money we’ll have at the end of 12 months if we put \$100 in the bank at a monthly interest rate of 0.5%, i.e. $s = 0.005$: $FV = \$100(1.005)^{12} \approx 106.17$. So a monthly interest rate of 0.5% actually corresponds to an annual interest rate of 6.17%.
- (c) The approximate monthly interest rate is 0.5%.

3.11 [“Imagine that you own. . .”]

- (a) Plug \$100 and 5% into the perpetuity formula to get a present value of \$2000.

- (b) Plug \$84 and 5% into the perpetuity formula to get a present value of \$1680. Adding this to the \$400 you get from catching 400 fish today and you get a present value of \$2080, which is higher than the present value of the maximum sustainable yield policy.

3.12 [“Assume that you’ve just bought...”]

- (a) No, this is a sunk cost.
- (b) Use the present value of a lump sum formula to get a present value of $\frac{\$500}{1.05} \approx \476.19 .
- (c) Use the present value of a perpetuity formula to get a present value of $\frac{\$20}{.05} = \400 . So it’s better to rent.
- (d) “Jack, I disagree with you. Instead of paying \$450 today to buy a steam-cleaner, I’d rather put that \$450 in the bank and ‘live off the interest’. At the end of every year I’d have \$22.50 in interest, which would pay for the annual rental of a steam-cleaner *and* leave me with \$2.50 left over for wild parties.” (Alternately, you could put \$50 towards a wild party today and put the remaining \$400 in the bank; the interest payments would then be \$20 per year, exactly enough to rent a steam-cleaner.)

3.13 [“Geothermal energy...”]

- (a) Plug \$1000, 0.13, and 30 years into the annuity formula to get a present value of about \$749.57 for benefits. Since the present value of costs is \$1000, the costs are greater than the benefits.
- (b) The interest rate needs to go down. Lower interest rates will increase the present value of the benefits without increasing the present value of the costs. Alternatively, you can imagine borrowing the \$1000 and having to pay it back plus interest. The lower the interest rate, the easier it will be for you to pay back the money with the revenues generated by the power plant.
- (c) Plug \$100 and 0.13 into the present value of a perpetuity formula to get \$769.23.
- (d) Put \$769.23 in the bank at 13% interest and each year you’ll get \$100 in interest. By “living off the interest”, you can generate payments of \$100 at the end of each year forever with an initial investment of only \$769.23.

3.14 It means that putting \$4.3 trillion in the Social Security trust fund *now* will keep the system solvent for the next 75 years.

3.15 [“Consider a choice...”]

- (a) Higher interest rates favor the lump sum payment (“money today”) because higher interest rates make the future less important: you need to put less money in the bank today in order to get \$20 in 10 years if the interest rate goes up from 7% to 10%.
 - (b) No: you can use the bank to transfer money between time periods. If the annuity has a higher value, you should choose the annuity and then borrow against it (or sell it) in order to have access to money today.
- 3.16 If you put \$0.0100857 in the bank today at 5% interest, at the end of 250 years you’ll have about \$2000 and can then “live off the interest”, getting interest payments of \$100 every year forever.
- 3.17 [“For the sake of simplicity...”]
- (a) Divide 9% by 12 to get 0.75%. To see how much you’ll have after a year if you put \$100 in the bank at a 0.75% monthly interest rate, plug \$100, 12 months, and 0.75% into the future value formula to get \$109.38. The extra \$0.38 is from *compound interest*.
 - (b) Plug \$100, 12 months, and 0.72% into the future value formula to get \$108.99, very close to \$109.
 - (c) Plug \$494.29, 48 months, and 0.72% into the annuity formula to get a present value of \$20,000.
 - (d) Plug \$ x (the monthly amount we’re trying to solve for), 48 months, and 0.72% into the annuity formula to get $PV = x \cdot 40.462$. We want the present value to equal \$30,000, so solving for x gives us monthly payments of about \$741.43.

Chapter 4

Risk

Motivating question: What is the maximum amount $\$x$ you would pay to play a game (see Figure 4.1) in which you flip a coin and get $\$10$ if it comes up heads and $\$0$ otherwise?

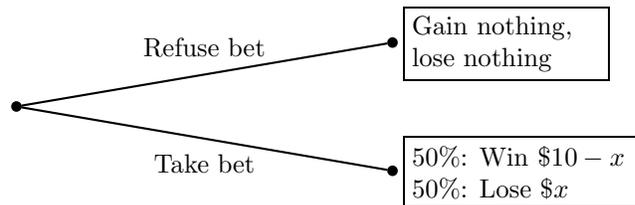


Figure 4.1: A decision tree involving risk

The important issue in this game is your attitude toward **risk**. People who are **risk-averse** buckle their seatbelts, buy insurance, and otherwise try to avoid risk; if you are this type of person you will be unwilling to pay $\$5$ or more to play this game. People who are **risk-loving** go skydiving, drive recklessly, or engage in other risk-seeking behaviors; if you are this type of person, you might be willing to pay more than $\$5$ to play this game. People who are **risk-neutral** are ambivalent about risk; if you are this type of person, you'd be willing to play this game for less than $\$5$, you'd avoid playing the game for more than $\$5$, and you would be indifferent about playing the game for exactly $\$5$.

What's the deal with $\$5$?

The deal is that $\$5$ is the **expected value** of this game. Since you have a 50% chance of getting $\$10$ and a 50% chance of getting $\$0$, it makes some sense to say that *on average* you'll get $\$5$. (Section 4.1 fleshes out this idea.) The concept of expected value provides a valuable perspective by condensing this risky situation

into a single number. Note, however, that it conveys only one perspective; we'll soon see that different risky situations can have the same expected value, and to help differentiate them we'll introduce another concept called **variance**.

Mathematically, an expected value calculation weighs each possible outcome by its likelihood, giving more weight to more likely outcomes and less weight to less likely outcomes. In the game described above, the two possible outcomes are heads (H) and tails (T), so the expected value is

$$\begin{aligned} EV &= \Pr(H) \cdot (\$10) + \Pr(T) \cdot (\$0) \\ &= \frac{1}{2} \cdot (\$10) + \frac{1}{2} \cdot (\$0) \\ &= \$5. \end{aligned}$$

More generally, we can write the formula for expected value as:

$$\text{Expected Value} = \sum_{\text{Outcomes } i} \text{Probability}(i) \cdot \text{Value}(i).$$

The Greek letter \sum (“sigma”) is the mathematical notation for summation, e.g., $\sum_{y=1,2,3} y^2 = 1^2 + 2^2 + 3^2 = 14$.

Example: Fair and unfair bets

A **fair bet** is a bet with an expected value of zero. Flipping a coin and winning $\$x$ if it comes up heads and losing $\$x$ if it comes up tails is a fair bet. If $x = 5$, this is identical to the game above in which you pay me $\$5$ and then I pay you $\$10$ if the coin comes up heads and nothing if the coin comes up tails.

An **unfair bet** is one with an expected value less than zero. If you go to a casino and play roulette, for example, you will see that the roulette wheel has 38 numbers (the numbers 1–36, plus 0 and 00). If you bet $\$1$ and guess the right number, your payoff is $\$35$; if you guess the wrong number, your payoff is $-\$1$, i.e., you lose your dollar. So if you bet $\$1$ on, say, number 8, then the expected value of the amount you'll get back is

$$\Pr(8) \cdot (\$35) + \Pr(\text{Not } 8) \cdot (-\$1) = \frac{1}{38} \cdot (\$35) + \frac{37}{38} \cdot (-\$1) \approx -\$0.05.$$

4.1 Reducing risk with diversification

Playing roulette *once* is obviously a high-risk endeavor. (Presumably this is part of what makes it attractive to gamblers.) It would therefore seem that owning a casino would also be a high-risk endeavor. *But this is not the case.*

To see why, consider flipping a coin. Flip the coin once, and the actual percentage of heads will be either 0% or 100%. (In other words, the coin either comes up heads or it comes up tails.) Flip the coin twenty times, though, and the percentage of heads is likely to be pretty close to 50%. (Specifically, the

odds are better than 7 in 10 that the percentage of heads will be between 40% and 60%.) Flip the coin one hundred times and the odds are better than 95 in 100 that the percentage of heads will be between 40% and 60%. And flip the coin one thousand times and the odds are better than 998 in 1000 that the percentage of heads will be between 45% and 55%.

These results stem from the **law of large numbers**, a statistical result. In English, the law of large numbers says that if you flip a coin a large number of times, odds are that the proportion of heads will be close to 50%. (Details [online](#)¹ and [online](#)².) By extension, if each of a large number of coin flips pays \$10 for heads and \$0 for tails, odds are that you'll come close to averaging \$5 *per coin flip*. It is no coincidence that \$5 also happens to be the expected value of this bet: repeat any bet a large number of times and odds are that the *payout per bet* will be close to the expected value of the bet. This is the sense in which expected value can be thought of as the average payout of a bet.

Risk and roulette

To apply the law of large numbers to casinos, note that each gambler plays roulette only a few times, but that the casino plays roulette thousands of times each day. So the law of large numbers does not apply to individual gamblers, but it does apply to the casino. As a consequence, an individual gambler faces a great deal of risk, *but the casino does not*. Since the expected value from betting \$1 on roulette is $-\$0.05$, odds are extremely good that the casino will gain about \$0.05 for each dollar wagered. As long as the mob doesn't get involved, then, running a casino is not necessarily any riskier than running, say, a photocopy shop. (Go [online](#)³ for an amusing story that illustrates this point.)

Risk and the stock market

Another application of the law of large numbers is in the stock market, where investors are often advised to **diversify** their portfolios. Compared to owning one or two stocks, investors can reduce risk by owning many stocks. (One way to do this is to own shares of an **index fund** that buys a little bit of everything.)

To see how **diversification** can reduce risk, consider a coin flip that pays \$100 if it comes up heads (H) and \$0 if it comes up tails (T). This situation has an expected value of \$50, but the risk is clear: either you win the whole \$100 or you win nothing.

Now consider flipping two coins, each paying \$50 for heads and \$0 for tails. The expected value in this situation is also \$50:

$$\begin{aligned} EV &= \Pr(HH) \cdot (\$100) + \Pr(HT) \cdot (\$50) + \Pr(TH) \cdot (\$50) + \Pr(TT) \cdot (\$0) \\ &= \frac{1}{4} \cdot (\$100) + \frac{1}{4} \cdot (\$50) + \frac{1}{4} \cdot (\$50) + \frac{1}{4} \cdot (\$0) = \$50. \end{aligned}$$

¹<http://www.stat.berkeley.edu/~stark/Java/lln.htm>

²http://www.ruf.rice.edu/~lane/stat_sim/binom_demo.html

³<http://www.theonion.com/content/node/38825>

Compared to the one-coin flip, however, this two-coin flip is less risky: instead of always winning everything (\$100) or nothing (\$0), there is now a 50% chance of winning something in between (\$50). A risk-averse individual should prefer the two-coin flip over the one-coin flip because it has less variability. (In the language of statistics, the two-coin flip has lower **variance**. In English, the two-coin flip means that you aren't putting all your eggs in one basket.)

The risk inherent in coin-flipping (or stock-picking) can be reduced even further by spreading the risk out even more, i.e., through more diversification. Flip twenty coins, each independently paying \$5 for heads and \$0 for tails, and the probability of getting an outcome that is “in the middle” (say, between \$40 and \$60) is over 70%. Flip one hundred coins, each independently paying \$1 for heads and \$0 for tails, and the odds of ending up between \$40 and \$60 is over 95%. And flip one thousand coins, each independently paying \$0.10 for heads and \$0 for tails, and there is a 99.86% chance that you will end up with an amount between \$45 and \$55. The expected value in all of these situations is \$50; what diversification does is reduce variance, so that with one thousand coin flips you are virtually guaranteed to end up with about \$50.

These coin-flipping results follow from the law of large numbers. But there is an important difference between coin-flipping and stock-picking: while the outcome of one coin flip has no influence on the outcome of the next one, stocks have a tendency to go up or down together, i.e., coin flips are **independent** while stock prices are **correlated**.) The law of large numbers applies to risks that are independent, so diversification of a stock market portfolio can reduce or eliminate risks that are not correlated, e.g., the risk that a particular company will be ruined by a corrupt CEO. But the law of large numbers does not apply when risks are correlated: diversification cannot reduce the risk of stock market crashes or other systemic risks that affect the entire economy.

A final point is that diversification is not always painless. It's easy to diversify if you have no preference for one stock over another—and we will see in Chapter 5 that this is a reasonable position to take. But if you have favorites then you have to balance risks against rewards: investing your life savings in your favorite stock may give you the highest expected payoff, but it also exposes you to a great deal of risk that you could reduce by diversifying. The optimal behavior in such situations is the subject of **portfolio selection theory**, a branch of economics whose development helped win James Tobin the 1981 Nobel Prize in Economics. When Tobin died in 2002, his *New York Times* obituary (available [online](#)⁴) included this story:

After he won the Nobel Prize, reporters asked him to explain the portfolio theory. When he tried to do so, one journalist interrupted, “Oh, no, please explain it in lay language.” So he described the theory of diversification by saying: “You know, don't put your eggs in one basket.” Headline writers around the world the next day created some version of “Economist Wins Nobel for Saying, ‘Don't Put Eggs in One Basket’.”

⁴<http://bit.ly/2SoGs7>

4.2 *Algebra:* Indifference curves

A more mathematical treatment of this material allows us introduce **indifference curves**: we can ask when an individual is indifferent (i.e., has no preference) between one risky situation and another, or between a risky situation and a riskless one. For example, consider a coin-flipping bet in which someone wins \$100 for heads and \$0 for tails. Presumably this person would be indifferent between this bet and a symmetric bet in which he wins \$100 for tails and \$0 for heads. (If we refer to the first bet using the notation $(100, 0)$, then the notation $(0, 100)$ describes the second bet.) In fact, there are probably an infinite number of pairs (x, y) such that he'd be indifferent between the bet $(100, 0)$ and the bet (x, y) , i.e., a bet in which he wins $\$x$ for heads and $\$y$ for tails. If we graph all these pairs we get an indifference curve like the one in Figure 4.2. Using the utility-maximization idea described in Chapter 1, we can say that all the points on this indifference curve have the same utility.

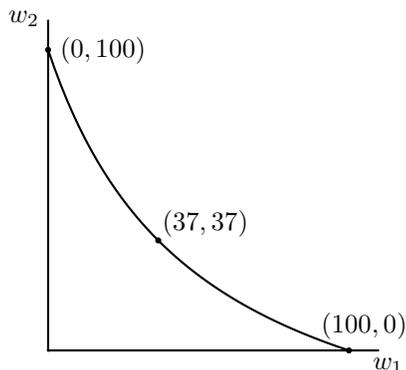


Figure 4.2: An indifference curve for a coin-flipping bet

Of particular interest on this indifference curve is the point where $x = y$, because at this point the individual is not exposed to any uncertainty: regardless of the outcome of the coin flip, he gets the same amount. This amount is called the **certainty equivalent wealth** corresponding to that indifference curve. For the curve in Figure 4.2, the certainty equivalent wealth is \$37, meaning that our individual is indifferent between getting \$37 *with certainty* and facing the uncertain bets—such as $(100, 0)$ and $(0, 100)$ —that are on the same indifference curve.

Note that the expected value of the bet $(100, 0)$ is \$50, which is *more* than the individual's certainty equivalent wealth. This implies that he is risk-averse; a risk-neutral person would have a certainty equivalent wealth of \$50 for the same original bet of $(100, 0)$. The difference between the certainty equivalent wealth of a risk-averse person (in our example, \$37) and that of a risk-neutral person (\$50) makes it possible for insurance companies to provide a benefit to society and still make a profit. If an insurance company offers to “buy” the bet $(100, 0)$ from our risk-neutral individual for, say, \$40, he will accept the offer

because he is indifferent between the bet $(100, 0)$ and having \$37 with certainty, and having \$40 with certainty is better than having \$37 with certainty. (In the language of economics, accepting the offer of \$40 will put him on a higher indifference curve, i.e., one corresponding to a higher utility level.) What's in it for the insurance company? Well, they spend only \$40 to buy a bet that has an expected value of \$50; by buying many such bets, the insurance company has an extremely good chance of making a profit.

Problems

Answers are in the next section. The first few words of each question (and each answer) are hyperlinked to allow online users to quickly move back and forth.

- 4.1 You roll a six-sided die and win that amount (minimum \$1, maximum \$6). What is the expected value of this game?
- 4.2 With probability $1/3$ you win \$99, with probability $2/3$ you lose \$33. What is the expected value of this game?
- 4.3 Imagine that you are taking a multiple-guess exam. There are five choices for each question; a correct answer is worth 1 point, and an incorrect answer is worth 0 points. You are on Problem #23, and it just so happens that Problem #23 is in Hungarian. (When you ask your teacher, she claims that the class learned Hungarian on Tuesday....)
 - (a) You missed class on Tuesday, so you don't understand any Hungarian. What is the expected value of guessing randomly on this problem?
 - (b) Now imagine that your teacher wants to discourage random guessing. To do this, she changes the scoring system, so that a blank answer is worth 0 points and an incorrect answer is worth x , e.g., $x = -\frac{1}{2}$. What should x be in order to make random guessing among five answers a fair bet (i.e., one with an expected value of 0)?
 - (c) Is this "fair bet" policy going to discourage test-takers who are risk-averse? What about those who are risk-loving?
 - (d) Your teacher ends up choosing $x = -\frac{1}{3}$, i.e., penalizing people $1/3$ rd of a point for marking an incorrect answer. How much Hungarian will you need to remember from your childhood in order to make guessing a better-than-fair bet? In other words, how many answers will you need to eliminate so that guessing among the remaining answers yields an expected value strictly greater than 0?
- 4.4 Two businesses that involve lots of gambling are the casino business and the insurance business. Are these businesses particularly risky to get involved in? Explain why or why not.

4.5 *Fun.* The Intergovernmental Panel on Climate Change reports that human activity (especially the burning of fossil fuels such as coal, oil, and natural gas) is warming the earth. (Note: With the exception of this fact, all of the numbers &etc in this question are completely made up.)

- (a) Assume that global warming will raise sea levels and increase the frequency of hurricanes, leading to damages of \$1 million ($= 10^6 = 1,000,000$) at the end of each year for the next seven years. What is the present value of that damage if the interest rate is 4%?
- (b) Next, assume that the full damages you've calculated above will only occur with probability 1/3. With probability 1/3 the damages will be only half as big, and with probability 1/3 the damages will be zero. What is the **expected value** of the damage caused by global warming? [Note: If you didn't answer part 4.5a above, just assume for this part that the total damage is \$1,000,000.]
- (c) Next, assume that the hurricanes &etc won't happen for 100 years. Using an interest rate of 4%, take the expected damages you calculated in part 4.5b and compute the present value of having that amount of damage occur 100 years in the future. [Note: If you didn't answer part 4.5b, just assume for this part that the total damage is \$1,000,000.]
- (d) What would be the present value of those damages if they won't occur for 500 years?

4.6 *Fun/Challenge.* (The Monty Hall Problem) This problem gets its name from the TV game show *Let's Make A Deal*, hosted by Monty Hall. The scenario is this: Monty shows you three closed doors. Behind one of these doors is a new car. Behind the other two doors are goats (or some other "non-prize"). Monty asks you to choose a door, but after you do he doesn't show you what's behind the door you chose. Instead, he opens one of the *other* doors, revealing a goat, and then offers you the opportunity to switch to the remaining unopened door. [As an example, say you originally pick Door #1. Monty opens up Door #2, revealing a goat, and then offers you the opportunity to switch from Door #1 to Door #3.] (You can go [online](http://en.wikipedia.org/wiki/Monty_Hall_problem)⁵ for a fun computer simulation of the Monty Hall problem, with accompanying discussion.) What should you do?

4.7 Howard Raiffa's book *Decision Analysis: Introductory Lectures on Choices Under Uncertainty* (1997) does an extensive analysis of variations on the following basic problem.

There are 1,000 urns. Eight hundred of them are of type U_1 ; each of these contain four red balls and six black balls. The remaining two hundred are of type U_2 ; each of these contain nine red balls and one black ball. One of these 1,000 urns is chosen at random and placed in front of you; you

⁵http://en.wikipedia.org/wiki/Monty_Hall_problem

cannot identify its type or see the balls inside it. Which one of the following options maximizes your expected value, and what is that expected value?

Option 1 Guess that the urn is of type U_1 . If you are correct, you win \$40.00. Otherwise, you lose \$20.00.

Option 2 Guess that the urn is of type U_2 . If you are correct, you win \$100.00. Otherwise, you lose \$5.00.

Option 3 Refuse to play the game.

- 4.8 *Challenge. (One variation)* Prior to choosing one of the three options described above, you can conduct *at most one* of the following investigations. (Note that you can also choose not to conduct any of these.) What strategy maximized your expected value, and what is that expected value?

Investigation 1 For a payment of \$8.00, you can draw a single ball at random from the urn.

Investigation 2 For a payment of \$12.00, you can draw two balls from the urn.

Investigation 3 For a payment of \$9.00, you can draw a single ball from the urn, and then (after looking at it) decide whether or not you want to pay \$4.50 to draw another ball. (Whether or not you want to replace the first ball before drawing the second is up to you.)

- 4.9 You're a bidder in an auction for an antique vase. If you lose the auction, you get nothing. If you win the auction, assume that your gain is the difference between the maximum amount you'd be willing to pay for the vase—say, \$100—and the actual amount that you end up paying. (So if you pay \$80, your gain is \$20.)

- (a) In a **first-price sealed bid auction**, you write down your bid b on a piece of paper and submit it to the auctioneer in a sealed envelope. After all the bids have been submitted, the auctioneer opens the envelopes and finds the highest bidder. That bidder gets the item, and pays a price equal to their bid. If the probability of winning with a bid of b is $\Pr(b)$, write down an expected value calculation for this auction.
- (b) In a **second-price sealed bid auction**, everything's the same except that the winning bidder (the person with the highest bid) pays a price equal to the *second-highest* bid. Write down an expected value calculation for this auction if the probability of winning with a bid of b is $\Pr(b)$ and the highest bid *less* than b is c .
- (c) *Challenge.* From your answers above, can you figure out what kind of strategy (i.e., what bid b) will maximize your expected value in the different auctions? In particular: should you bid your true value, $b = \$100$, or should you bid more or less than your true value? (We'll study auctions more in Chapter 9.)

- 4.10 Although many services are provided on a fee-up-front basis (for example, you pay before you enter the movie theater), some services are provided on the “honor system”. Parking meters are one example: instead of paying someone when you park your car, you are “on your honor” to put money in the meter. (Some cities also do this with their public transportation systems.) Of course, it isn’t just a matter of honor: there are also enforcement officers—“meter maids”—who show up from time to time and penalize rule-breakers. So:
- (a) What is the expected cost of “risking it” by putting nothing in the meter if there’s an 80% chance that you’ll get away with it (so your cost will be \$0) and a 20% chance that you’ll get a \$20 ticket?
 - (b) Imagine that the city managers want to save money by cutting in half the number of enforcement officers (so that the chance of getting a ticket is only 10%). Can you suggest a way to do this without increasing the attractiveness of cheating?
- 4.11 Choosing a president is undoubtedly a more important decision than (say) choosing what kind of car you’re going to buy. But many people spend hours deciding what kind of car to buy and only minutes deciding which presidential candidate to vote for. This problem tries to explain why.
- (a) Activity #1 pays \$1 million with probability 0.0001 and \$0 with probability 0.9999. Activity #2 pays \$1000 with probability 1. Calculate the expected value of each of these activities to show that Activity #2 has a higher expected value.
 - (b) Use the analysis above to explain why many people spend hours deciding what kind of car to buy and only minutes deciding which presidential candidate to vote for.
- 4.12 Imagine that you are a profit-maximizing forester. You currently own trees containing 100 board-feet of timber.
- (a) With probability 2%, a fire will destroy your trees, and you’ll have no harvestable timber. With probability 98%, your trees will grow and in one year you’ll have 5% more board-feet of timber. What is the expected number of board-feet of timber you’ll have next year?
 - (b) Explain (as if to a non-economist) why the interest rate at the bank matters in deciding to cut the trees down now or to cut them down in year.
 - (c) Continuing with the story from part (a) above, assume that the price of lumber is constant over time and that you’re a risk-neutral forester. In order for cutting the trees down next year to be a better choice than cutting the trees down now, the interest rate at the bank has to be **(circle one: higher lower)** than _____%.

- 4.13 Social Security benefits are adjusted for inflation, meaning that payments to retirees increase at the rate of inflation.
- How much will my grandmother (call her Grammy) be paid in one year, in two years, and in three years if her current benefit is \$1000 and inflation is 3%?
 - One way to calculate the present value of these three payments is to use brute force: determine the present value of each payment separately and then add them together. Go ahead and do this when the nominal interest rate is 5%.
 - Calculate the present value of receiving \$1000 at the end of each year for 3 years if the relevant interest rate is 2%. Compare with your answer from (b) and explain.

Answers

- 4.1 The expected value is $\frac{1}{6}(1) + \frac{1}{6}(2) + \frac{1}{6}(3) + \frac{1}{6}(4) + \frac{1}{6}(5) + \frac{1}{6}(6) = \frac{21}{6}$.
- 4.2 The expected value is $\frac{1}{3}(99) + \frac{2}{3}(-33) = 11$.
- 4.3 [“Imagine that you are taking...”]
- The expected value of guessing randomly is $\frac{1}{5}(1) + \frac{4}{5}(0) = \frac{1}{5}$.
 - If an incorrect answer is worth x , the expected value from guessing randomly is $\frac{1}{5}(1) + \frac{4}{5}(x) = \frac{1+4x}{5}$. If the teacher wants this expected value to equal zero, she must set $x = -\frac{1}{4}$.
 - Since this makes random guessing a fair bet, it will discourage risk averse students but not risk loving students.
 - If you can't eliminate any answers, the expected value of guessing randomly is $\frac{1}{5}(1) + \frac{4}{5}(-\frac{1}{3}) = -\frac{1}{15}$. If you can eliminate one answer, you have a 1 in 4 chance of getting the right answer if you guess randomly, so your expected value if you can eliminate one answer is $\frac{1}{4}(1) + \frac{3}{4}(-\frac{1}{3}) = 0$. If you can eliminate two answers, you have a 1 in 3 chance of getting the right answer if you guess randomly, so your expected value if you can eliminate two answers is $\frac{1}{3}(1) + \frac{2}{3}(-\frac{1}{3}) = \frac{1}{9}$. So you need to eliminate at least two answers in order to make random guessing yield an expected value greater than zero.
- 4.4 No, they are not particularly risky. This is because of the **law of large numbers**, discussed in Section 4.1. The individual bettor plays roulette only a few times, and so faces a lot of risk. The casino plays roulette thousands of times each day, and so has a very good idea of what the overall outcome will be; since each \$1 wager has an expected payoff of only \$.95, it can expect to gain about \$.05 for every dollar wagered.

Similarly, although insurance companies have no idea whether an individual driver is going to get into an accident this year, or whether an individual person is going to die this year, or whether an individual home is going to burn to the ground this year, the law of large numbers usually gives them a very good idea of the percentage of accidents or the percentage of deaths or the percentage of fires to expect from the hundreds of thousands of cars, lives, and homes they cover.

So casinos or insurance companies are not necessarily any riskier as business endeavors than, say, running a photocopy shop.

4.5 [“The Intergovernmental Panel on Climate Change...”]

- (a) Using the annuity formula we get a present value of about \$6 million.
- (b) The expected damages are $\frac{1}{3}(6) + \frac{1}{3}(3) + \frac{1}{3}(0) \approx \3 million.
- (c) Plug \$3 million into the present value of a lump sum formula to get a present value of \$59,400.
- (d) Using the present value of a lump sum formula, we get \$0.009, or about one penny.

4.6 **You should switch:** your odds of winning will increase from $\frac{1}{3}$ to $\frac{2}{3}$. A more extreme example may help provide some intuition behind this result: assume that there are 100 doors, only one of which leads to a car; after you pick a door, Monty opens up 98 of the other doors to reveal goats and then offers you the opportunity to switch to the remaining unopened door. Doing so will increase your odds of winning from $\frac{1}{100}$ to $\frac{99}{100}$.

4.7 **Your expected value** from Option 1 is $.8(40) + .2(-20) = 28$. Your expected value from Option 2 is $.8(-5) + .2(100) = 16$. Your expected value from Option 3 is 0. So Option 1 maximizes your expected value.

4.8 **This is a difficult problem.** For more information on it, read Raiffa’s book or do some research on <http://www.cs.ubc.ca/~murphyk/Bayes/bayesrule.html> Bayes’s Rule.

4.9 **Your expected value** from bidding b in either type of auction is

$$\text{Prob}(b \text{ wins}) \cdot \text{Value}(b \text{ wins}) + \text{Prob}(b \text{ loses}) \cdot \text{Value}(b \text{ loses}).$$

- (a) In a first-price auction, $\text{Value}(b \text{ wins}) = 100 - b$ and $\text{Value}(b \text{ loses}) = 0$; so your expected value is

$$\text{Prob}(b \text{ wins}) \cdot (100 - b) + 0.$$

- (b) In a second-price auction, $\text{Value}(b \text{ wins}) = 100 - c$, where c is the highest bid less than b , and $\text{Value}(b \text{ loses}) = 0$. So your expected value is

$$\text{Prob}(b \text{ wins}) \cdot (100 - c) + 0.$$

(c) Chapter 9 discusses this in more detail.

4.10 [“Although many services...”]

- (a) The expected cost is $(.80)(0) + (.20)(\$20) = \4 .
- (b) The city could double the amount of the ticket from \$20 to \$40. This would mean that the expected value of risking it is still \$4: $(.90)(0) + (.10)(\$40) = \4 .

4.11 [“Choosing a president...”]

- (a) The expected values are $(.0001)(\$1000000) + (.9999)(\$0) = \$100$ and $(1)(\$1000) = \1000 .
- (b) There is a very low probability that your vote will influence the outcome of the election, so the expected value from thinking about whom to vote for can be low even though the overall outcome of the election itself is important.

4.12 [“Imagine that you are...”]

- (a) We have $(0.1)(0) + (.98)(105) = 102.9$.
- (b) To maximize your present value you need to compare the return you’ll get from “investing in the trees” with the return you’ll get from investing in the bank. Investing in the bank means cutting down the trees and putting the proceeds in the bank. Investing in the trees means letting the trees grow so there will be more lumber next year.
- (c) The interest rate at the bank has to be lower than 2.9%.

4.13 [“Social Security benefits...”]

- (a) Plug the inflation rate (.03) into the future value formula to get a payment of \$1030 in one year, \$1060.90 in two years, and 1092.73 in three years.
- (b) Use the nominal interest rate and the lump sum formula to get a present value of approximately

$$\$980.95 + \$962.27 + \$943.94 = \$2887.16.$$

- (c) Plug \$1000, .02, and 3 years into the annuity formula to get a present value of \$2883.88. This is very close to the answer from above! (If you use 1.94% as a better estimate of the real interest rate, the annuity formula gives you a present value of \$2887.26, which is very close indeed to the \$2887.16 figure above.) The punch line here is that you can use the real interest rate to determine the present value of inflation-adjusted annuities.

Chapter 5

From one to some

An economics professor and an economics student meet on their way to class. The student spies a \$10 bill on the ground and reaches down to pick it up. The professor says, “Don’t bother. If it really was a \$10 bill, somebody else would have already picked it up.”

This chapter connects the first part of this book (individual optimization) and the second part (strategic interactions between optimizing individuals) by looking at **arbitrage**, which investorwords.com defines as “[an attempt] to profit by exploiting price differences of identical or similar financial instruments, on different markets or in different forms. The ideal version is **riskless arbitrage**.”

An example of riskless arbitrage is the simultaneous purchase and sale of some asset at different prices. If, for example, Microsoft stock is trading on the New York Stock Exchange (NYSE) for \$100 and on the NASDAQ stock exchange for \$105, an individual trader can engage in arbitrage by simultaneously “buying low” and “selling high”: purchasing Microsoft stock on the NYSE and selling it on NASDAQ. The net result is a riskless profit of \$5 per share.

More generally, the concept of arbitrage applies to assets that are *similar* rather than *identical*. It can also be used in contexts that don’t involve financial instruments: a driver in rush-hour traffic who switches from a slow-moving lane into a faster-moving one might reasonably be said to be engaging in arbitrage.

5.1 No arbitrage

Like the preceding chapters in this book, the concept of arbitrage is centered on the optimizing individual; it is the individual stockbroker, for example, who engages in arbitrage by simultaneously buying and selling some asset. The upcoming chapters in this book look at interactions *between* optimizing individuals. One of the key questions will be: “What does a world full of optimizing individuals look like?”

It will take the rest of this chapter to develop a completely correct answer to that question in the context of arbitrage. There is, however, a *nearly* correct answer that is remarkably simple: *in a world full of optimizing individuals, there would be no opportunities for arbitrage*. This “no arbitrage” result predicts that Microsoft stock will trade at the same price on the New York Stock Exchange and on NASDAQ. It also predicts that switching lanes during rush hour will be an exercise in futility because all the lanes will be moving at the same speed.

These predictions help explain the title (*Hidden Order*) of David D. Friedman’s 1996 book about economics. Especially interesting is the fact that these social patterns emerge without conscious effort: stockbrokers don’t intend to equate the price of Microsoft stock on different exchanges, and nobody is “in charge” of different lanes of rush-hour traffic. To the contrary, each stockbroker is only interested in making money, and each driver is only interested in getting home. Despite this, the aggregate result of these myopic activities is not chaos but order. Finding such “hidden order” is one of the pleasures of economics.

5.2 Rush-hour arbitrage

A *completely* correct statement about rush-hour traffic is that we should *expect* different lanes to travel at *approximately* the same speed. The reason is that **rush-hour arbitrage is self-eliminating**. If one lane does happen to be moving a little bit faster, some drivers—the “explorers”—will change lanes, thereby gaining a few seconds. But all those explorers shifting into the fast-moving lane will slow it down, and all those explorers shifting out of slow-moving lanes will speed them up. The end result is that the explorers—in their efforts to gain precious seconds—act to equalize the traffic speeds on the different lanes.

A curious result here is that the existence of explorers can help drivers who are not explorers. These other drivers—the “sheep”—prefer to just pick a lane and stick with it. In the absence of explorers, this behavior would be risky: the other lane really might be moving faster. But the presence of explorers—and the fact that arbitrage is self-eliminating—means that the sheep can relax, comfortable in the knowledge that the lane they choose is expected to move at approximately the same speed as any other lane.

Having examined the need for the “approximately” in the statement that we should *expect* different lanes to travel at *approximately* the same speed—explorers *can* get ahead, but only by a little—we now turn our attention to the “expect”. This word refers to expected value from Chapter 4, meaning that the focus is on what happens *on average* rather than what happens *on any particular occasion*. For example, if you have an even bet (flip a coin, win \$1 if it comes up heads, lose \$1 if it comes up tails), the **expected value** of that bet is \$0. In hindsight—i.e., after the coin is flipped—it is easy to say whether or not you should have taken the bet. Before the fact, though, all that you can say is that *on average* you will neither gain nor lose by taking the bet.

Returning to rush hour, imagine that two identical freeways connect two cities. Despite the existence of explorers listening to radio traffic reports in the

hopes of finding an arbitrage opportunity, we cannot say that the travel times on the two freeways will *always* be approximately equal: if there is an accident on one freeway while you're on it, you may get home hours later than the lucky drivers on the other freeway. What we can say is that the *expected* travel times on the two freeways will be approximately equal. In hindsight, of course, it is easy to say that you should have taken the other freeway. But before the fact all that can be said is that *on average* the freeways will have equal travel times.

5.3 Financial arbitrage

A nice analogy connects traffic decisions and financial decisions: think of the different lanes of the freeway as different stocks or other investment options, and think of travel speed as measuring the rate of return of the different investments, so that a fast-moving lane corresponds to a rapidly increasing stock price. In both cases optimizing individuals are looking for the fastest lane, so it is not surprising that the economic analyses of these two phenomena are similar. A *completely* correct statement about financial arbitrage is that we should *expect comparable investments to have comparable rates of return*.

The word “comparable” is necessary because there is a lot more variety to be found in financial investments than in the lanes of traffic on the freeway. Before we discuss this in detail, note that—as with rush-hour arbitrage—economic reasoning about financial arbitrage focuses on *expectations*. Just because different stock prices have the same *expected* rate of return doesn't mean that they will have the same *actual* rate of return. If you look at the prices of different stocks over time, you'll find winners and losers, but using hindsight to look back and say that you should have bought or sold XYZ stock five years ago is like saying that you should have picked the winning lottery numbers prior to the drawing.

Another parallel with rush-hour arbitrage is that **financial arbitrage is self-eliminating**. The “explorers” in stock markets are investment managers and others who spend a lot of time and effort trying to find stocks that are likely to go up in value faster than others. By driving up the stock price of companies that are expected to do well—and driving down the stock price of companies that are expected to perform poorly—these explorers equalize the attractiveness of “comparable” stocks. And, as in the case of rush-hour arbitrage, the presence of financial “explorers” allows other investors—the “sheep”—to just pick some stocks and stick with them. This is a good thing because many people can't or don't want to spend all their time researching stocks.

The fact that it takes time and effort to look for financial arbitrage opportunities is important because it allows us to paint a complete picture of how the whole process works. The *nearly* correct “no arbitrage” reasoning at the beginning of this chapter suggests that the reason there are no arbitrage opportunities is that someone would have taken advantage of them if there were. But if there are no arbitrage opportunities, who would bother to look for them? And if nobody bothers to look for them, arbitrage opportunities might well exist, as

in the joke at the beginning of this chapter.

A complete picture is that there *are* small discrepancies that occasionally appear between the price of Microsoft stock on the NYSE and on NASDAQ; and there *are* people who profit from such discrepancies. But it is not right to think of these people as getting “free money” any more than it is right to think of other workers as getting free money when they get paid for a day’s work. Indeed, looking for arbitrage opportunities can be thought of as a type of investment. As such, we have good reason to think that the expected rate of return from this investment should be comparable to that from comparable investments such as working elsewhere on Wall Street.

What do we mean by “comparable investments”?

For one thing, we mean investments that have similar **liquidity**: if you lock up your money in a 5-year Certificate of Deposit (CD) that carries penalties for early withdrawal, you can expect a higher rate of return than if you put your money in a savings account that you can liquidate at any time.

Another issue is **risk**: many people are risk-averse, so we should expect high-risk investments to have higher expected rates of return than low-risk investments. For example, junk bonds sold by a company threatened with bankruptcy should have higher expected interest rates than U.S. government bonds because they involve more uncertainty: if the company selling those bonds goes bankrupt, the bonds will be worthless. (See Problem 5.2.) If the junk bonds and U.S. government bonds have expected returns of 6% and 4%, respectively, then the **risk premium** associated with the junk bonds is the difference between the expected rates of return: $6\% - 4\% = 2\%$. To compensate investors for the added risk of lending money to a troubled company, the company must offer a higher *expected* rate of return.

Curiously, the risk premium associated with stocks seems to be undeservedly large. Although stocks go up and down, over time they mostly seem to go up, and they seem to go up quite fast: investing in the stock market has been a much better investment than investing in bonds or in the Bank of America. Consider this information from the *Schwab Investor*: If you invested \$2,000 each year from 1981-2000 in U.S. Treasury bonds, you would have ended up with \$75,000; if you invested that same amount in the stock market, you would have ended up with \$282,000. (Even an unlucky investor, one who invested at the worst time each year, would have ended up with \$255,000.) In order to make bonds or other investments an appealing alternative to stocks, the risk premium associated with stocks would have to be quite large. This whole issue is something of a mystery to economists, who call it the **equity premium puzzle**. Figure it out and you’ll win the Nobel Prize in Economics.¹

¹Harvard economist Martin Weitzman may have beaten you to it. His idea is that the risk premium is due to small possibilities of very bad returns in the stock market; the fact that we haven’t seen any in the past could just be because stock markets haven’t been around very long. See [Weitzman 2007](#) (“Subjective Expectations and Asset-Return Puzzles”, *American Economic Review* 97: 1102–30).

Problems

Answers are in the next section. The first few words of each question (and each answer) are hyperlinked to allow online users to quickly move back and forth.

- 5.1 Explain (as if to a non-economist) why comparable investments should have comparable expected rates of return.
- 5.2 Consider a company that has a 10% chance of going bankrupt in the next year. To raise money, the company issues junk bonds paying 20%: if you lend the company \$100 and they *don't* go bankrupt, in one year you'll get back \$120. Of course, if the company *does* go bankrupt, you get nothing. What is the *expected* rate of return for this investment? If government bonds have a 3% expected rate of return, what is the **risk premium** associated with this investment?
- 5.3 Economic reasoning indicates that comparable investments should have comparable expected rates of return. Do the following examples contradict this theory? Why or why not?
- (a) Microsoft stock has had a much higher rate of return over the last twenty years than United Airlines stock.
 - (b) Oil prices are not going up at the rate of interest.
 - (c) Junk bonds pay 20% interest while U.S. Treasury bonds pay only 4% interest.
 - (d) Banks in Turkey (even banks that have no risk of going bankrupt) pay 59% interest while banks in the U.S. only pay 4% interest.
- 5.4 “Opportunities for arbitrage are self-eliminating.”
- (a) Explain this statement in the context of lanes of traffic on a congested freeway.
 - (b) Many economists believe that a similar logic holds when it comes to attempts by “active investors” to pick stock-market winners. Describe the investment advice from these economists, or otherwise explain.

Answers

- 5.1 This is answered in the text.
- 5.2 If you lend the company \$100, the expected value of your investment is $.90(120) + .10(0) = \$108$, meaning that your expected rate of return is 8%. If government bonds have a 3% expected rate of return, the risk premium associated with the junk bonds is $8\% - 3\% = 5\%$.

5.3 [“Economic reasoning indicates...”]

(a) The *actual* rates of return turned out to be different, but it could still be true that at any point in time over the last twenty years the *expected* rates of return were the same. Economic reasoning suggests that they should have been; otherwise investors would not have been acting optimally. In terms of its effect on this logic, the fact that the actual rates of return turned out to be different is no more cause for concern than the fact that some lottery tickets turn out to be winners and some turn out to be losers.

(b) Again, actual rates of return can be different even though expected rates of return are the same at any moment in time. For example, there may be unexpected advances (or setbacks) in the development of electric cars or other alternatives to gasoline-powered cars.

An alternative explanation is that the cost of extracting oil has gone down over time; as a result, the profits from each barrel of oil may be going up at the rate of interest even though the price of oil may not be steady or falling.

(c) Companies that issue junk bonds are financially troubled; if they go bankrupt, you lose both the interest and the principal. So even though junk bonds may offer 20% interest, their expected rate of return is much less than 20%, and therefore much closer to the expected rate of return of U.S. Treasury bonds. Also, the U.S. Treasury isn't likely to go bankrupt, meaning that two assets don't have comparable risk. So it is likely that there is a risk premium associated with the junk bonds.

(d) The answer here is that the rate of inflation is much higher in Turkey than in the U.S. So even though the two banks pay different nominal interest rates, their expected real interest rates may be equal.

One way to see the effect of inflation is to imagine that you'd invested 100 U.S. dollars in a Turkish bank in June 2000. The exchange rate then was about 610,000 Turkish lira to \$1, so your \$100 would have gotten you 61 million Turkish lira. After one year at a 59% interest rate, you would have 97 million lira. But when you try to change that back into U.S. dollars, you find that the exchange rate has changed: in June 2001 it's 1,240,000 lira to \$1, so your 97 million lira buys only \$78.23. You actually would have lost money on this investment!

5.4 [“Opportunities for arbitrage...”]

(a) Drivers who move out of slow lanes and into fast lanes slow down the fast lanes and speed up the slow lanes, thereby equalizing traffic speed in the different lanes.

(b) These economists argue that “you can't beat the market” (formally, this is called the efficient market hypothesis) and therefore that you

should invest in an index fund that buys a little bit of everything. The idea here is that this sort of passively managed fund will have lower management costs than an actively managed fund that tries to beat the market.

Part II

Strategic interactions

Chapter 6

Cake cutting

A recently divorced woman finds a washed-up bottle on the beach, and when she rubs it a Genie comes out and says, “You have three wishes, but I must warn you: whatever you wish for, your ex-husband will get ten times as much.” The woman thinks for a while and then says, “First, I’d like to be beautiful.” “Fine,” the Genie replies, “but remember—your ex-husband will be ten times more beautiful.” “That’s okay, I’ll still be gorgeous.” “Very well,” the Genie says, and makes it so. “And your second wish?” “I’d like \$10 million.” “Remember,” says the Genie, “your ex-husband will get \$100 million.” “That’s fine,” the woman replies, “I’ll still be able to buy whatever I want.” “Alright,” the Genie replies. “And your third wish?” “I’d like to have a mild heart attack.”

Having studied how optimizing individuals *act*, we now begin the study of how they *interact*. The branch of economics that studies strategic interactions between individuals is called **game theory**.

Here are some of the key issues to consider with strategic interactions:

The players Who’s involved and what are their motivations? (In general we’ll assume that the players care only about themselves, but it is possible to include other-regarding preferences such as altruism.)

The strategies What options does each player have? These strategies often depend on other players’ actions: “If they do X, then I can do Y or Z. . . .”

The payoffs What are the players’ outcomes from the various combinations of strategies? (One early focus of game theory was **zero-sum games**, where a “good” outcome for one player means a “bad” outcome for another.)

The timeline Who gets to act when? (In chess, for example, the player who gets to play first is said to have an edge—a **first-mover advantage**.)

The information structure Who knows what and when do they know it?

6.1 Some applications of game theory

One obvious application of game theory is to games. Here, for instance, is the introduction to Ken Binmore’s analysis of poker in his 1991 book *Fun and Games*:

Poker is a card game for which one needs about seven players for the variants normally played to be entertaining. Neighborhood games are often played for nickels and dimes, but there is really no point in playing Poker except for sums of money that it would really hurt to lose. Otherwise, it is impossible to bluff effectively, and bluffing is what Poker is all about.

Information is also key in poker: I don’t know what cards you have, and you don’t know what cards I have, and I know that you don’t know what cards I have, and you know that I know that you don’t know what cards I have. . . .

Poker is of course not the only interesting game of strategy. Others include:

Evolutionary games Why are some species aggressive and others evasive?

Why do species like elephant seals—where a few males mate with all the females and many males don’t mate at all—have equal numbers of male and female offspring? These are topics from **evolutionary game theory**, one of the most successful applications of game theory.

Auctions How much should you bid in an auction? If you’re auctioning something off, what kind of auction should you use? (More in Chapter 9.)

Sports Should you go for the two-point conversion or the one-point conversion? Should you give in to a player’s (or an owner’s) contract demands or hold out for a better deal?

Business How should you go about invading another company’s turf, or stopping another company from invading your turf?

War The development of game theory in the 20th century slightly preceded the Cold War, and the ideas of game theory helped make John von Neumann, one of its pioneers, into a proponent of “preventative war”, i.e., nuking the Soviets before they could develop the atomic bomb and reply in kind. (He was fond of saying, “If you say why not bomb them tomorrow, I say why not today? If you say today at five o’clock, I say why not one o’clock?”) Game theory can help explain key features of the Cold War, such as the 1972 ABM Treaty in which the U.S. and the U.S.S.R. agreed not to pursue *defenses* against missile attacks. (In addition to perhaps playing a role in averting nuclear catastrophe, the underlying doctrine of Mutually Assured Destruction has one of the world’s best acronyms.)

Fair division The problem that we are going to focus on in this chapter is a deceptively silly one: that of fairly dividing cake. Note, however, that

issues of fair division arise in many important circumstances—for example, in divorce or estate settlements, where the property to be divided can be thought of as the “cake”—and that cake and pie analogies are widely used to discuss resource allocation problems, from George W. Bush’s campaign promise in 2000 to “make the pie higher” to this quote from John Maynard Keynes’s 1920 essay *The Economic Consequences of the Peace*:

[The economic system in Europe before World War I] depended for its growth on a double bluff or deception. On the one hand the laboring classes accepted... a situation in which they could call their own very little of the cake that they and nature and the capitalists were co-operating to produce. And on the other hand the capitalist classes were allowed to call the best part of the cake theirs and were theoretically free to consume it, on the tacit underlying condition that they consumed very little of it in practice. The duty of “saving” became nine-tenths of virtue and the growth of the cake the object of true religion.

6.2 Cake-cutting: The problem of fair division

The basic cake-cutting problem is for “Mom” to figure out a way to fairly divide a cake among two or more kids. If you’d prefer a more real-world example, think of a divorce settlement or the settling of an estate of somebody who died without a will; in these cases, the couple’s assets, or the items in the estate, are comparable to the cake that must be divided, and the judge is “Mom”.

Three aspects of the cake-cutting problem are of particular interest. First, the cake is not necessarily homogenous: there may be parts with lots of frosting and parts without frosting, parts with lots of chocolate and parts without chocolate, etc. (In a divorce or estate settlement, the different parts of the cake correspond to the different items to be divided: cars, houses, children, etc.) Second, the kids dividing the cake may have different values: one may love chocolate, another may hate chocolate. Finally, these values may not be known to Mom (or to the judge dividing the assets). Even if they were, what parent wants to be in the middle of a cake-cutting argument? So what Mom is looking for is a mechanism through which the kids can divide up the cake fairly and without rancor. In other words, Mom’s problem is one of **mechanism design**.¹

Motivating question: How can you get two kids to divide a piece of cake so that each kid gets—in his or her opinion—at least 50% of the cake?

One answer: There’s a time-honored procedure called Divide and Choose: one kid cuts the cake into two pieces, and the other kid chooses between them. Barring mistakes, each kid is guaranteed at least 50% of the cake.

¹The 2007 Nobel Prize in Economics went to three of the founders of mechanism design: Leonid Hurwicz, Eric S. Maskin, and Roger B. Myerson.

Another answer is the Moving Knife procedure: Pass a long knife from left to right over the cake. When one kid calls out “Stop”, cut the cake at that point and give the slice to the left of the knife to the kid that called out. The remainder of the cake goes to the kid who was silent.

Question: Imagine you’re one of the kids participating in Divide and Choose. Do you want to be the cutter or the chooser?

Answer: It might depend on how well you can handle a knife and how fairly you can eyeball the cake, but on a deeper level it depends on how much you know about the other child’s preferences, and how much they know about your preferences. . . and how much you know they know about your preferences!

If you know a lot about what the other kid likes, you can take advantage of this information by choosing to cut the cake. You can then cut it in such a way that your sibling thinks the cake is split 51–49 but you think the cake is cut 40–60; if your sibling is optimizing, she’ll take the first piece—which she values as 51% of the cake—and you can make off with what is (in your valuation) 60% of the cake. In the divorce analogy, if you know your soon-to-be-ex-spouse will do anything to keep the 1964 Mustang convertible, you should do the cutting: put the convertible in one pile and everything else in the other pile!

On the other hand, if you don’t know a lot about your sibling’s preferences, you might want to be the chooser. As long as your sibling doesn’t know a lot about your preferences either, she’ll have to cut the cake so that it’s split 50–50 according to her valuation. If your valuation is different, you’ll be able to choose a piece that you value at more than 50%.

Question: What if Mom has 3 or more kids? And what if the object that you are dividing up is not a *good* but a *bad*, e.g., chores rather than cake?

Answer: It’s possible to extend the Divide and Choose and Moving Knife procedures to 3 or more kids. There are also a few algorithms for what is called dirty-work fair division. You can read more about them in a neat book called *Cake-Cutting Algorithms: Be Fair If You Can*.²

Question: How do you divide a cake between 3 kids in an **envy-free** manner, i.e., so that no kid covets another kid’s piece?

Answer: This problem went unsolved for many decades, but a variety of moderately complicated algorithms are now known. Figure 6.1 shows one.

What does “fair” mean?

The issue of envy-free division leads us to an important question: Does a division have to be envy-free to be fair? More broadly, what do we mean by “fair”?

²Jack Robertson and William Webb, *Cake-Cutting Algorithms: Be Fair If You Can* (1998). Some parts of this book are readable but the mathematics quickly gets pretty intense.

<p>Envy-Free Algorithm for Three Players (Attributed to John Conway, Richard Guy, and/or John Selfridge)</p> <p>Step 1 Have Player A cut the cake into three pieces.</p> <p>Step 2 Have Player B rank the pieces from best to worst—X_1, X_2, X_3—and then trim off a slice $E \geq 0$ from X_1 so that $X_1' = X_1 - E$ and X_2 have the same value for Player B.</p> <p>Step 3 Have Player C choose from among X_1', X_2, and X_3.</p> <p>Step 4 If C did not choose X_1', give X_1' to B and the remaining piece to A. If C did choose X_1', have B choose between X_2 and X_3; again, give the remaining piece to A.</p> <p>Step 5 Either B or C received X_1'. Call that person P_1 and the other person P_2.</p> <p>Step 6 Have P_2 cut E into three pieces, and have those pieces chosen in the order P_1, A, P_2.</p>
--

Figure 6.1: An envy-free algorithm for three players, adapted from Section 1.4 of *Cake-Cutting Algorithms*. (No this is not fair game for exams!)

In their books *The Win-Win Solution* and *Fair Division*,³ Steven Brams and Alan Taylor identify the following criteria for fairness in division problems:

Proportionality If, for example, there are 3 kids, then each kid should get at least 1/3rd of the cake (according to his or her estimation).

Envy-free No kid should covet another kid's share.

Equitability “Equitability [means that] both parties think they received the same fraction of the total, as each of them values the different items. For example, equitability would hold if the husband believed he got 70% and the wife believed she got 70%, which is quite a different story from his believing that he got 51% and her believing that she got 90%.”

Pareto efficiency There is no other division of the cake that would make at least one kid better off and not make any kid worse off.

Brams and Taylor then go on to describe their (patented!) Adjusted Winner algorithm for achieving these goals and to apply this algorithm to real-world situations such as the Israeli-Palestinian conflict.

³Brams, Steven J. and Alan D. Taylor, *The Win-Win Solution: Guaranteeing Fair Shares to Everybody* (New York: W.W. Norton, 1999) is a very easy-to-read popular guide by two game theorists, one a political scientist, the other a mathematician. A theoretical treatment by the same folks is *Fair Division: From Cake-Cutting to Dispute Resolution* (Cambridge University Press, 1996).

6.3 The importance of trade

Observation of sandwich-swapping in any elementary school lunchroom will confirm that just because “Mom” divides the cake in a certain way doesn’t mean that things will stay that way.

The possibility of trade is of fundamental importance in economics. The reason is the close relationship between trade and the fourth criterion on Brams and Taylor’s list: Pareto efficiency, which specifies that there is no other allocation of resources that would make at least one kid better off without making anyone else worse off. If this criterion is *not* met, then it *is* possible to reallocate resources in such a way that at least one kid would be better off and nobody would be worse off. In such **Pareto inefficient** situations, at least one kid would have a strong incentive to push for such a reallocation, e.g., by trading with or bribing the other kids to accept that alternative allocation.

To put it more simply: there is an inherent tension between trade and Pareto inefficient situations. Because of this antagonism, we should expect the possibility of trade to naturally lead towards Pareto efficient allocations of resources. This is the essence of the **Coase Theorem**, which says that Pareto efficiency will always be attained as long as there is costless bargaining.⁴

For another perspective on the Coase Theorem, consider what blues musician B.B. King said in 2000 about Napster, the pioneering music-swapping website. Commenting on the copyright-infringement lawsuits filed against Napster by recording labels and artists, B.B. said that “copyright and things of that sort are something that will have to be worked out and they will be worked out. I remember when they didn’t want you to have a VCR, but they worked it out and I think for the best. *Smart people always get together and work it out.*”⁵

Although B.B. has yet to win a Nobel Prize in Economics, his words get at the heart of the Coase Theorem: if there’s nothing stopping people from trading, nothing should stop people from trading until they reach a Pareto efficient allocation of resources.⁶ One implication is that attaining Pareto efficiency can be surprisingly simple. For example, here is a solution to the cake-cutting problem if all you are concerned about is Pareto efficiency: divide up the cake however you like, and then allow the children to trade with each other!

The Coase Theorem also suggests that the idea of trade is deep. A surprising example here is **comparative advantage**, described in problem 6.3.

⁴Ronald Coase won the 2001 Nobel Prize in Economics in part for discussing this issue.

⁵*Yahoo Entertainment News*, Sept. 13, 2000, emphasis added. Incidentally, the “they” who didn’t want you to have a VCR was the movie industry, which was afraid that people would stop going out to the movies. For details, read Carl Shapiro and Hal Varian’s *Information Rules* (1998).

⁶If bargaining is costly, of course, there is something stopping people from trading; so one implication of the Coase Theorem is that we should focus on impediments to bargaining. With Napster, a variety of nasty issues led them into bankruptcy instead of into a successful bargain with the music industry. But the words of an earlier (pre-iTunes) version of this textbook still ring true: “it seems likely that smart people will eventually figure out a method for digital distribution of music over the internet.”

Economics and Pareto efficiency

Of the four elements of “fairness” described by Brams and Taylor, economics focuses almost exclusively on Pareto efficiency. This focus on Pareto efficiency leads economists to think rather differently than non-economists, who tend to focus on equity. (The next chapter, which discusses Pareto efficiency in more depth, presents examples such as taxation and monopoly.)

It is important to remember that Pareto efficiency is only one element of “fairness”. For example, giving all of the cake to one kid satisfies the criterion of Pareto efficiency—it is not possible to make any other kid better off without making the lucky kid worse off—but that does not necessarily make it a “good” or “fair” way to cut cake.

It is also important to understand *why* economists focus on Pareto efficiency. One reason, as discussed above, is that Pareto efficiency can be surprisingly easy to attain: simply let people trade. This is certainly not the case with other aspects of “fairness”.

A second reason also provides a strong contrast between Pareto efficiency and other aspects of “fairness”: underneath their attractive exteriors, issues such as “proportionality” or “equitability” are fraught with difficulty. (What if one kid is bigger? What if one kid has been misbehaving?) Pareto efficiency, on the other hand, has a claim to universal appeal. Although it is not easy to argue that all Pareto efficient allocations of resources are *good*—would you want to live in a world where one person owned everything?—it is relatively easy to argue that all Pareto inefficient allocations of resources are in some meaningful sense *bad*: if it’s possible to make someone better off without making anyone else worse off, why not do it?

A third and final reason is political feasibility. The real world has a given distribution of resources, and if you propose something that will make some people worse off, those people are likely to fight you tooth and nail. So economics largely concerns itself with squeezing the most value out of the existing situation: we take the initial distribution of resources as given and see how we can improve upon it. In most (but not all) cases, opportunities for improvement are closely related to making it easier for people to trade.

Problems

Answers are in the next section. The first few words of each question (and each answer) are hyperlinked to allow online users to quickly move back and forth.

- 6.1 [“Differences in opinion](#) make fair division harder.” Do you agree or disagree? Explain why.
- 6.2 Explain (as if to a non-economist) the Coase Theorem and its implications for the cake-cutting problem. In other words, explain why the

economist's solution to the cake-cutting problem hinges on allowing the children to trade after the initial allocation has been made.

- 6.3 (Specialization and Gains from Trade) In this chapter we're examining the mechanisms of trade and the benefits of allowing people to trade. Here is one (long, but not difficult) numerical example about trade, based on what is sometimes called the **Robinson Crusoe model** of an economy.

Imagine that Alice and Bob are stranded on a desert island. For food, they must either hunt fish or gather wild vegetables. Assume that they each have 6 hours total to devote to finding food each day, and assume that they really like a balanced diet: at the end of the day, they each want to have equal amounts of fish and vegetables to eat. We are going to examine the circumstances under which they can gain from trade.

Story #1: Imagine that Alice is better than Bob at fishing (she can catch 2 fish per hour, and he can only catch 1 per hour) and that Bob is better than Alice at gathering wild vegetables (he can gather 2 per hour, and she can only gather 1). Economists would say that Alice has an **absolute advantage** over Bob in fishing and that Bob has an absolute advantage over Alice in gathering vegetables. Intuitively, do you think they can gain from trade? (Just guess!) Now, let's find out for sure:

- (a) If Alice and Bob could not trade (e.g., because they were on different islands), how many hours would Alice spend on each activity, and how much of each type of food would she end up with? How many hours would Bob spend on each activity, and how much of each type of food would he end up with? (Hint: Just play with the numbers, remembering that they each have six hours and want to get equal amounts of fish and vegetables.)
- (b) Now, imagine that Alice and Bob can trade with each other. Consider the following proposal: Alice will specialize in fishing, and Bob will specialize in gathering vegetables. After they each devote six hours to their respective specialties, they trade with each other as follows: Alice gives half her fish to Bob, and Bob gives half his vegetables to Alice. How many fish and how many vegetables will they each end up with in this case?
- (c) Are Alice and Bob both better off than when they couldn't trade (question 6.3a)?

Story #2: Now, imagine that Alice is better than Bob at fishing (she can catch 6 fish per hour, and he can only catch 1 per hour) and that Alice is also better than Bob at gathering wild vegetables (she can gather 3 per hour, and he can only gather 2). Economists would say that Alice has an absolute advantage over Bob in both fishing and gathering vegetables. Intuitively, do you think they can gain from trade? (Just guess!) Now, let's find out for sure:

- (d) If Alice and Bob could not trade (e.g., because they were on different islands), how many hours would Alice spend on each activity, and how much of each type of food would she end up with? How many hours would Bob spend on each activity, and how much of each type of food would he end up with?
- (e) Now, imagine that Alice and Bob can trade with each other. Consider the following proposal: Alice will specialize in fishing, increasing the amount of time that she spends fishing to 3 hours (leaving her with 3 hours to gather vegetables); and Bob will specialize in gathering vegetables, increasing the amount of time that he spends gathering vegetables to 5 hours (leaving him 1 hour to fish). After they each devote six hours as described above, they will trade with each other as follows: Alice gives 5 fish to Bob, and Bob gives 4 vegetables to Alice. How many fish and how many vegetables will they each end up with in this case?
- (f) Are Alice and Bob both better off than when they couldn't trade (question 6.3d)?

Now, forget about possible trades and think back to Alice and Bob's productive abilities.

- (g) What is Alice's cost of vegetables in terms of fish? (In other words, how many fish must she give up in order to gain an additional vegetable? To figure this out, calculate how many minutes it takes Alice to get one vegetable, and how many fish she could get in that time. Fractions are okay.) What is Alice's cost of fishing in terms of vegetables?
- (h) What is Bob's cost of fishing in terms of vegetables? What is Bob's cost of vegetables in terms of fish?
- (i) In terms of vegetables, who is the least-cost producer of fish?
- (j) In terms of fish, who is the least-cost producer of vegetables?

The punch line: Having each party devote more time to their least-cost product is the concept of **comparative advantage**.

Answers

- 6.1 Arguably, *differences in opinion* make fair division easier, not harder. For example, if one child only likes vanilla and the other child only like chocolate, then cake division is, well, a piece of you-know-what.
- 6.2 The *Coase Theorem* says that people who are free to trade have a strong incentive to trade until they exhaust all possible gains from trade, i.e., until they complete all possible Pareto improvements and therefore reach

a Pareto efficient allocation of resources. The implication for the cake-cutting problem is that a “mom” whose sole concern is efficiency can divide the cake up however she wants—as long as the children can trade, they should be able to reach a Pareto efficient allocation regardless of the starting point. For example, if you give the chocolate piece to the kid who loves vanilla and the vanilla piece to the kid who loves chocolate, they can just trade pieces and will end up at a Pareto efficient allocation.

6.3 [“Specialization and Gains from Trade...”]

- (a) Alice would spend 4 hours gathering veggies and 2 hours fishing, providing her with 4 veggies and 4 fish. Bob would do exactly the opposite (4 hours fishing, 2 hours gathering veggies) and would also end up with 4 of each.
- (b) If they specialize, Alice spends 6 hours fishing, so she gets 12 fish; Bob spends 6 hours hunting, so he gets 12 veggies. Then they split the results, so each gets 6 fish and 6 veggies, a clear Pareto improvement over the no-trade situation.
- (c) Yes.
- (d) They would allocate their time as before, but now Alice would get 12 fish and 12 veggies and Bob would get 4 fish and 4 veggies.
- (e) If they specialize as described in the problem, Alice ends up with 18 fish and 9 veggies, and Bob ends up with 1 fish and 10 veggies. After they trade, Alice ends up with 13 fish and 13 veggies, and Bob ends up with 6 fish and 6 veggies, another Pareto improvement!
- (f) Yes.
- (g) Alice must give up 2 fish to get one vegetable, and must give up 0.5 veggies to get one fish.
- (h) Bob must give up 0.5 fish to get one vegetable, and 2 veggies to get one fish.
- (i) Alice.
- (j) Bob. When they concentrate on the items for which they are the least-cost producer, they can both benefit from trade even though Alice has an absolute advantage over Bob in both fishing and gathering veggies. This is the concept of **comparative advantage**.

Chapter 7

Pareto efficiency

A priest, a psychologist, and an economist are playing a round of golf when they get stuck behind two painfully slow golfers who, despite the assistance of a caddy, are four-putting every green and taking all day to line up their shots. The resulting frustration erupts on the 8th hole: the priest yells, “Holy Mary, I pray that you should take some lessons before you play again”; the psychologist hollers, “What happened in your childhood that makes you like to play golf so slowly?”; and the economist shouts, “I didn’t expect to spend this much time playing a round of golf!”

By the 9th hole they’ve had it, so they approach the caddy and demand to be allowed to play through. “OK,” replies the caddy, “but you shouldn’t be so cruel. These are retired firemen who lost their eyesight saving people in a fire.” The priest is mortified: “Here I am, a man of the cloth, and I’ve been swearing at the slow play of two blind men.” The psychologist is also mortified: “Here I am, trained to help others with their problems, and instead I’ve been complaining selfishly.” The economist ponders the situation and finally says, “OK, but next time why don’t they play at night?”¹

By suggesting a course of action, the economist in this joke is delving into **welfare economics**, which attempts to make value judgments and policy recommendations about what is best for society as a whole. Welfare economics is rooted in the ideas in Chapter 1 about decision trees: its objective is to consider all the options and then pick the best one. Instead of focusing on a single individual, however, welfare economics seeks the best outcome for society as a whole.

The problem, of course, is that “best”—like “fair”—is hard to define. Social norms play an important role in defining “best”, which explains why welfare

¹This joke is a modification of one on the JokEc webpage, available [online](http://netec.wustl.edu/JokEc.html) at <<http://netec.wustl.edu/JokEc.html>>.

economics is also called **normative economics**. In contrast, **positive economics** seeks to make objective predictions about the ramifications of various choices. The difference between positive and normative economics is the difference between *will* and *should*: positive economics deals with what *will* happen if we adopt a given policy; normative economics deals with whether we *should* adopt that policy.

This chapter explores two definitions of “best”, beginning with perhaps the most well-known application of welfare economics: cost-benefit analysis.

7.1 Cost-benefit analysis

The basic idea of cost-benefit analysis is to put dollar values on the costs and benefits of each option under consideration and then add them up to get the **net benefits** of each option. The “best” option is the one with the largest net benefits, i.e., the biggest difference between benefits and costs. Note that many cost-benefit analyses limit the number of options under consideration to two: the existing policy and a proposed alternative. If the benefits of switching to the new policy outweigh the costs, the new policy is said to be a **cost-benefit improvement over** the existing policy.²

For an example, consider a year 2000 cost-benefit analysis by the U.S. Environmental Protection Agency (EPA) of a proposal to tighten regulations on arsenic in drinking water. In order to compare the existing regulations with the proposed alternative, the analysis put dollar values on everything from the cost of new water treatment plants to the benefits of reduced death and illness from arsenic-induced bladder cancer. (The EPA study valued a “statistical life” at \$6.1 million in 1999 dollars.) The conclusion: the proposed alternative was a cost-benefit improvement over the existing regulations.

This example highlights one of the controversial aspects of cost-benefit analysis: the practice of assigning dollar values to things that aren’t usually valued in monetary terms. In addition to inviting ethical and philosophical debate, determining things like “the value of life” also poses theoretical and practical challenges for economists who actually do cost-benefit analysis.³

Perhaps less obvious is another controversial aspect of cost-benefit analysis: by focusing on the *quantities* of costs and benefits it ignores the *distribution* of those costs and benefits. Distribution is in fact irrelevant to the analysis: a “Robin Hood” policy that helps the poor and hurts the rich receives the same treatment as a policy that does the reverse.

One response to distributional concerns is to argue that cost-benefit analysis is based on the principle of “a dollar is a dollar”, meaning that one person’s losses can be stacked up against another person’s gains in what economists call an **interpersonal utility comparison**. It is less easy, but still possible, to argue

²Economists often refer to this as a **Kaldor-Hicks improvement**; Nicholas Kaldor and John Hicks were economists who wrote key articles about this topic in 1939.

³You can learn more about these economic black arts by studying cost-benefit analysis, health economics, or environmental economics.

that cost-benefit analysis is *not* based on this principle and does not involve interpersonal utility comparisons. One such argument is that gains and losses will average out over time, so that following the prescriptions of cost-benefit analysis will, in the long run, make everybody better off. Perhaps more subtle is the argument that economists should emphasize the *possibility* of everyone being better off and leave decisions about *actual* distribution to elected officials.

These are difficult issues, and it is in large part because of them that cost-benefit analysis, though useful in practice, is unattractive in theory. Fortunately, there is an alternative concept, one based on the work of Italian economist Vilfredo Pareto (1848-1923). The key difference between cost-benefit analysis and the Pareto perspective is that—as explained in the next section—the latter does not allow for there to be any losers. It is the Pareto perspective, to which we now turn, that is the foundation of welfare economics.

7.2 Pareto

Consider again the comparison of an existing policy with some proposed alternative, and recall that the new policy is a cost-benefit improvement over the existing policy if the benefits of switching to the new policy outweigh the costs. In contrast, the new policy is a **Pareto improvement over** the existing policy if switching makes at least one person better off *and makes nobody worse off*.

This can be thought of as combining the requirement that the benefits outweigh the costs with the requirement that nobody can be made worse off. Alternately, it can be thought of as requiring a separate cost-benefit analysis for each person involved. Only if the benefits for each person are greater than or equal to the costs for that person can the new policy be described as a Pareto improvement over the existing policy.

Note that the term “Pareto improvement over” can only be used to compare two options, e.g., by asking whether option X is a Pareto improvement over option Y. The concept of Pareto improvement, like the concept of “taller”, is a comparative one. *It makes no sense to say “Option X is a Pareto improvement”, just as it makes no sense to say “Maria is taller”.*

Of course, it does make sense to say “Maria is the tallest student”, meaning that no other student is taller than her, or “Maria is not the tallest student”, meaning that there is some other student who *is* taller than her. The related Pareto concepts are Pareto efficient and Pareto inefficient. Some option (call it A) is **Pareto efficient** if there is no other option B that is a Pareto improvement over A; similarly, option A is **Pareto inefficient** if there *is* some other option B that is a Pareto improvement over A.⁴

It is also possible to define these terms without specifically mentioning Pareto improvement. An allocation of resources is Pareto efficient if there is no other allocation of resources that would make least one person better off and make nobody worse off. Similarly, an allocation of resources is Pareto inefficient if

⁴Pareto efficient allocations are sometimes referred to as being Pareto optimal or simply as efficient or optimal.

there *is* another allocation of resources that would make at least one person better off and make nobody worse off.

The Pareto concepts take some effort to apply properly. Remember that the terms “Pareto efficient” and “Pareto inefficient” apply to *individual* allocations of resources—every allocation of resources is either Pareto inefficient or Pareto efficient, depending on whether there is or is not a Pareto improvement over it—while the term “Pareto improvement” *compares* two allocations of resources. It makes no sense to say that allocation X is a Pareto improvement, just like it makes no sense to say that Las Vegas is southwest. (Southwest of what?) What does make sense is to say that Las Vegas is southwest of Toronto, or that allocation X is a Pareto improvement over allocation Y.

This analogy between directions and Pareto improvements brings up another important issue: comparing two allocations of resources is not like comparing two numbers x and y , where either $x \geq y$ or $y \geq x$. If X and Y are two allocations of resources, it is *not* true that either X is a Pareto improvement over Y or Y is a Pareto improvement over X. For example, if X is the allocation in which the first child gets all the cake and Y is the allocation in which the second child get all the cake, neither is a Pareto improvement over the other, and in fact, both are Pareto efficient outcomes. Again, the analogy with directions makes sense: comparing two allocations in terms of Pareto improvement is like comparing two cities to see if one is southwest of the other; it is possible that *neither* is southwest of the other. (One analogous result here is that there can be multiple Pareto efficient outcomes, not just one!)

7.3 Example: Taxes

As we saw in the last chapter, Pareto efficiency is only one component of “fairness”; the fact that it is the center of attention in economics produces a distinct perspective on the world. For example, Pareto efficiency in the cake-cutting problem can be attained simply by letting the children trade, so economists are likely to pay less attention than most people to the cake-cutting process itself and pay more attention to the possibility of trade after the cake is cut.

The distinct economic perspective is also evident in discussions of taxes. What bothers most people about taxes is having to pay them, but *this is not what bothers economists about taxes*. The reason is simple: a policy that results in some people handing money over to the government does not inherently result in Pareto inefficiency. Consider, for example, a government policy that imposes a **lump sum** tax of \$500 on Alice and gives Bob a \$500 lump sum payment. Such a policy *does not* produce a Pareto inefficient outcome because it is not possible to make Alice better off without making Bob worse off.

What bothers economists is that the vast majority of taxes are *not* lump sum taxes. Rather, they are taxes on behaviors such as purchasing (sales taxes), working (income and payroll taxes), and investing (business taxes). These bother economists because people can avoid some or all of these taxes by changing their behavior, e.g., by reducing the amount that they buy, work, or invest,

and it is these changes in behavior that lead to Pareto inefficiencies.

For an example, imagine that Alice is willing to pay up to \$1500 for one vacation to Hawaii each year, and up to \$1000 more for a second vacation to Hawaii each year. Happily, vacations cost only \$700, so she takes two trips each year. Now imagine that the government imposes a \$500 travel tax that raises the cost of each vacation to \$1200; it gives the resulting revenue to Bob as a lump sum payment. Note, however, that Alice will respond to the tax by taking only one trip each year. *It is the “lost trip” that results in a Pareto inefficient outcome*, a lose-lose situation that economists call a **deadweight loss**. A Pareto improvement over this outcome would be to lower the tax to \$250 per trip: Alice would then take two trips to Hawaii, so she would be better off, and Bob would still get a \$500 lump sum payment.⁵

Example: Monopoly pricing and price discrimination

A firm has a **monopoly** when it is the only seller in a market. What bothers most people about monopolies is that they charge too much, but *this is not what bothers economists about monopolies*. The reason is simple: high prices—even exorbitant prices—are not evidence of Pareto inefficiency. Imagine, for example, a movie theater that engages in **personalized pricing** by charging each person their maximum willingness-to-pay to watch a new movie. CEOs might have to pay \$100 per ticket, students might pay \$10, and it is not clear how to make consumers better off without making someone else—namely, the monopolist—worse off.

What bothers economists about monopolies is not high prices *per se*, but the fact that high prices often cause buyers to leave the market before all gains from trade have been exhausted; it is these unrealized gains from trade that are the source of inefficiency. Consider the case of **uniform pricing**, where the monopolist charges everyone the same price. (This may be because the monopolist may not know enough about different consumers to charge them different prices, or because the monopolist cannot prevent **resale** that would undermine efforts to price discriminate.) In this case it might be profit-maximizing for the monopolist to charge everybody \$100 per ticket even though this leaves the movie theater half-empty. The Pareto inefficiency here is clear: some people would be willing to pay something—not \$100, but more than nothing—to watch the movie, and it wouldn’t cost the monopolist anything to let them in. So a Pareto improvement over the uniform pricing outcome is for the monopolist to continue to charge \$100 to CEOs but to charge students a lower price: nobody is worse off and some people (the students, the monopolist, or both) are better off.

⁵In this example, lowering the tax rate per trip from \$500 to \$250 does not lower government revenue because Alice responds by buying more stuff. The dreamy idea that this is true in reality—an idea associated with the **Laffer curve** and **supply-side economics**—doesn’t hold much water, but more reasonable concepts like **dynamic scoring** are based on the same premise, namely that tax rates affect economic performance that in turn affect tax revenues.

As an aside, it is worth noting that in between charging everyone a different price (personalized pricing) and charging everyone the same price (uniform pricing) are two intermediate types of **price discrimination**. One is **group pricing**, a less extreme version of personalized pricing in which particular groups—students, seniors, etc.—are charged different prices. The other is **versioning** or **screening**, whereby sellers offer all potential customers the same menu of options but design the menu in such a way that different types of customers will choose different options. Mail-in rebates, for example, are often used instead of instant rebates because they present customers with a menu of two options—cut out the UPC code and mail it in to get the lower price, or pay more and walk away with no strings attached—that result in different customers choosing different options and the monopolist increasing its profits through price discrimination.⁶ Versioning may also help explain the dramatic differences between flying coach and first class, and anyone who flies coach can get cold comfort from Jules Dupuit’s 1849 analysis of similar circumstances in the railroad industry in France:

It is not because of the few thousand francs which have to be spent to put a roof over the third-class carriages or to upholster the third-class seats that some company or other has open carriages with wooden benches. What the company is trying to do is to prevent the passengers who pay the second class fare from traveling third class; it hits the poor, not because it wants to hurt them, but to frighten the rich.

Problems

Answers are in the next section. The first few words of each question (and each answer) are hyperlinked to allow online users to quickly move back and forth.

7.1 [Explain, as if to a non-economist](#), the following concepts, and use each in a sentence.

- (a) Pareto inefficient
- (b) Pareto improvement
- (c) Pareto efficient

⁶The terms “personalized pricing”, “versioning”, and “group pricing”—also known as first-degree, second-degree, and third-degree price discrimination, respectively—come from Carl Shapiro and Hal Varian’s excellent 1998 book *Information Rules: A Strategic Guide to the Network Economy*. The book has great stories in it, such as the one about versioning in IBM laser printers where the only difference between the cheap-but-slow E-series printers and the fast-but-expensive F-series printers was found to be a chip in the E-Series printers that made them pause at the end of each line.

- 7.2 “A Pareto efficient outcome may not be good, but a Pareto inefficient outcome is in some meaningful sense bad.”
- Give an example or otherwise explain, as if to a non-economist, why “a Pareto efficient outcome may not be good.”
 - Give an example or otherwise explain, as if to a non-economist, why “a Pareto inefficient outcome is in some meaningful sense bad.”
- 7.3 “If situation A is Pareto efficient and situation B is Pareto inefficient, situation A must be a Pareto improvement over situation B.” Do you agree with this claim? If so, explain. If not, provide a counter-example.
- 7.4 Consider a fair division problem such as the division of cake or the allocation of fishing quotas.
- Economists tend to place a great deal of importance on providing opportunities to trade (e.g., allowing the buying and selling of fishing quotas). Briefly explain why this is.
 - “Even if there are opportunities to trade, the initial allocation of resources (e.g., the determination of who gets the fishing quotas in an ITQ system) is important because it helps determine whether or not we reach *the* Pareto efficient allocation of resources.”
 - Is there such a thing as “*the* Pareto efficient allocation of resources”? Explain briefly.
 - Do you agree that initial allocations are important in order to achieve Pareto efficiency, or do you think that they’re important for a different reason, or do you think that they’re not important? Support your answer with a brief explanation.

Answers

- 7.1 [“Explain, as if to a non-economist. . .”]
- Pareto inefficient means that it’s possible to make someone better off without making anyone else worse off; in other words, there’s a “free lunch”. Example: It is Pareto inefficient to give Tom all the chicken and Mary all the veggies because Tom’s a vegetarian and Mary loves chicken.
 - A Pareto improvement is a reallocation of resources that makes one person better off without making anyone else worse off. Example: giving Tom the veggies and Mary the chicken is a Pareto improvement over giving Tom the chicken and Mary the veggies.
 - Pareto efficient means that there is no “free lunch”, i.e., it’s not possible to make someone better off without making anyone else worse off. Example: Giving Tom the veggies and Mary the chicken is a Pareto efficient allocation of resources.

7.2 [“A Pareto efficient outcome. . .”]

- (a) A Pareto efficient allocation of resources may not be good because of equity concerns or other considerations. For example, it would be Pareto efficient for Bill Gates to own everything (or for one kid to get the whole cake), but we might not find these to be very appealing resource allocations.
- (b) A Pareto inefficient allocation is in some meaningful sense bad because it’s possible to make someone better off without making anybody else worse off, so why not do it?

7.3 The claim that any Pareto efficient allocation is a Pareto improvement over any Pareto inefficient allocation is not true. For example, giving one child the whole cake is a Pareto efficient allocation, and giving each child one-third of the cake and throwing the remaining third away is Pareto inefficient, but the former is not a Pareto improvement over the latter.

7.4 [“Consider a fair division problem. . .”]

- (a) When people trade they bring about Pareto improvements—why would any individual engage in a trade unless it made him or her better off? Pareto improvements are a good thing in and of themselves, and if you get enough of them then you end up with a Pareto efficient allocation of resources.
- (b)
 - i. No. There are multiple Pareto efficient allocations.
 - ii. Initial allocations are a matter of equity; economists tend to focus on efficiency. As long as there are opportunities to trade, a Pareto efficient outcome will result *regardless of the initial allocation*.

Chapter 8

Simultaneous-move games

In contrast to sequential-move games, in which players take turns moving, **simultaneous-move games**—such as “Rock, Paper, Scissors”—involve players moving at the same time. The best way to analyze these games is to use a **payoff matrix**. An example is the **duopoly game** shown in Figure 8.1.

A duopoly is a market with only two sellers, in this case Coke and Pepsi. The two firms simultaneously choose “high” or “low” prices. If they both choose the high price, they each make profits of \$4 million. If they both choose the low price, they each make smaller profits of \$3 million. (Note that both firms choosing the high price produces larger profits for each of them than both firms choosing the low price!) But if one firm chooses the high price and one chooses the low price, customers will flock to the low-priced product: the low-pricing firm will make \$5 million and the high-pricing firm will make only \$1 million.

The payoff matrix in Figure 8.1 summarizes this information. It lists Coke’s choices in different rows, Pepsi’s choices in different columns, and the payoffs from different combinations of their choices at the intersection of the appropriate row and column, e.g., (4,4) if both firms choose the high price. By convention, the row player is player 1 and the column player is player 2, so an outcome of (5,1) indicates that player 1 (Coke) gets \$5 million and player 2 (Pepsi) gets \$1 million.

		Pepsi	
		High price	Low price
Coke	High price	4,4	1,5
	Low price	5,1	3,3

Figure 8.1: The duopoly game

8.1 Strictly dominant strategies

In the duopoly game, you might think that both players have a strong incentive to choose the high price—after all, this produces large profits for both of them—but in fact *the opposite is true when you consider the players’ individual choices*. Coke can always make more money by choosing a low price than by choosing a high price: if Pepsi chooses a low price, Coke gets \$3 million by choosing a low price and only \$1 million by choosing a high price; if Pepsi chooses a high price, Coke gets \$5 million by choosing a low price and only \$4 million by choosing a high price. No matter what Pepsi does, choosing the low price produces a higher payoff for Coke than choosing the high price, meaning that “low price” is a **strictly dominant strategy** for Coke.

Similarly, we can see that “low price” is a strictly dominant strategy for Pepsi: no matter what Coke does, Pepsi can make more money by choosing the low price instead of the high price. Since both firms act as profit-maximizing individuals, we can therefore predict that both firms will choose the low price, yielding profits of \$3 million for each firm.

8.2 The Prisoners’ Dilemma

The duopoly game is one version of the Prisoners’ Dilemma, an important game whose name comes from the version shown in Figure 8.2. The story is this: You are visiting an unnamed foreign country, waiting in line to buy a snack from a street vendor, when you and the guy in front of you in line are arrested and put into separate jail cells. A police officer comes into your cell and tells you that you and the other guy are under suspicion for armed robbery. The officer tells you that they can definitely convict both you and the other guy for some minor crime—say, littering—but that they can’t convict either of you for armed robbery unless one of you agrees to testify against the other.

The officer would really like to have an armed robbery conviction to brag about, so he offers you a deal: if you confess and the other guy doesn’t, you walk free and the other guy gets locked up for 20 years. The officer also tells you that he’s making the same offer to the other guy: if the other guy confesses and you don’t, *he* walks free and *you* get 20 years in the clink. What if you both confess, or you both keep quiet? If you both confess then you each get a

		Player 2	
		Confess	Keep quiet
Player 1	Confess	-5,-5	0,-20
	Keep quiet	-20,0	-1,-1

Figure 8.2: The Prisoners’ Dilemma

5-year sentence for armed robbery, and if you both keep quiet then you each get a 1-year sentence for littering. (This information is summarized in Figure 8.2; note that the payoffs are negative because having to spend time in jail is bad.)

You might think that you and the other guy both have a strong incentive to keep quiet—after all, this produces small jail times for both of you—but as with the duopoly game the opposite is true because confessing is a strictly dominant strategy for each player. For example, if the other guy chooses to confess then you will get a 5-year sentence by confessing versus a 20-year sentence by keeping quiet; and if the other guy chooses to keep quiet then you can get out immediately by confessing versus a 1-year sentence by keeping quiet. Of course, there is another outcome which both of you prefer to the outcome in which you both confess: you'd both be better off—i.e., it would be a Pareto improvement over both confessing—if you both kept quiet. But the incentive structure is such that you both confess!

The Prisoners' Dilemma game has an amazing number of applications, all involving what are called **collective action problems**. For example:

- You and everybody else might have shorter commutes if you all take the bus or carpool, but you can always get to work faster by driving alone. (See problem 8.1.)
- Your firm and a rival firm might both make more money if you agreed to set higher prices, but if your rival sets higher prices you can make even more money by cheating on the deal and undercutting your rival's prices. (This is the lesson from the duopoly game in Figure 8.1.)
- The different nations in OPEC (the Organization of Petroleum Exporting Countries) might all be better off if they agree to restrict the supply of oil—thereby raising the price, generating higher profits for all of them—but each of them is tempted to secretly cheat on the deal and produce more than their allotted quota.
- You and everybody else might like to burn wood in your fireplace, but if everybody does that we would have horrible air quality problems that would hurt us all.
- You and another classmate are assigned a joint project. The best outcome is for each of you to do a decent amount of work, resulting in good grades for both of you. The temptation for both of you is to be a **free-rider** by slacking off and forcing your classmate to do all the work. Of course, if both of you do this you'll both get bad grades.
- The actual Prisoners' Dilemma game in Figure 8.2 is a reasonable representation of some of the dangers present in the criminal justice system. Suspects might agree to falsely testify against others out of their own self-interest, potentially resulting in unjust convictions.

Supplemental material

Chapter 17 looks at the tragedy of the commons in the specific context of fisheries, and Chapter 18 looks at sequential-move games, games like chess in which players take turns moving.

Problems

Answers are in the next section. The first few words of each question (and each answer) are hyperlinked to allow online users to quickly move back and forth.

- 8.1 [Everybody in City X](#) drives to work, so commutes take two hours. Imagine that a really good bus system could get everybody to work in 40 minutes if there were no cars on the road. There are only two hitches: (1) If there are cars on the road, the bus gets stuck in traffic just like every other vehicle, and therefore (2) people can always get to their destination 20 minutes faster by driving instead of taking the bus (the extra 20 minutes comes from walking to the bus stop, waiting for the bus, etc.).
- (a) If such a bus system were adopted in City X and each resident of City X cared only about getting to work as quickly as possible, what would you expect the outcome to be?
 - (b) Is this outcome Pareto efficient? Explain briefly.
 - (c) “The central difficulty here is that each commuter must decide what to do without knowing what the other commuters are doing. If you knew what the others decided, you would behave differently.” Do you agree with this argument?
 - (d) What sort of mechanism do you suggest for reaching the optimal outcome in this game? Hint: Make sure to think about enforcement!
- 8.2 [\(The Public/Private Investment Game\)](#) You are one of ten students in a room, and all of you are greedy income-maximizers. Each student has \$1 and must choose (without communicating with the others) whether to invest it in a private investment X or a public investment Y. Each dollar invested in the private investment X has a return of \$2, which goes entirely to the investor. Each dollar invested publicly has a return of \$10, which is divided equally among the ten students (even those who invest privately). So if six students invest publicly, the total public return is \$60, divided equally among the ten students; the four students who invested privately get an additional \$2 each from their private investment.
- (a) What outcome do you predict if all the students must write down their investment decisions at the same time?

- (b) Is this outcome Pareto efficient? If not, identify a Pareto improvement.
- (c) "The central difficulty here is that the students must decide without knowing what the other students are doing. If you knew what the other students decided, you would behave differently." Do you agree with this argument?
- (d) If communication were possible, what sort of mechanism do you suggest for reaching the optimal outcome in this game? Hint: Make sure to think about enforcement!

8.3 **Catalytic converters** are devices that reduce the amount of pollution produced by motor vehicles. Imagine that each of the 500,000 residents of X-ville (including you) owns a car without a catalytic converter, and that each of you has to decide whether or not to purchase one. Imagine further that (1) it will cost you \$100 to purchase and install a catalytic converter; (2) *each* car that does not have a catalytic converter results in extra pollution that imposes health costs of one-tenth of one penny (\$0.001) on you and every other resident of the city; and (3) like your fellow X-villians, you just want to do whatever has the lowest cost for you personally.

- (a) If you and other X-ville residents are each allowed to choose whether or not to purchase a catalytic converter, what outcome does game theory predict?
- (b) Is this outcome Pareto efficient? Explain briefly, e.g., by identifying a Pareto improvement if the outcome is Pareto inefficient.
- (c) "The central difficulty here is that each resident must decide what to do without knowing what the other residents are doing. If you knew what the others decided, you would behave differently." Do you agree with this argument?
- (d) What sort of mechanism might you suggest for reaching the optimal outcome in this game? Hint: Make sure to think about enforcement!

8.4 **During the 2008 election campaign** there was a gay bar across the street from the headquarters of the No on Prop 8 campaign. Prop 8 was a ballot measure banning gay marriage, so as you can imagine the gay community was strongly opposed to it. Nonetheless, the headquarters of No on Prop 8 had very few people in it and the gay bar was packed. Write a paragraph discussing this phenomenon in the context of class topics such as the tragedy of the commons.

8.5 **Narrowly defined**, a "Prisoners' Dilemma" situation involves the following: (1) a symmetric, simultaneous-move game featuring two players; (2) the existence of a dominant strategy for each player; and (3) a predicted outcome that is Pareto inefficient.

- (a) Draw a payoff matrix that describes such a situation. (It may help to remember the following conventions about payoff matrices: player 1 chooses the row, player 2 chooses the column, and an outcome of $(4, 3)$ indicates that player 1 gets 4 and player 2 gets 3.)
- (b) Now draw a payoff matrix that has (1) and (2) but not (3), i.e., has an outcome that *is* Pareto efficient. (Such a game might be called the “Prisoners’ Delight”.)
- (c) Now draw a payoff matrix that has (1) but not (2) or (3). (Such a game does not have a predicted outcome according to what we’ve learned in class; if you’re curious about solution concepts for such concepts, google “Nash equilibrium”.)
- (d) A slightly broader definition of “Prisoners’ Dilemma” would include “Tragedy of the Commons” situations featuring more than two players. Provide an example of one such situation—you can describe one we’ve discussed in class, or make up your own—and briefly explain what the strategies are, what the predicted outcome is, and what would be a Pareto improvement over that predicted outcome.

8.6 What is the relationship between the Prisoners’ Dilemma and the Tragedy of the Commons?

Answers

8.1 [“Everybody in City X...”]

- (a) A good prediction is that everybody would drive to work because driving is a dominant strategy: no matter what everybody else does, you always get there 20 minutes faster by driving.
- (b) This outcome is not Pareto efficient because the commute takes 2 hours; a Pareto improvement would be for everybody to take the bus, in which case the commute would only take 40 minutes.
- (c) The central difficulty is *not* that you don’t know what others are going to do; you have a dominant strategy, so the other players’ strategies are irrelevant for determining your optimal strategy.
- (d) A reasonable mechanism might be passing a law that everybody has to take the bus or pay a large fine.

8.2 [“The Public/Private Investment Game...”]

- (a) A good prediction is that everybody will invest in the private good because it’s a dominant strategy: no matter what everybody else does, you always get \$1 more by investing privately.

- (b) This outcome is not Pareto efficient because each player only gets a return of \$2; a Pareto improvement would be for everybody to invest in the public good, in which case each player would get a return of \$10.
- (c) The central difficulty is *not* that you don't know what others are going to do; you have a dominant strategy, so the other players' strategies are irrelevant for determining your optimal strategy.
- (d) A reasonable mechanism might be passing a law that everybody has to invest in the public good or pay a large fine.

8.3 ["Catalytic converters..."]

- (a) A good prediction is that everybody would choose to not purchase a catalytic converter. For any given driver, purchasing the device would cost \$100; doing without it would impose health costs *on that driver* of only \$.001.
- (b) This outcome is not Pareto efficient. With each resident bearing health costs of \$.001 for each of the 500,000 cars in Seattle, the total health cost for each resident is \$500. A Pareto improvement would be for everyone to buy the catalytic converters, in which case each resident would only bear \$100 in costs.
- (c) The central difficulty is *not* that you don't know what others are going to do; you have a dominant strategy, so the other players' strategies are irrelevant for determining your optimal strategy.
- (d) A reasonable mechanism might be passing a law that everybody has to purchase a catalytic converter or pay a large fine.

8.4 A classic *tragedy of the commons* situation, it might have been better for the gay community as a whole if all those people put their time and money into the No on Prop 8 campaign, but each individual might still have a dominant strategy (or close to it) of going to the bar. The reason is that each individual is very unlikely to affect the outcome of the election, so each individual has an incentive to focus on personal enjoyment (going to the bar) instead of contributing to the No on Prop 8 campaign.

8.5 ["Narrowly defined..."]

- (a) There are a number of examples in the text.
- (b) See Figure 8.3.

		Player 2	
		Confess	Keep quiet
Player 1	Confess	-1,-1	0,-20
	Keep quiet	-20,0	-5,-5

Figure 8.3: The Prisoners' Delight

(c) See Figure 8.4.

		Player 2	
		Confess	Keep quiet
Player 1	Confess	-1,-1	-10,-20
	Keep quiet	-20,-10	-5,-5

Figure 8.4: The Prisoners' Confusion

(d) Anything from the traffic problem to the pollution problem to the public-private investment game to a multi-player version of the original prisoners' dilemma which gives the problem its name.

8.6 **The Tragedy of the Commons** is an extension of the Prisoners' Dilemma to situations with more than two players. Both of them feature situations in which each player has a dominant strategy and the result of all the players playing their dominant strategies is a Pareto inefficient outcome.

Chapter 9

Auctions

José robs a bank in Texas and hightails it across the border to Mexico with the Texas Rangers in hot pursuit. They finally catch him in a small town in Mexico, but the money is nowhere to be found. And the Rangers realize that they have another problem: they don't speak Spanish, and José doesn't speak English.

Eventually, they find a translator: “Tell José we want the money!” The translator tells this to José, and José replies (in Spanish) “Tell them to go to hell.” After the translator relates this, the Rangers pull out their guns and point them at José's head: “Tell him that if he doesn't give us the money, we're going to kill him!” The translator dutifully tells this to José, who begins to quake and says, “Tell them I hid the money under the bridge.” The translator turns back to the Rangers: “José says he is not afraid to die.”¹

The moral of this story is that people don't always do what we want them to do, and in particular people don't always tell the truth. This moral is at the heart of our next topic: auctions.

There are at least three reasons to use an auction to sell something.² First, auctions are a fast way to sell miscellaneous goods, hence their use by the police in disposing of impounded cars and by collectors for selling knickknacks on eBay. But speed is also important for commodities with a short shelf life, such as wholesale markets for flowers or other perishables.

Second, auctions can prevent dishonest behavior. For example, let's say I'm a procurement manager for a company that needs 100 new copying machines. It just so happens that my brother sells copying machines, and I end up deciding to buy the machines from him. Now, my manager—and my company's shareholders—might have some questions about whether my brother's

¹This joke is a modification of one in David D. Friedman's book *Hidden Order*.

²These ideas come From Elmar Wolfstetter, *Topics in Microeconomics: Industrial Organization, Auctions, and Incentives* (Cambridge University Press, 1999), p. 184.

copier was really the best choice and whether I got the best deal I could. Using an auction makes the whole process transparent and reduces opportunities for collusion.

Finally, auctions can reveal information about buyers' valuations. This is perhaps the most common reason for using auctions. If you want to sell an item but don't know how much buyers are willing to pay for it, you can use an auction to force buyers to compete against each other. A natural byproduct of this competition is that buyers "show their hands" to the seller, revealing at least something (and, as we will see below, sometimes everything!) about their valuation of the item being sold.

9.1 Kinds of auctions

There are four standard auctions. The most common is the **ascending price open auction**, also called an English auction. In this auction the bidding starts at a low value and bidders raise each other in small increments. ("Five dollars", "Six dollars", "Seven dollars", and so on, until there are no further bids: "Going once, going twice, going three times, sold!") This is the type of auction used most commonly on eBay.

A second type is the **descending price open auction**, also called a Dutch auction or a reverse auction or a reverse Dutch auction. In this auction the bid starts out at a *high* value and the auctioneer *lowers* the bid in small increments until someone calls out "Mine!", pushes a button, or otherwise indicates a willingness to buy at the specified price. This type of auction is used in the Aalsmeer Flower Auction, which completes 50,000 auctions each morning—each in a matter of seconds—sells \$1 billion in flowers every year, and holds the Guinness record for the world's largest commercial building. (You can get a virtual tour [online](#).³)

Descending price auctions are often conducted with the help of a **Dutch auction clock**. The auction starts after some (high) asking price is posted, and every tick of the clock reduces the asking price; the auction ends as soon as some buyer stops the clock by pushing their button, thereby agreeing to pay whatever price is displayed on the clock. (Here's a neat visualization for comparing ascending and descending price auctions. In an ascending price auction, the price starts out at a low level and goes up; if we imagine that each bidder starts out standing up and then sits down when the price exceeds their willingness to pay, the winner will be the last bidder standing. In a descending price auction, the price starts out at a high level and goes down; if we imagine that each bidder starts out sitting down and then stands up when they're willing to claim the object, the winner will be the first bidder standing.)

A third type of auction is the **first-price sealed bid auction**. In this auction each bidder writes down their bid and places it in a sealed envelope. The auctioneer gathers up all the bids, opens the envelopes, and awards the item to the highest bidder. That bidder pays the highest bid, i.e., their own

³<http://www.vba.nl>

bid. So if you bid \$10 and are the highest bidder then you win the auction and pay \$10.

Finally, there is the **second-price sealed bid auction**. As in the first-price sealed-bid auction, each bidder writes down their bid and places it in a sealed envelope, and the auctioneer awards the item to the highest bidder. The winning bidder, however, pays only the *second-highest* bid price. So if the highest bid is \$10 and the second-highest bid is \$7, the person who bid \$10 wins the auction, but pays only \$7.

An obvious question here is: If you want to auction off an item, which kind of auction should you use? The answer is not easy, so we'll start with a simpler question: What kind of dummy would run a second-price sealed-bid auction instead of a first-price sealed bid auction?

9.2 Bid-shading and truth-revelation

It turns out that first-price sealed-bid auctions are not obviously superior to second-price sealed-bid auctions because bidders do not have identical strategies in the two types of auctions; in fact, bidders will bid *more* in second-price sealed bid auctions! To see why, consider a buyer whose **true value** for some object is \$100, meaning that the buyer is indifferent between having \$100 and having the object.

In a first-price sealed bid auction, buyers have a strong incentive to **shade their bid**, i.e., to bid less than their true value. If you're indifferent between having \$100 and having the object, bidding \$100 for the object makes no sense: even if you win the auction you haven't really gained anything. The only way you stand a chance of gaining something is by bidding less than your true value, in which case you "make a profit" if you have the winning bid. (See problem 9.5 for mathematical details.) How much to shade your bid is a difficult question, since it depends on how much you think other people will bid, and how much they bid depends on how much they think you'll bid. . . .

In contrast, second-price sealed bid auctions are **truth-revealing**, meaning that the incentive for each buyer is to bid their true value. In fact, bidding your true value is a **weakly dominant strategy** in a second-price sealed bid auction. (The difference between strict and weak dominance is not important here, but if you're interested read this footnote.⁴)

To see why second-price sealed-bid auctions are truth-revealing, consider the highest bid *not including your own*. If this bid is *less* than your true value—say, a bid of \$80 when your true value is \$100—you cannot do better than bidding your true value: any bid you make that's more than \$80 would produce the same outcome (you win the auction and pay \$80), and any bid you make that's

⁴A **strictly dominant strategy** means that your payoff from that strategy will always be *greater* than your payoff from any other strategy; a **weakly dominant strategy** means that your payoff from that strategy will always be *greater than or equal to*—i.e., at least as large as—your payoff from any other strategy. For our purposes what's important is that in both cases you have no incentive to choose any other strategy.

less than \$80 would produce a worse outcome (you'd end up losing an auction that you would have liked to win.)

A similar argument applies if the highest bid not including your own is *more* than your true value—say, a bid of \$120 when your true value is \$100—because once again you can do no better than bidding your true value: any bid you make that's less than \$120 would produce the same outcome (you lose the auction), and any bid you make that's more than \$120 would produce a worse outcome (you'd end up winning an auction that you would have liked to lose).

In short: all you get out of bidding less than your true value in a second-price sealed bid auction is the risk of losing an auction that you'd rather win, and all that you get out of bidding more than your true value is the risk of winning an auction that you'd rather lose. So you cannot do better than bidding your true value in a second-price sealed bid auction. Our conclusion is that bidders have an incentive to shade their bids below their true value in a first-price sealed bid auction, but to truthfully reveal their true values in a second-price sealed bid auction. So it's no longer clear that first-price sealed-bid auctions will yield higher profits than second-price sealed-bid auctions!

9.3 Auction equivalences

Imagine that you're a bidder in a descending price auction, and that nature calls just as the auction is about to start. If you're lucky enough to have a trusted friend with you, what do you need to tell her so that she can bid for you? Simple: all you need to tell her is the “stand-up” price indicating when you'd call out or push the button to win the auction.

Now imagine that a case of food poisoning strikes *all* of the bidders just as the auction is about to start, so that you all need to head off to the loo. Imagine further that your friend is a trusted friend of all the other bidders, too, and that they all tell her their stand-up prices. Armed with all this information, your friend could participate in the auction on behalf of everyone.

Of course, your friend wouldn't actually have to go through the auction process in order to determine the outcome. Just by looking at all the bidders' stand-up prices, she can tell who's going to win the auction: the bidder with the highest stand-up price. And she can tell what price that winning bidder is going to pay: a price equal to that bidder's stand-up price.

But this looks exactly like a first-price sealed bid auction! Indeed, an auctioneer tiring of a descending price auction could simply ask the bidders to write down their stand-up prices and then award the item to the highest bidder in exchange for that bidder's stand-up price. The punch line is that descending price auctions are strategically equivalent to first-price sealed bid auctions! Bidders should have identical strategies in the two auctions, and the outcomes of the two auctions should be identical.

A similar story shows that ascending price auctions are strategically equivalent to second-price sealed bid auctions. A bidder needing to use the loo could simply tell a friend the “sit-down” price beyond which they're not willing to

Let's say you find something on eBay that you want... You're willing to pay \$25.00 for it, but the current bid price is only \$2.25. You could take the long route and sit at your computer, outbidding each new bid until you reach \$25.00.

Luckily, there's a better way. Here's how it works:

- 9.1 Decide the **maximum** you're willing to pay and enter this amount.
- 9.2 eBay will now confidentially bid **up** to your maximum amount. In this way, you don't have to keep an eye on your auction as it unfolds.
- 9.3 If other bidders outbid your maximum at the end of the auction, you don't get the item. But otherwise, you're the winner—and the final price might even be less than the maximum you had been willing to spend!

Remember: eBay will use only as much of your maximum bid as is necessary to maintain your position as high bidder. Winning was never easier!

Figure 9.1: eBay's Proxy Bidding feature

continue bidding. If that same friend could see these sit-down prices for all of the bidders, she could anticipate the outcome of the auction: the bidder with the highest sit-down price would win and would pay an amount essentially equal to the *second-highest* sit-down price. (If the highest sit-down price was \$100 and the second-highest was \$20, everybody except for the highest bidder would drop out at a price of about \$20, at which point the auction would be over.⁵) For further supporting evidence, Figure 9.1 duplicates the [online](#)⁶ description of the Proxy Bidding feature on eBay, which runs ascending price auctions. Note that Proxy Bidding effectively turns ascending price auctions into second-price sealed-bid auctions!

In conclusion, the four auctions we began with can be divided into two pairs of strategically equivalent auctions. The auctions in each pair—one pair consisting of ascending price auctions and second-price sealed bid auctions, the other pair consisting of descending price auctions and first-price sealed bid auctions—share the same essential properties. For example, we showed earlier that bidding your true value is a weakly dominant strategy in a second-price sealed bid auction. Bidders in an ascending price auction also have a weakly dominant strategy: to continue bidding as long as the asking price is less than your

⁵The qualifying terms “essentially” and “about” are necessary because the exact determination of the final price depends on the minimum bid increment—the minimum raise required over the previous bid—and the order of the bidding. For present purposes, it helps to think of the minimum bid increment as being \$.01 or some other tiny amount.

⁶<http://pages.ebay.com/help/buyerguide/bidding-prxy.html>

true value. Note also that the existence of dominant strategies eliminates the strategic tension from these auctions.

In contrast, strategic tension is quite evident in the other pair of auctions, which have no dominant strategies. There is an obvious element of suspense in a descending price auction: the winning bidder wants to hold off on bidding until just before another bidder is going to bid. Less obvious is the identical tension in first-price sealed bid auctions: the winning bidder here wants to bid only slightly more than the next-highest bidder.

We can now return to one of our earlier questions: If you want to auction off an item, which kind of auction should you use? We've already seen that ascending-price auctions are equivalent to second-price sealed bid auctions and that descending-price auctions are equivalent to first-price sealed bid auctions. We have also seen that a first-price sealed-bid auction is not—as it seems at first—clearly superior to a second-price sealed bid auction because bidders will shade their bids in a first-price sealed-bid auction and reveal their true values in a second-price sealed-bid auction.

A deep and remarkable result called the **Revenue Equivalence Theorem** says that in many cases all of these auctions yield the same expected revenue: if you have an item to auction off, your expected revenue is identical regardless of the type of auction you choose! This result is comforting to economists because it helps explain the variety of auctions used in the world: if seller could get more revenue from a particular type of auction, one would expect that type of auction to be the dominant form, but if all auctions yield the same expected revenue then sellers have no reason to gravitate towards one type.

9.4 Auction miscellany

There are many other fascinating topics and results about auctions. Here are a handful of examples to whet your appetite.

All-pay auctions

In an all-pay auction, the bidders all submit bids, and the object goes to the highest bidder, but all bidders pay their bid price. So if you bid \$10 and I bid \$7, you win the object and pay \$10, and I don't win the object and pay \$7. Examples include:

Political lobbying Each lobbying group “donates” money to politicians, hoping to win the favor of those politicians. The politician may respond by agreeing to the requests of the highest-paying donor. But the losing donors don't get their money back!

Queuing (waiting in line) Buyers hoping to get tickets to crowded shows pay in part with their time—after all, time is money—but the buyers who get turned away when the tickets are all sold do not get their time back.

Patent races In a research and development (R&D) contest, different firms invest money in the hopes of being the first firm to develop a particular drug or gizmo. Those sums are lost even if the firm loses the race.

The winner's curse

Until this point we have been implicitly assuming that the bidders have **independent values**, meaning that my value of the object being auctioned (e.g., a painting) isn't related to your value of the object. When the object is worth the same amount to all the bidders—for example, a jar of coins or an oil field, or maybe even a star baseball player—we have a **common values** auction, and we can sometimes see a phenomenon called the **winner's curse**.

Common values auction are interesting because none of the bidders know what the common value is. For example, the oil companies bidding for the right to drill under a given plot of land don't know how much oil is really there; instead, they must rely on estimates. It is likely that some of these estimates are higher than the true amount and some are lower, but we can imagine that *on average* they are correct, i.e., that their estimates fluctuate around the true amount. The winner's curse materializes if each bidder bids his or her own estimate, because then the winner will be the bidder with the highest estimate. . . an estimate that almost certainly exceeds the true amount! The company that wins the right to drill for oil ends up finding less oil than they estimated, meaning that they she might very well lose money in the end.⁷

Multi-unit auctions

The Federal Communications Commission (FCC) and similar bodies in other countries have and continue to conduct spectrum auctions to allocate various wavelengths, e.g., for use in wireless phones. These auctions have generated billions of dollars in revenue and are truly complicated. One complication arises because different areas are connected to each other: if you own spectrum rights in Tacoma, your value for spectrum rights in Seattle is likely to increase, and vice versa. What the FCC has done is use a multi-unit auction to auction off all this airspace at the same time.

Appropriately designing these auctions to maximize revenue is a thorny problem for microeconomists, and experience has shown that serious strategic behavior can arise in these auctions. One example, a phenomenon called **code bidding**, is described in Figure 9.2 and the accompanying text, both excerpted from a working paper⁸ on the 1996-1997 FCC auction:

⁷A neat article on this topic involving jars of pennies and business students is Bazerman, M.H. and Samuelson, W.F., "I won the auction but don't want the prize", *Journal of Conflict Resolution* 27:618-634 (1983).

⁸Peter Crampton and Jesse A. Schwartz, "Collusive Bidding in the FCC Spectrum Auctions", November 24, 1999.

	Marshalltown, IA 283 E		Rochester, MN 378 D		Waterloo, IA 452 E		
Round	McLeod	USWest	McLeod	USWest	AT&T	McLeod	USWest
24	56,000					287,000	
...					
46				568,000			
52			689,000				
55				723,000			
58			795,000				
59				875,000			313,378
60						345,000	
62			963,000				
64		62,378		1,059,000			
65	69,000						
68					371,000		

Figure 9.2: Code bidding in the FCC spectrum auction

[Figure 9.2] shows all of the bids that were made on Marshalltown, block E and Waterloo, block E after round 24, and all of the bids on Rochester, block D after round 46. USWest and McLeod were contesting Rochester, trading bids in rounds 52, 55, 58, and 59. Rather than continue to contest Rochester, raising the price for the eventual winner, USWest bumped McLeod from Waterloo in round 59 with a code bid, \$313,378. The “378” signified market 378—Rochester. USWest’s bid revealed that McLeod was being punished on Waterloo for bidding on Rochester. In round 60, McLeod retook Waterloo, bidding \$345,000, \$58,000 more than its round 24 bid. But McLeod did not yet concede Rochester—it placed another bid on Rochester in round 62. USWest then used the same technique in round 64, punishing Marshalltown instead. USWest’s bid in round 64 on Rochester won the license.

Problems

Answers are in the next section. The first few words of each question (and each answer) are hyperlinked to allow online users to quickly move back and forth.

- 9.1 *Fun/Challenge.* The website freemarkets.com runs procurement auctions: companies in need of supplies post information about their purchasing needs (e.g., so and so many sheets of such and such kind of glass) and the maximum amount they're willing to pay for those purchases; bidders then bid the price down, and the lowest bidder receives that price for the specified products. The ads for freemarkets.com say things like, "At 1pm, Company X posted a request for 1 million springs, and indicated that it was willing to pay up to \$500,000. By 4pm, the price was down to \$350,000."
- (a) Explain how an auction can help Company X get a low price on springs.
 - (b) Is the example ad above impressive? Is it susceptible to gaming (i.e., strategic manipulation)?
- 9.2 You're a bidder in a second-price sealed-bid auction. Your task here is to explain (as if to a mathematically literate non-economist) why you should bid your true value.
- (a) Explain (as if to a non-economist) why you cannot gain by bidding *less* than your true value.
 - (b) Explain (as if to a non-economist) why you cannot gain by bidding *more* than your true value.
- 9.3 You're a bidder in a first-price sealed bid auction.
- (a) Should you bid your true value, more than your true value, or less than your true value? Explain briefly, as if to a mathematically literate non-economist.
 - (b) Do you have a dominant strategy, i.e., is your optimal bid always the same regardless of others' bids? Explain why or why not.
- 9.4 Your mathematically literate but non-economist friend Jane owns one of the few original copies of *Send This Jerk the Bedbug Letter!*, a best-selling book about playing games with giant corporations. She decides to auction off the book to raise money for her new dot.com venture. She tells you that she's going to use a first-price sealed bid auction. You ask her why she doesn't use a second-price sealed bid auction, and she looks at you like you're nuts: "Look, dummy, I'm trying to make as much money as I can. Why would I charge the second-highest bid price when I can charge the highest bid price!?" Write a response.

9.5 We can use the expected value calculations from Chapter 4 to get another perspective on bidding in first- and second-price sealed bid auctions.

- (a) The first step in calculating expected values is determining the different possible outcomes. So: what are the possible outcomes of bidding $\$x$ in an auction?
- (b) Next: write down and simplify an expression for the expected value of bidding $\$x$ in an auction. Use $\text{Value}(\text{Winning})$ to denote the value of winning the auction. Assume that the value of losing the auction is zero.
- (c) Write down an expression for the expected value of bidding $\$x$ in a first-price sealed bid auction. Assume that your gain or “profit” from winning an auction is the difference between your true value for the item and the price you actually have to pay for it. Can you use this expected value expression to highlight the issues faced by a bidder in such an auction? For example, can you show mathematically why bidders should shade their bids?
- (d) Write down an expression for the expected value of bidding $\$x$ in a second-price sealed bid auction. (Again, assume that your gain or “profit” from winning an auction is the difference between your true value for the item and the price you actually have to pay for it.) Can you use this expected value expression to highlight the issues faced by a bidder in such an auction? For example, can you show mathematically why bidders should bid their true value?

9.6 It just so happens that eBay is currently running an auction for a collection of *all five* *NSYNC bobblehead dolls. Imagine that your value for such a collection is $\$20$, meaning that you are indifferent between having the dolls and having $\$20$.

- (a) In a first-price sealed bid auction, should you bid an amount b that is (less than equal to more than) your true value ($\$20$)? Explain briefly. It may help to write down an expected value calculation.
- (b) In a second-price sealed bid auction, explain why it makes sense to bid your true value (i.e., $\$20$). *Hint:* Consider the highest bid *excluding* your own bid. If that bid is more than $\$20$, can you do better than bidding your true value? If that bid is less than $\$20$, can you do better than bidding your true value?
- (c) Your friend Ed needs some cash, so he decides to auction off his prized collection of *NSYNC bobblehead dolls. You suggest a second-price sealed bid auction, to which he says, “Second price? Why should I accept the *second-highest* price when I can do a first-price sealed bid auction and get the *first-highest* price?” Write a response. *Hint:* Think about your answers to the first two auction questions above.

Answers

9.1 [“The website freemarkets.com...”]

- (a) Auctions pit different suppliers against each other, and their individual incentives lead them to drive down the price. This helps ensure that Company X will not be paying much more for springs than it costs the suppliers to produce them.
- (b) The example in the ad above may not be as impressive as it sounds because of the potential for gaming: if Company X knows that a number of firms can produce the springs for about \$350,000, it has essentially nothing to lose by indicating a willingness-to-pay of \$500,000—or even \$1,000,000—because the auction dynamics will drive the price down toward \$350,000. An analogy may help: say I want to purchase a \$20 bill. As long as there are enough competitive bidders, I can more-or-less fearlessly say that I’m willing to pay up to \$1,000 for that \$20 bill; competitive pressures will force the winning bid down to about \$20.

9.2 [“You’re a bidder...”]

- (a) The intuition can be seen from an example: say you’re willing to pay up to \$100, but you only bid \$90. Let y be the highest bid not including your bid. If $y < 90$ then you win the auction and pay y ; in this case, bidding \$90 instead of \$100 doesn’t help you or hurt you. If $y > 100$ then you lose the auction and would have lost even if you bid \$100; again, bidding \$90 instead of \$100 doesn’t help you or hurt you. But if y is between \$90 and \$100 (say, $y = 95$) then bidding \$90 instead of \$100 actively hurts you: you end up losing the auction when you would have liked to have won it. (You had a chance to get something you value at \$100 for a payment of only \$95, but you didn’t take it.)
- (b) Again, the intuition can be seen in the same example in which you’re willing to pay up to \$100. Assume that you bid \$110 and that y is the highest bid not including your bid. If $y < 100$ then you win the auction and pay y ; in this case bidding \$110 instead of \$100 doesn’t help you or hurt you. If $y > 110$ then you lose the auction; again, bidding \$110 instead of \$100 doesn’t help you or hurt you. But if y is between \$100 and \$110 (say, $y = 105$) then bidding \$110 instead of \$100 actively hurts you: you end up winning the auction when you would have liked to have lost it. (You pay \$105 for something you only value at \$100.)

9.3 [“You’re a bidder...”]

- (a) You should bid less than your true value. If your true value is, say, \$100, then you are indifferent between having the object and having

\$100. If you bid \$100, winning the auction won't make you better off; if you bid more than \$100, winning the auction will actually make you worse off. The only strategy that makes it possible for you to be better off is for you to bid less than \$100.

- (b) You do *not* have a dominant strategy in this auction. If the highest bid other than yours is \$90, you want to bid \$91. If the highest bid other than yours is \$20, you want to bid \$21. Since your optimal bid depends on what others bid, you do not have a dominant strategy; we know that you want to shade your bid, but we don't know by how much without considering the other players.

9.4 A *reasonable response* might start off by noting that bidders will behave differently in the two auctions: bidders will shade their bids in a first-price auction, but not in a second-price auction. So in a first-price auction you get the highest bid from among a set of relatively low bids, and in a second-price auction you get the second-highest bid from among a set of relatively high bids. It's no longer clear which auction has the higher payoff. (In fact, there is a deeper result in game theory, called the Revenue Equivalence Theorem, which predicts that both types of auctions will yield the same expected payoff.)

9.5 ["We can use..."]

- (a) There are two possible outcomes: either $\$x$ is the highest bid and you win the auction, or $\$x$ isn't the highest bid and you lose the auction.
- (b) Your expected value from bidding $\$x$ in the auction is

$$\begin{aligned} \text{EV}(\text{Bidding } \$x) &= \text{Pr}(\text{Your } \$x \text{ bid wins}) \cdot \text{Value}(\text{Winning}) \\ &\quad + \text{Pr}(\text{Your } \$x \text{ bid loses}) \cdot \text{Value}(\text{Losing}). \end{aligned}$$

Since the value of losing is zero (you get nothing, you pay nothing), the second term disappears. So your expected value boils down to something like

$$\text{EV}(\text{Bidding } \$x) = \text{Pr}(\text{Your } \$x \text{ bid wins}) \cdot \text{Value}(\text{Winning}).$$

- (c) The expression above simplifies to

$$\text{EV}(\text{Bidding } \$x) = \text{Pr}(\text{Your } \$x \text{ bid wins}) \cdot (\text{Value of object} - \$x).$$

Here we can see that bidding your true value is a bad idea: your expected value will never be greater than zero! We can also see the tension at work in first-price sealed bid auctions: by reducing your bid, you lower the probability that you will win, but you increase the value of winning. (Optimal bidding strategies in this case are complicated. How much to shade your bid is a difficult question, since it depends on how much you think other people will bid...)

- (d) Your expected value of bidding $\$x$ reduces to

$$EV(\text{Bidding } \$x) = \Pr(\text{Your } \$x \text{ bid wins}) \cdot (\text{Value of object} - \$y)$$

where $\$y$ is the *second-highest* bid. Since the price you pay is not determined by your own bid, shading your bid below your true value doesn't help you. It only increases the probability that you will lose the bid when you would like to have won it. (The same is true for bidding over your true value. This only increases the probability that you will win the object and be forced to pay an amount greater than your true value.) You maximize your expected value by bidding your true value.

9.6 ["It just so happens..."]

- (a) You should bid less than your true value. Otherwise your expected value from the auction will never be more than zero (and will be less than zero if you bid more than your true value):

$$EV = \text{Prob}(\text{Win}) \cdot (20 - b) + \text{Prob}(\text{Lose}) \cdot (0).$$

- (b) If the highest bid excluding your own bid is $x > \$20$, you cannot do better than bid $\$20$ (and lose the auction); the only way to win the auction is to bid more than x , but if you do that then you'll end up paying x , which is more than your true value. On the other hand, if the highest bid excluding your own is $x < \$20$, you cannot do better than bid $\$20$ (and win the auction, paying $\$x$); raising your bid cannot help you, and lowering your bid doesn't reduce the amount you'll pay, but does increase your risk of losing the auction when you would have liked to have won it.
- (c) Yes, in a first-price sealed bid auction you'll get the first-highest price; but we showed above that bidders will bid less than their true value. In contrast, bidders will bid an amount equal to their true value in a second-price sealed bid auction. So even though you only get the second-highest bid, the bid values will be higher than in a first-price auction. (A deeper result here is the revenue equivalence theorem, which says that these two types of auctions have the same expected payoff for seller.)

Chapter 10

From some to many

Mr. Smeds and Mr. Spats

Mr. Spats / Had twenty-one hats,
And none of them were the same.
And Mr. Smeds / Had twenty-one heads
And only one hat to his name.
Now, when Mr. Smeds / Met Mr. Spats,
They talked of the / Buying and selling of hats.
And Mr. Spats / Bought *Mr. Smeds'* hat!
Did you ever hear anything / Crazier than that?
– Shel Silverstein,
A Light in the Attic (1981)

The chapter on auctions highlights the power of competition: in situations with only one seller and lots of buyers, the seller can benefit by getting the buyers to compete against each other. Even if the seller would be willing to sell for a low price, *and even if the potential buyers know that the seller would be willing to sell for a low price*, forcing the potential buyers to compete against each other in an auction can generate a sale price much higher than the seller's minimum price.

A symmetric result occurs in situations with only one buyer and lots of sellers. The buyer can then use an auction to get sellers to compete against each other.¹ Even if the buyer would be willing to pay a very high price, *and even if the potential sellers know that the buyer would be willing to pay a high*

¹For example, a firm might hold an auction for some mechanical parts that it needs, or a government might auction off a contract to build a bridge; in a first-price sealed bid auction, the bidder with the *lowest* price would win the contract and would receive a payment each to its bid.

price, forcing the potential sellers to compete against each other in an auction can generate a sale price much lower than the buyer's maximum price.

In sum: with one seller and many buyers, competition helps the seller; with one buyer and many sellers, competition helps the buyer. An obvious question, then, is: what happens if there are many buyers *and* many sellers? This is the subject matter of the next part of this text.

Our focus in Part II has been on game theory, i.e., interactions between two or more optimizing individuals. Part III focuses more specifically on (1) *market* interactions, i.e., interactions between buyers and sellers; and (2) *competitive* interactions in which there are many buyers and many sellers, each small in relation to the whole. Such situations are known as **competitive markets**; the branch of economics that studies competitive markets is **price theory**.

Price theory give us powerful tools for analyzing interactions between optimizing individuals. But it is important to recognize the limitations of these tools. In particular, they are only appropriate in competitive markets: there must be many buyers, each small in relation to all the buyers together; and there must be many sellers, each small in relation to all the sellers together. For a visual analogy, picture a sunny beach with no big rocks or boulders, just small grains of sand, each tiny in relation to the whole beach.

The theoretical ideal is sometimes (and somewhat redundantly) called a **perfectly competitive market**. Of course, the theoretical ideal is just that: an ideal. In reality, competitiveness is a matter of degrees. Some markets—such as those for wheat, engineers, or houses—come quite close to the perfectly competitive idea, while others—such as the market for airplanes, which is dominated by Boeing and Airbus—are clearly not competitive. Another group of markets—such as the market for coffee, in which Starbucks is a significant buyer—come somewhere in between, and it is not always easy to determine whether or not the competitive model applies. In such situations, it can help to think about the purpose of the competitive market restrictions.

That purpose is to ensure that any individual buyer or seller cannot affect the market price or any other “big picture” variable. Each individual therefore takes the market price as given; such an individual is called a **price-taker**. The key assumption in price theory is that all the buyers and sellers are price-takers.

As we will see, the importance of the price-taking assumption is that it eliminates opportunities for strategic behavior. If strategic behavior is possible, we must use the tools from Part II. The absence of meaningful strategic behavior in competitive markets is what allows us to use the powerful tools from Part III.

10.1 Monopolies in the long run

Question: So monopolies can get big profits. What's wrong with that?

For one thing, consumers “pay too much”. This may not be inefficient, but voters (and hence politicians) may not be thrilled about big rich companies raking in money hand over fist from poor consumers. Also, we have seen two sources

of inefficiency from monopolies: a monopoly may engage in inefficient behavior (such as mail-in coupons) in order to price discriminate, and a monopoly that cannot price discriminate may set a price that is inefficiently high (i.e., charge \$25 instead of \$10 even though the product costs nothing to produce and there is one consumer who is willing to pay \$10 and one consumer who is willing to pay \$25).

But there's also something very *right* about monopolies making big profits: those profits entice other firms to enter the market. The first company that made scooters made a pile of money, but then other firms saw that there was money to be made in scooters, and they started making scooters, too. So profits serve as a signal to the market, and *in getting high profits monopolies are sowing the seeds of their own demise*.

To make an analogy, recall that in Chapter 5 we compared different investments to the various lanes of traffic on a congested bridge. We concluded that financial arbitrage should yield comparable expected returns for comparable investments, just like transportation arbitrage should yield comparable expected travel times for all the different lanes. Now imagine that some adventurous driver builds a new lane, or discovers a new route. At first this driver gets to go faster than everybody else, but it won't be long before other drivers notice and follow along. Eventually, the new lane should have the same travel time as all the other lanes.

In other words: If a company finds a new market and makes monopoly profits, other businesses will try to enter that market. If and when they do, prices will fall and rates of return will be equalized with comparable investments. This is the topic of the coming chapters on **competitive markets**.

10.2 Barriers to entry

If a company is to maintain a monopoly (and monopoly profits) in the long run, then, there must be something preventing other firms from entering the market. In other words, there must be one or more **barriers to entry**. These could be legal barriers, such as patents or copyrights that prevent others from copying an innovation or work of art. (Question: Why do governments establish and enforce such barriers to entry if they lead to monopolies and monopoly pricing?) There might also be economic barriers to entry, e.g., control over a key asset or economies of scale.

Part III

Market interactions

Chapter 11

Supply and demand

If we go to the grocery store and stand in the milk aisle, we can observe the **quantity transacted** (the number of gallons of milk bought and sold) and we can observe the **market price** (the amount the buyers pay the sellers for each gallon). Figure 11.1a graphs these observables, p and q . The figure also shows an area representing their product, pq , which is **total revenue** and also **total expenditure**. (If the market price p is \$2.00 per unit and the quantity transacted q is 1,000 units, then the total amount of money that changes hands—the total revenue for sellers, and the total expenditure for buyers—is \$2,000. Graphically we get a rectangle of height p , width q , and area pq .)

If we go back to the grocery store on other occasions, we can see the observables p and q changing over time. (Graphically, we get Figure 11.1b.) The story of supply and demand attempts to explain what's going on with p and q , i.e., with the dot we see moving around Figure 11.1b.

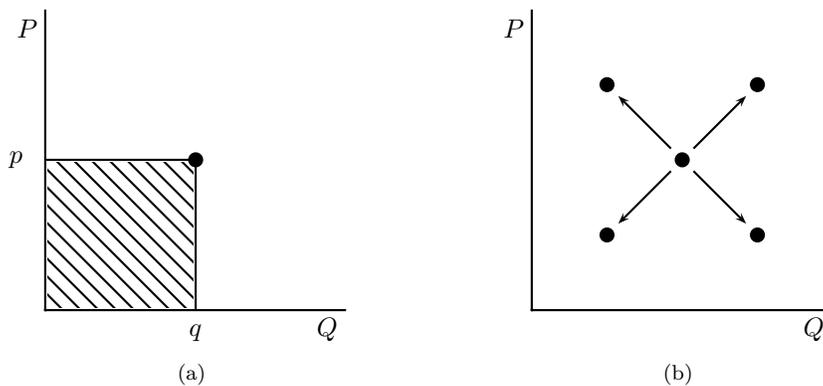


Figure 11.1: (a) the observables p and q Total revenue and/or total expenditure is the shaded box with area pq ; (b) the observables changing over time.

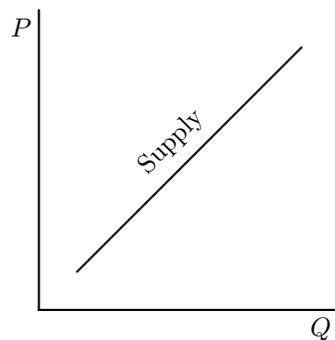
11.1 The story of supply and demand

The story of supply and demand is that the dot we see is really the intersection of two curves: the market supply curve and the market demand curve. When the dot moves around, it is because these curves are moving.

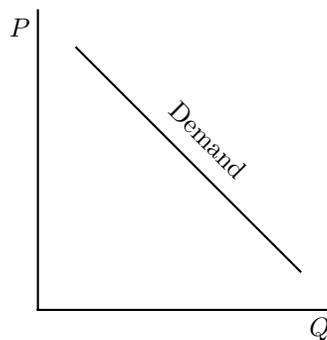
The **market supply curve** (see Figure 11.2a) answers *hypothetical* questions like these: *If* the market price were \$6 per unit, how many units of this good would sellers want to sell? *If* the market price were \$7 per unit, how many units of this good would sellers want to sell? Put all these hypotheticals together, and you get the market supply curve.¹

Similarly, the **market demand curve** (see Figure 11.2b) is a graphical presentation of *hypothetical* questions like these: *If* the market price were \$6 per unit, how many units of this good would buyers want to buy? *If* the market price were \$7 per unit, how many units of this good would buyers want to buy? Put all these hypotheticals together, and you get the market demand curve.

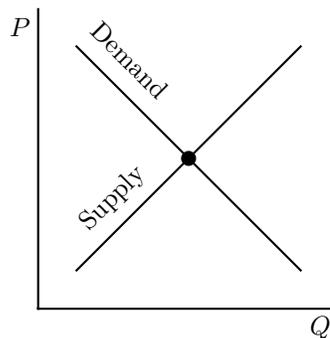
¹It might seem backwards to have the market price on the y -axis and the quantity on the x -axis. This counterintuitive set-up is explained further in Chapter 13.



(a) A market supply curve



(b) A market demand curve



(c) The market equilibrium

Figure 11.2: Supply and demand

Note from Figure 11.2 that supply and demand curves slope in opposite directions. The supply curve is *upward sloping*, indicating that as the price per unit rises sellers want to sell more units, or, equivalently, that as the price per unit falls sellers want to sell fewer units. The demand curve is *downward sloping*, indicating that as the price per unit rises buyers want to buy fewer units, or, equivalently, that as the price per unit falls buyers want to buy more units. (This phenomenon is known as the **Law of Demand**.) We will examine these slopes in more detail in Chapter 20.

Question: Why should we expect the observables p and q to show up at the intersection of the supply and demand curves?

Answer: Because at that intersection point the following equality holds: the quantity that buyers want to buy at that price is equal to the quantity that sellers want to sell at that price. It is because of this equality that this price is called the **market-clearing price** and the intersection point is called the **market equilibrium**.

Any outcome other than this market equilibrium is incompatible with individual optimization. At any price higher than the market-clearing price, the quantity that buyers want to buy is *less* than the quantity that sellers want to sell; this creates incentives for individual sellers to lower their prices and for individual buyers to seek out lower prices. At any price lower than the market-clearing price, the quantity that buyers want to buy is *more* than the quantity that sellers want to sell; this creates incentives for individual buyers to increase their offers and for individual sellers to seek out higher prices. Only at the intersection of the supply and demand curves does the quantity that sellers want to sell at a certain price equal the quantity that buyers want to buy at that same price.

11.2 Shifts in supply and demand

A good way to think about the market equilibrium is to imagine that the demand curve is blue, that the supply curve is yellow, and that the only color we can see in the real world is green. The market equilibrium comes at the point where the two curves intersect, and the punch line—yellow and blue makes green!—carries an important lesson: the green dot has no independent existence of its own, and it doesn't move unless either the yellow line or the blue line moves. In other words, the observables p and q do not change unless either the demand curve or the supply curve changes.

In most cases, outside events will affect only one of the two curves: a late frost in Florida will affect the supply curve for orange juice but not the demand curve; a news report on the health benefits of orange juice will affect the demand curve for orange juice but not the supply curve. With only one of the two curves changing, there are four basic patterns, as shown in Figure 11.3.

Demand increases Buyers want to buy more than before at any given price, so the demand curve shifts to the *right*. The equilibrium price and quantity both increase. (See Figure 11.3a.)

Demand decreases Buyers want to buy less than before at any given price, so the demand curve shifts to the *left*. The equilibrium price and quantity both decrease. (See Figure 11.3b.)

Supply increases Sellers want to sell more than before at any given price, so the supply curve shifts to the *right*. The equilibrium price decreases and the equilibrium quantity increases. (See Figure 11.3c.)

Supply decreases Sellers want to sell less than before at any given price, so the supply curve shifts to the *left*. The equilibrium price increases and the equilibrium quantity decreases. (See Figure 11.3d.)

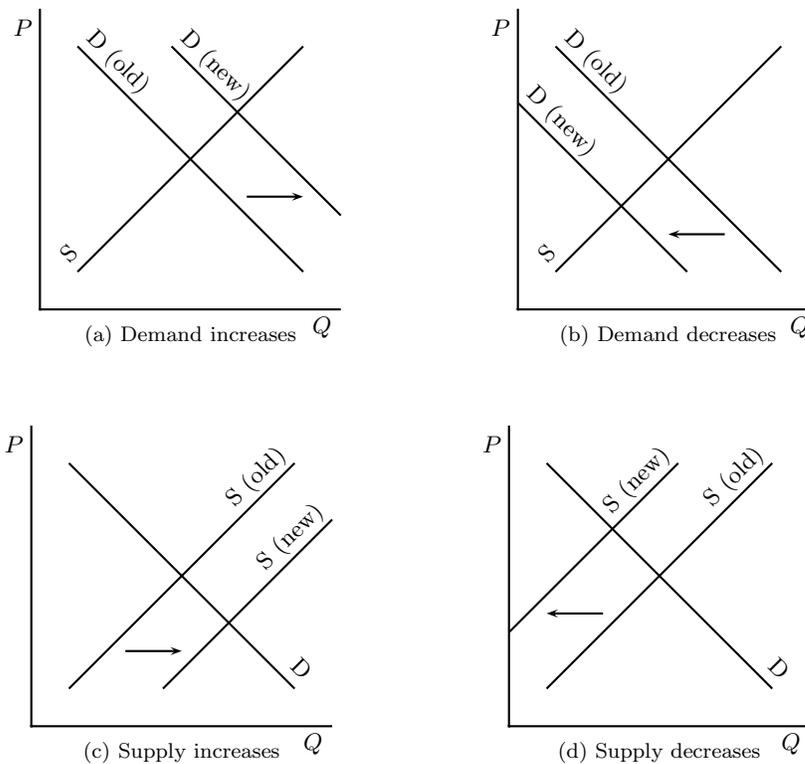


Figure 11.3: The four basic types of shifts in supply and demand. Note that an increase in one of the curves results in a shift to the *right*, and that a decrease results in a shift to the *left*. So increases *do not* shift curves up, and decreases *do not* shift curves down.

Note that an increase in either supply or demand shifts the relevant curve to the *right* (at any price people want to buy or sell more) and that a decrease in either supply or demand shifts the relevant curve to the *left* (at any price people want to buy or sell less). The curves move left or right, *not up or down!* This counterintuitive behavior stems from the fact—to be discussed more in Chapter 13—that we have the quantity Q on the x -axis and the market price P on the y -axis.

11.3 Algebra: The algebra of markets

The market equilibrium occurs at the intersection of the supply curve and the demand curve, and therefore lies on both of these curves. Algebraically, this means that we can find the market equilibrium by simultaneously solving the equations for the supply and demand curves. If the supply curve is $q = S(p)$ and the demand curve is $q = D(p)$, then the market equilibrium is the point (p^*, q^*) such that $q^* = S(p^*)$ and $q^* = D(p^*)$.

For an example, consider the supply curve $q = 15 + 2p$ and the demand curve $q = 20 - 3p$. We can solve these equations simultaneously:

$$20 - 3p = 15 + 2p \implies 5p = 5 \implies p = 1.$$

We can then use this value of p to find q from either the supply curve ($q = 15 + 2p = 17$) or the demand curve ($q = 20 - 3p = 17$). So the market equilibrium occurs at a price of $p = 1$ per unit and a quantity of $q = 17$ units.

Supplemental material

Chapter 20 looks at some details of supply and demand and shows how market supply and demand curves are built from individual supply and demand curves, which in turn are built from decision trees and individual optimization by individuals and firms.

Problems

Answers are in the next section. The first few words of each question (and each answer) are hyperlinked to allow online users to quickly move back and forth.

- 11.1 Explain, as if to a non-economist, why the intersection of the market supply curve and the market demand curve identifies the market equilibrium.
- 11.2 For each item, indicate the likely impact on the supply and demand for wine. Then indicate the effect on the equilibrium price and quantity. It may help to use a graph.

- (a) The legal drinking age for wine is lowered to 18.
 - (b) A fungus destroys part of the grape harvest. (Grapes are **inputs** in wine-making, as are labor, machinery, and glass.)
 - (c) The price of cheese increases. (Wine and cheese are **complements** or **complementary goods**, as are skis and ski boots; monitors and keyboards; and peanut butter and jelly.)
 - (d) The price of beer falls. (Beer and wine are **substitutes**, as are eyeglasses and contact lenses; burritos and hamburgers; and pens and pencils.)
- 11.3 For each item, indicate the likely impact on the supply and demand for popsicles in Hawaii. Then indicate the effect on the equilibrium price and quantity. It may help to use a graph.
- (a) More tourists visit Hawaii.
 - (b) An arsonist burns down half of the popsicle factories in Hawaii.
- 11.4 For each item, indicate the likely impact on the supply and demand for codfish. Then indicate the effect on the equilibrium price and quantity. It may help to use a graph.
- (a) News reports that cod contains lots of omega-3 fatty acids, which are great for your health.
 - (b) Overfishing drastically reduce the fish population.
- 11.5 For each item, indicate the likely impact on the supply and demand for paperback books. Then indicate the effect on the equilibrium price and quantity. It may help to use a graph.
- (a) The invention (and widespread use) of the printing press.
 - (b) The invention (and widespread use) of the television.
 - (c) The invention (and widespread use) of "book lights" (the small clip-on lights that allow people to read at night without disturbing their spouses/partners/etc.)
 - (d) News reports that reading books is a cure for stress and high blood pressure.
 - (e) A decrease in the price of paper.
- 11.6 For each item, indicate the likely impact on the supply and demand for bulldozer operators and other skilled construction workers. (It may help to think for a moment about who the suppliers and demanders are for these services.) Then indicate the effect on the equilibrium price and quantity. It may help to use a graph.

- (a) The elimination of vocational programs that teach people how to use bulldozers.
- (b) A huge increase in the number of well-paying service-sector jobs such as computer programming.
- (c) A fall in the price of bulldozers and other construction equipment. (To state the obvious: bulldozers and bulldozer operators are complements, like bread and butter or computers and monitors.)
- (d) An increase in the wage for unskilled laborers. (To state the less obvious: skilled labor (e.g., workers who can use bulldozers) and unskilled labor (e.g., workers who can only use shovels) are substitutes, as are tea and coffee and planes, trains, and automobiles.)

11.7 Read the following excerpt from the *New York Times* of October 5, 2000.

The energy proposals that Mr. Bush, the Republican presidential candidate, brought out last week—including opening part of the Arctic National Wildlife Refuge to exploration and incentives to promote coal and nuclear power—could test the willingness of Americans to rebalance environmental and energy priorities in the face of higher prices. For his part, Vice President Al Gore, the Democratic presidential candidate, favors investments in mass transit and incentives to encourage the use of more fuel-efficient vehicles and alternative energy sources.

The “energy crisis” was a big topic in the presidential race. (It might be interesting to investigate how the real price of gasoline has changed over the last 30 or so years.) For each item, indicate the likely impact on the supply and demand for oil. Then indicate the effect on the equilibrium price and quantity. It might help to use a graph. Please note that, in addition to being refined to make gasoline for cars, oil is also used to heat homes and to produce electricity; coal and nuclear power are also used to produce electricity.

- (a) Opening part of the Arctic National Wildlife Refuge to oil exploration.
- (b) Government incentives to promote coal and nuclear power.
- (c) Government investments in mass transit.
- (d) Government incentives to encourage the use of solar-powered vehicles.
- (e) Will all of these policies reduce the price of oil? Yes No (Circle one)
- (f) Will all of these policies reduce the consumption of oil? Yes No (Circle one)
- (g) Is it correct that Bush’s proposals all address the supply side of the problem?

- (h) Is it correct that Gore's proposals all address the demand side of the problem?

11.8 **Let's look a little more closely** at one of now-President Bush's energy proposals: opening up the Arctic National Wildlife Refuge (ANWR) to oil drilling.

When you answered the previous question, you probably assumed that that oil would become available immediately, i.e., that oil companies could immediately begin extracting and selling that oil. (I wanted you to assume that, so do not go back and rethink your answer above!) It turns out that life is more complicated than that: it takes time to build pipelines and to drill oil wells into pristine Arctic wilderness, so any oil that comes from ANWR will not reach the market for something like 5 years. This fact became a source of contention during the presidential campaign, with Al Gore arguing that opening ANWR would have no effect on current gasoline prices because of this 5-year time lag, and George W. Bush arguing. . . well, I don't remember what his argument was, but it probably had something to do with how when people take actions there have to be consequences.

Unlike the majority of the American public, you now understand how supply and demand works, and you should be able to assess the validity of Al's argument. You should try to do this on your own; otherwise (or once you try it on your own), the questions below can serve to guide and/or confirm your thinking.

- (a) Think ahead five years into the future (to 2006), when that oil from ANWR will finally reach the market. Indicate the effect this will have on the market for oil five years from now. (You should draw a supply and demand graph.)
- (b) Next: looking at your graph, what is the effect on the market price for oil in 2006? Will it be higher, lower, or the same?
- (c) Next: Come back to the year 2001. We need to figure out the impact of that future price change on the market for oil today. So: imagine that you own a bunch of oil. You're trying to decide whether to invest in the bank (by extracting and selling the oil and putting the money in the bank) or to "invest in the oil" (by leaving the oil in the ground until, say, 2006). Does your answer to the previous question make investing in oil look more attractive or less attractive?
- (d) Next: As a result, are you likely to sell more oil this year or less oil?
- (e) Finally, think about what this means in terms of your individual supply curve, and remember that all the oil companies are thinking just like you. So: use a supply and demand graph to determine the effect on oil prices today of opening up ANWR for oil drilling. Does today's price go up, down, or stay the same?

Answers

- 11.1 This is the price at which the amount that buyers want to buy equals the amount that sellers want to sell. At a higher price, sellers want to sell more than buyers want to buy, creating incentives that push prices down toward the equilibrium. At a lower price, buyers want to buy more than sellers want to sell, creating incentives that push prices up toward the equilibrium.
- 11.2 [“...supply and demand for wine...”]
- (a) Demand increases. Equilibrium price up, quantity up.
 - (b) Supply decreases. Equilibrium price up, quantity down.
 - (c) Demand decreases. Equilibrium price down, quantity down.
 - (d) Demand decreases. Equilibrium price down, quantity down.
- 11.3 [“...supply and demand for popsicles...”]
- (a) Demand increases. Equilibrium price up, quantity up.
 - (b) Supply decreases. Equilibrium price up, quantity down.
- 11.4 [“...supply and demand for codfish...”]
- (a) Demand increases. Equilibrium price up, quantity up.
 - (b) Supply decreases. Equilibrium price up, quantity down.
- 11.5 [“...supply and demand for paperback books...”]
- (a) Supply increases. Equilibrium price down, quantity up.
 - (b) Demand decreases. Equilibrium price down, quantity down.
 - (c) Demand increases. Equilibrium price up, quantity up.
 - (d) Demand increases. Equilibrium price up, quantity up.
 - (e) Supply increases. Equilibrium price down, quantity up.
- 11.6 [“...supply and demand for bulldozer operators...”]
- (a) Supply decreases. Equilibrium price up, quantity down.
 - (b) Supply decreases. Equilibrium price up, quantity down.
 - (c) Demand increases. Equilibrium price up, quantity up.
 - (d) Demand increases. Equilibrium price up, quantity up.
- 11.7 [“...supply and demand for oil...”]
- (a) Supply increases. Equilibrium price down, quantity up.

- (b) Demand decreases. Equilibrium price down, quantity down.
- (c) Demand decreases. Equilibrium price down, quantity down.
- (d) Demand decreases. Equilibrium price down, quantity down.
- (e) Yes.
- (f) No. Opening the Arctic National Wildlife Refuge would increase the equilibrium quantity.
- (g) No. Promoting substitutes to oil (e.g., coal and nuclear power) is a demand-side strategy.
- (h) Yes.

11.8 [“...drilling in the future...”]

- (a) Supply increases. Equilibrium price down, quantity up.
- (b) Lower.
- (c) Less attractive.
- (d) More oil.
- (e) Supply increases. Equilibrium price down, quantity up.

Chapter 12

Taxes

We're going to devote considerable energy to the analysis of taxes, both because of their importance in real life and because their impact on the market can be quantified: we can figure out *exactly* what happens to the supply and demand curves. The example we're going to use is the hypothetical market for gasoline shown in Figure 12.1.

12.1 Per-unit taxes on the sellers

Imagine that the government imposes a **per-unit tax** of \$0.40 per gallon on the sellers of gasoline. One thing is obvious: the demand curve for gasoline is not going to change because of a tax on the sellers. So if p and q are going to change—and your intuition hopefully suggests that they should—then something has to happen to the supply curve.

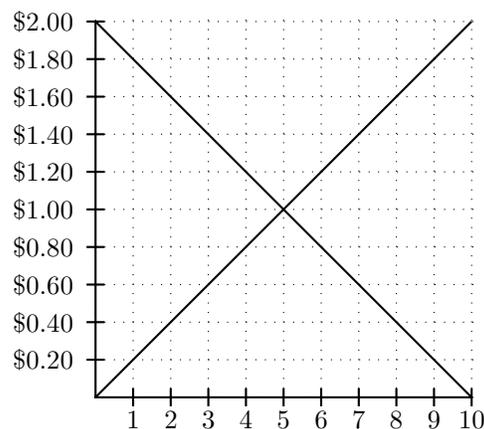


Figure 12.1: A hypothetical market for gasoline, with price in dollars per gallon and quantity in millions of gallons per day.

Consider how much sellers want to sell at a price of \$1.40. Figure 12.1 indicates that, without a tax, sellers want to sell 7 million gallons of gasoline at that price. What we're trying to figure out is how much the sellers want to sell at a price of \$1.40 when there's a tax of \$0.40 per gallon.

But we can figure this out from the supply curve! If the sellers receive \$1.40 per gallon but have to pay a tax of \$0.40 per gallon, what they end up with is \$1.00 per gallon. So the sellers should want to sell exactly as much at \$1.40 per gallon with a \$0.40 tax as they wanted to sell at \$1.00 per gallon without the tax. That amount, as we can see from Figure 12.1, is 5 million gallons.

The same logic applies along the entire supply curve. With a tax of \$0.40, sellers should want to sell exactly as much at a price of x as they wanted to sell at a price of $(x - .40)$ without the tax. The resulting shift in the market supply curve is shown in Figure 12.2. To see the effect on the market equilibrium, we now look to see where the new supply curve intersects the demand curve. (Recall that nothing changes on the demand side because the tax is on the sellers). Before the tax, the market-clearing price was \$1.00 per gallon; after the tax, the market-clearing price is \$1.20 per gallon.

This result has a number of curious features:

- Even though the amount of the tax is \$0.40, the market price doesn't change by \$0.40. The market price only changes by \$0.20.
- Even though the \$0.40 per gallon tax is levied on the sellers, the sellers do not end up worse off by \$0.40 per gallon. The sellers receive \$1.20 per gallon from the buyers, so after paying the tax of \$0.40 per gallon the sellers end up with \$0.80 per gallon. This is only \$0.20 per gallon less than the \$1.00 per gallon that they received before the tax.
- Even though the tax is levied on the sellers, the buyers do not come away unscathed. Instead of paying \$1.00 per gallon, the buyers are paying

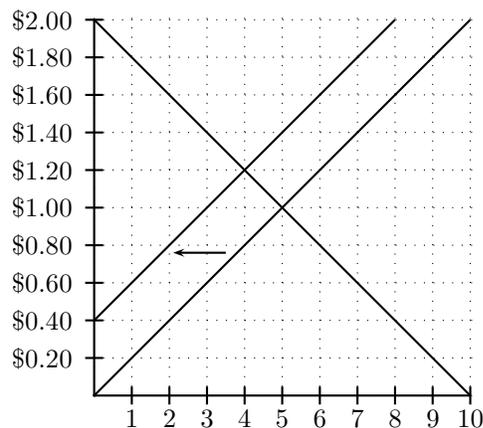


Figure 12.2: The effect of a \$0.40 per gallon tax on the sellers

\$1.20 per gallon; this is \$0.20 more per gallon than they paid before the imposition of the tax.

- Figure 12.2 shows that a per-unit tax of \$0.40 on the sellers shifts the market supply curve *up* by \$0.40, the amount of the tax. Although we have emphasized that supply and demand curves shift left or right—not up or down—it turns out that this is not a coincidence. We will return to this issue in Chapter 13.

Sales taxes and subsidies

The logic described above applies in a wide variety of situations. To analyze **ad valorem taxes** on the sellers—e.g., a 10% sales tax on sellers—note that at a price of \$1.00, sellers pay \$0.10 in tax, so they should want to sell the same quantity they wanted to sell at a price of \$0.90 without the tax. To analyze **subsidies**—either per-unit or ad valorem—note that subsidies are just negative taxes. For example, with a per-unit subsidy of \$0.40, sellers should want to sell at a market price of \$1.00 the same amount that they wanted to sell at a price of \$1.40 without the subsidy.

12.2 Per-unit taxes on the buyers

Let's do the \$0.40 per gallon gasoline tax again, but this time with the tax on the buyers. A key point here is this: *The market price p is what the buyer pays the seller.* So the buyer pays the seller the market price p and then pays the government the tax on top of that.

With that in mind, imagine the government imposes a \$0.40 per gallon tax on the buyers of gasoline. The supply curve for gasoline is not going to change because of a tax on the buyers, so if p and q are going to change, something has to happen to the demand curve.

Consider how much buyers want to buy at a price of \$1.40. Figure 12.1 indicates that, without a tax, buyers want to buy 3 million gallons of gasoline at that price. What we're trying to figure out is how much the buyers want to buy at a price of \$1.40 when there's a tax of \$0.40 per gallon.

This time we can figure it out by looking at the demand curve. If the buyers pay the sellers \$1.40 per gallon but have to pay an additional tax of \$0.40 per gallon, what they end up paying is \$1.80 per gallon. So the buyers should want to buy exactly as much at \$1.40 per gallon with a \$0.40 tax as they wanted to buy at \$1.80 per gallon without the tax. That amount, as we can see from Figure 12.1, is 1 million gallons.

The same logic applies along the entire demand curve. With a tax of \$0.40, buyers should want to buy exactly as much at a price of $\$x$ as they wanted to buy at a price of $\$(x + .40)$ without the tax. The resulting shift in the market demand curve is shown in Figure 12.3. To see the effect on the market equilibrium, we now simply add a supply curve. (Recall that nothing changes

on the supply side because the tax is on the buyers.) Before the tax, the market-clearing price was \$1.00 per gallon; after the tax, the market-clearing price is \$0.80 per gallon.

This result has curious features that parallel those on the supply side:

- Even though the amount of the tax is \$0.40, the market price doesn't change by \$0.40. The market price only changes by \$0.20.
- Even though the \$0.40 per gallon tax is levied on the buyers, the buyers do not end up worse off by \$0.40 per gallon. The buyers pay \$0.80 per gallon to the sellers and \$0.40 per gallon in taxes, so in the end the buyers end up paying \$1.20 per gallon. This is only \$0.20 per gallon more than the \$1.00 per gallon that they paid before the tax.
- Even though the tax is levied on the buyers, the sellers do not come away unscathed. Instead of receiving \$1.00 per gallon, the sellers receive only \$0.80 per gallon; this is \$0.20 less per gallon than they received before the imposition of the tax.
- Figure 12.3 shows that a per-unit tax of \$0.40 on the buyers shifts the market demand curve *down* by \$0.40, the amount of the tax. Again, this is not a coincidence, and we will return to this issue in Chapter 13.

Sales taxes and subsidies

As on the supply side, the logic described above applies to sales taxes on and subsidies for buyers. To analyze a 10% sales tax on buyers, note that at a price of \$1.00, buyers pay \$0.10 in tax to the government on top of the \$1.00 they pay the sellers, so they should want to buy the same quantity they wanted to buy at a price of \$1.10 without the tax. For a per-unit subsidy of \$0.40, buyers should

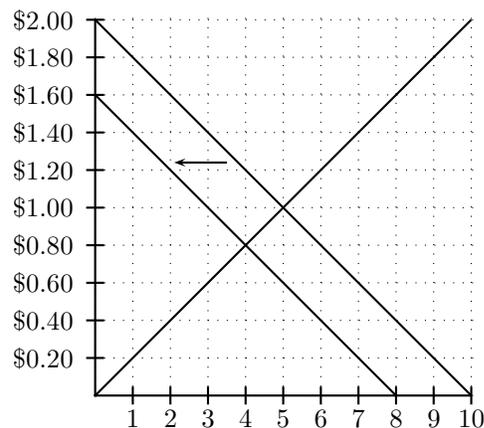


Figure 12.3: Effect of a \$0.40 per gallon tax on the buyers

want to buy at a market price of \$1.00 the same amount that they wanted to buy at a price of \$0.60 without the subsidy.

12.3 Tax equivalence

The curious results above lead to two important questions: What determines who bears the ultimate economic burden of the tax? And which is better for buyers (or for sellers), a tax on buyers or a tax on sellers?

To answer these questions, recall that the original (pre-tax) equilibrium was at 5 million gallons, with buyers paying \$1.00 per gallon and sellers receiving \$1.00 per gallon. Here is a summary of the results for a \$0.40 per-unit tax:

When the tax is on the sellers The equilibrium quantity falls to 4 million gallons and the market price rises from \$1.00 to \$1.20 per gallon. After the imposition of the tax, the buyers pay \$1.20 for each gallon of gasoline. The sellers receive \$1.20 per gallon, but then pay \$0.40 in tax, so what the sellers really get is \$0.80 per gallon.

When the tax is on the buyers The equilibrium quantity falls to 4 million gallons and the market price falls from \$1.00 to \$0.80 per gallon. After the imposition of the tax, the sellers receive \$0.80 for each gallon of gasoline. The buyers pay \$0.80 per gallon to the seller, plus \$0.40 to the government, so what the buyers really pay is \$1.20 per gallon.

We can conclude that the impact of the tax is the same regardless of whether it's on the buyers or on the sellers! When the tax is on the sellers, the sellers push some of the tax onto the buyers; when the tax is on the buyers, the buyers push some of the tax onto the sellers. The ultimate result of this battle is independent of who does the pushing: in both cases the economic impact of the tax is shared by the two parties. Shifting the **legal incidence** of the tax from the buyer to the seller—e.g., removing a tax on the buyer and imposing a similar tax on the seller—has no effect on the **economic incidence** of the tax. This is the **tax equivalence** result.

12.4 Tax incidence

In the example above, the buyers and sellers end up with equal shares of the economic incidence of the tax: buyers end up paying \$1.20 per gallon, \$0.20 more than before, and sellers end up getting \$0.80 per gallon, \$0.20 less than before.¹ Do buyers and the sellers always share the tax burden equally?

¹Note that the *sum* of the buyers' and sellers' tax burdens always equals the amount of the tax. In the examples above, the buyers and sellers combined are worse off by \$0.40 per gallon, which is the amount of the tax. Confirming this summation result is a good way to double-check your work. It should also make intuitive sense; after all, the government gets \$0.40 per gallon, and that money has to come from somewhere.

The answer is No. The guiding principle turns out to be that *taxes are paid by those who are least able to avoid them*. The basic idea—to be formalized in the next chapter—is that the distribution of the tax burden depends on the sensitivity of buyers and sellers to price changes. In the example above, buyers and sellers are equally sensitive to price changes, so they share the tax burden equally.

More generally, whichever party—buyers or sellers—is more sensitive to price changes will bear *less* of the tax burden, and whichever party is less sensitive to price changes will bear *more* of the tax burden. When a tax is imposed, the party with the stronger sensitivity to price changes can essentially say, “Hey, if you put the tax burden on me, I’ll cut back by a lot!” The party with the weaker sensitivity to price changes can only say, “Well, if you put the tax burden on me, I’ll cut back by a little.” Since the new equilibrium has to equalize the amount that buyers want to buy with the amounts that sellers want to sell, the party with the strongest sensitivity to price changes effectively pushes the lion’s share of the tax burden onto the other side.

12.5 Algebra: The algebra of taxes

The key fact in the algebraic treatment of taxes is that the market price is the amount the buyer pays the seller. So if the market price is p and there’s a tax on the sellers of \$1 per unit, what sellers really get is only $p - 1$. For example, if the supply curve before the imposition of the tax is $q = 6000 + 800p$, the supply curve after the imposition of a \$1 per-unit tax will be

$$q = 6000 + 800(p - 1) \implies q = 5200 + 800p.$$

Similarly, if there’s a tax on the sellers of 20% and the market price is p , what sellers really get is $100 - 20 = 80\%$ of p , i.e., $0.8p$; if the before-tax supply curve is $q = 6000 + 800p$, the after-tax supply curve will be

$$q = 6000 + 800(.8p) \implies q = 6000 + 640p.$$

The same logic applies to taxes on the buyers. If the market price is p and there’s a tax on the buyers of \$1 per unit, what buyers really have to pay is $p + 1$. (Remember that the market price is only what the buyer pays the seller—the buyer’s tax payment is extra!) For example, if the demand curve before the imposition of the tax is $q = 7500 - 500p$, the demand curve after the imposition of a \$1 per-unit tax will be

$$q = 7500 - 500(p + 1) \implies q = 7000 - 500p.$$

Similarly, if there’s a tax on the buyers of 20% and the market price is p , what buyers really have to pay is $1.2p$ (120% of p). If the before-tax demand curve is $q = 7500 - 500p$, the after-tax demand curve will be

$$q = 7500 - 500(1.2p) \implies q = 7500 - 600p.$$

Problems

Answers are in the next section. The first few words of each question (and each answer) are hyperlinked to allow online users to quickly move back and forth.

- 12.1 Explain, as if to a mathematically literate non-economist, why taxes shift the supply and/or demand curves the way they do. (The same answer works for sales taxes and per unit taxes, as well as for taxes on buyers and taxes on sellers.)
- 12.2 Consider the hypothetical market for oranges shown in Figure 12.4. Use it (and the replicas on the following pages) to answer the questions in the remaining problems in this chapter.
- What is the equilibrium price and quantity? (Use correct units!)
 - Calculate the slope of the supply curve and the slope of the demand curve. (Recall that slope is rise over run, e.g., $S_D = \frac{\Delta p}{\Delta q}$.) Calculate the ratio of the slopes $\left(\frac{S_D}{S_S}\right)$.
- 12.3 Suppose that the government imposes an excise tax of \$0.60 per pound on the sellers of oranges.
- Show the impact of this tax on the supply and demand curves.
 - At the new equilibrium, how many oranges will people eat? (Use correct units!)

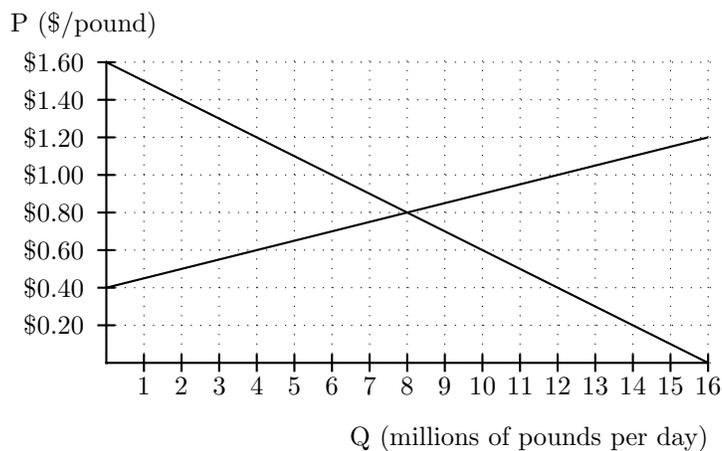


Figure 12.4: A hypothetical market for oranges

- (c) Calculate the total tax revenue for the government from this tax. (Use correct units!)
- (d) How much do the buyers pay for each pound of oranges?
- (e) How much after-tax revenue do the sellers receive for each pound of oranges?
- (f) Compare the tax burden ratio with the ratio of the slopes from problem 12.2b. Can you explain the intuition behind this result?
- 12.4 Answer the same questions as in problem 12.3, but now suppose that the government imposes a per-unit tax of \$0.60 per pound on the buyers of oranges. (Recall that the buyer now has to pay the government *in addition to* paying the seller.)
- 12.5 How do your answers in problem 12.4 compare with those in problem 12.3? What does this suggest about the difference between a per-unit tax on buyers and a per-unit tax on sellers?
- 12.6 Answer the same questions as in problem 12.3, but now suppose that the government imposes a sales tax of 50% on the sellers of oranges. (With a sales tax, if sellers sell a pound of oranges for \$1, they get to keep \$0.50 and have to pay the government \$0.50; if they sell a pound of oranges for \$2, they get to keep \$1 and have to pay the government \$1.)
- 12.7 Answer the same questions as in problem 12.3, but now suppose that the government imposes a sales tax of 100% on the buyers of oranges. (If buyers buy a pound of oranges for \$1, they have to pay the seller \$1 and the government \$1; if they buy a pound of oranges for \$2, they have to pay the seller \$2 and the government \$2.)

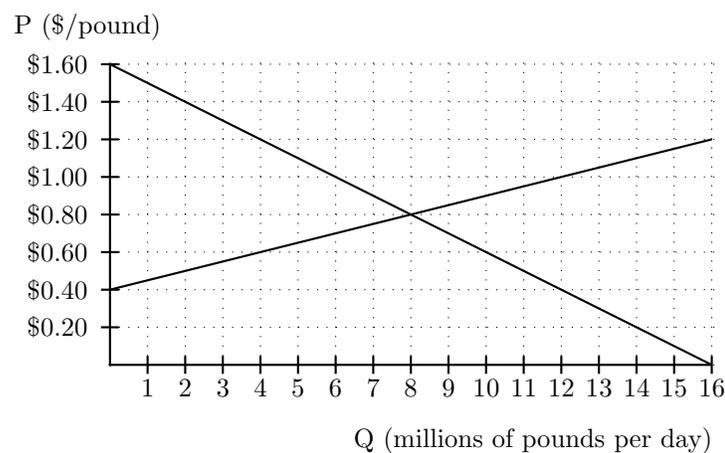


Figure 12.5: The same hypothetical market for oranges

- 12.8 How do your answers in problem 12.7 compare with those in problem 12.6? What does this suggest about the difference between a sales tax on buyers and a sales tax on sellers?

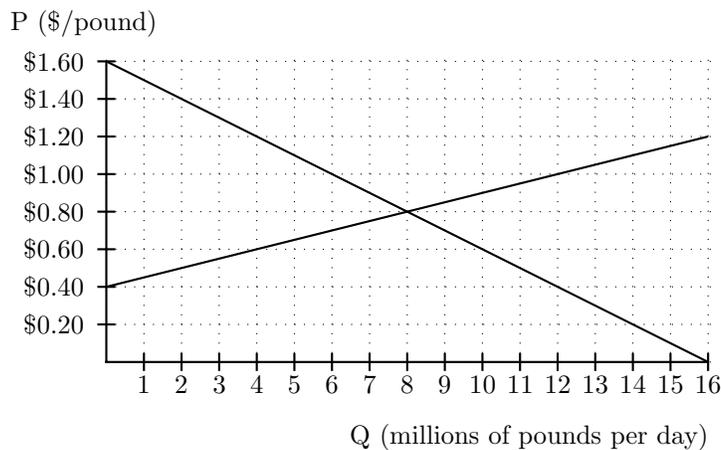


Figure 12.6: The same hypothetical market for oranges

Algebra problems

- 12.9 Consider a market with demand curve $q = 16 - 10p$ and supply curve $q = -8 + 20p$. (Here q is in millions of pounds per day and p is in dollars per pound.)
- Determine the market equilibrium price and quantity and the total revenue in this market.
 - Calculate the price elasticity of demand and the price elasticity of supply at the market equilibrium.
 - Imagine that the government imposes a \$0.60 per-unit tax on the buyers. Write down the new market supply and demand curves, and find the new market equilibrium price and quantity. How much of the tax burden is borne by the buyers, and how much by the sellers? Calculate the ratio of these tax burdens, and compare with the ratio of the elasticities calculated above.
 - Now imagine that the government instead decides to impose a \$0.60 per-unit tax on the sellers. How will this change things? Write down the new market supply and demand curves, find the new market equilibrium price and quantity, and compare with your answer from above (where the tax is on the buyer).
 - Now imagine that the government instead decides to impose a sales tax of 50% on the sellers. How will this change things? Write down

the new market supply and demand curves and find the new market equilibrium price and quantity.

- (f) Now imagine that the government instead decides to impose a 100% sales tax on the buyers. How will this change things? Write down the new market supply and demand curves, find the new market equilibrium price and quantity, and compare with your answer from above (where the tax is on the sellers).
- 12.10 Consider a market with demand curve $q = 500 - 20p$ and supply curve $q = 50 + 25p$.
- What is the original market equilibrium price and quantity?
 - What do the demand and supply curves look like with a \$2 per-unit tax on the buyers?
 - With a \$2 per-unit tax on the sellers?
 - With a 20% sales tax on the buyers?
 - With a 20% sales tax on the sellers?
- 12.11 Consider a world with market demand curve $q = 50 - 6p$ and market supply curve $q = 20p - 28$.
- What is the market equilibrium price and quantity?
 - How would the equations for the supply and demand curves change if the government imposed a tax of \$0.50 per unit on the buyers?
 - How would the equations for the supply and demand curves change if the government imposed a sales tax of 20% on the sellers?

Answers

- 12.1 This is the logic identified in the text, e.g., "At a price of x with a \$0.40 tax, buyers should be willing to buy exactly as much as they were willing to buy at a price of $\$(x + .40)$ without the tax." For per-unit taxes, you can also use the ideas of marginal cost and marginal benefit: a tax on the sellers increases marginal costs by the amount of the tax, and a tax on the buyers reduces marginal benefits by the amount of the tax. (Applying this marginal approach is a bit tricky for ad valorem taxes such as sales taxes. You need an additional assumption here about firm profits in equilibrium. . . .)
- 12.2 ["... a hypothetical market for oranges..."]
- The equilibrium price is \$0.80 per pound; the equilibrium quantity is 8 million pounds per day.

- (b) To find the slope of the supply curve, pick any two points on it, say, (8, .80) and (12, 1.00). Then the supply curve's slope is

$$S_S = \frac{\text{rise}}{\text{run}} = \frac{1.00 - .80}{12 - 8} = \frac{.20}{4} = .05.$$

Similarly, to find the slope of the demand curve, pick any two points on it—say (8, .80) and (12, .40). Then the slope of the demand curve is

$$S_D = \frac{\text{rise}}{\text{run}} = \frac{.40 - .80}{12 - 8} = \frac{-.40}{4} = -.1.$$

So the ratio of the two slopes is $\left(\frac{S_D}{S_S}\right) = \frac{-.1}{.05} = -2$.

12.3 [“Suppose that the government...”]

- (a) The supply curve shifts up by \$0.60. See Figure 12.7.

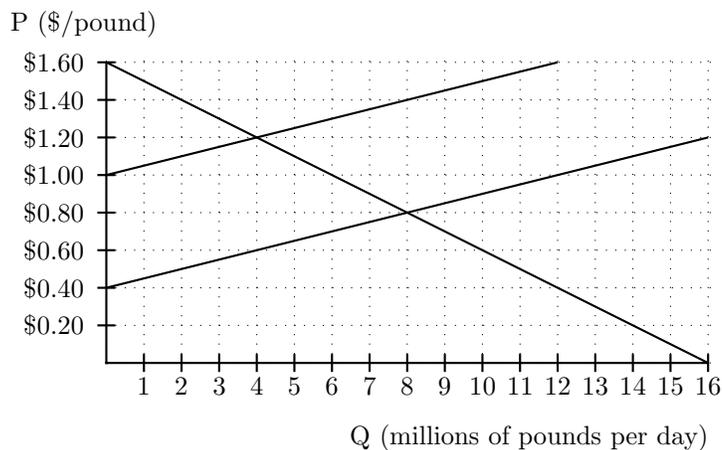


Figure 12.7: The hypothetical market for oranges

- (b) The new equilibrium features a price of \$1.20 per pound and a quantity of 4 million pounds per day.
- (c) A tax of \$0.60 per pound levied on 4 million pounds per day yields revenue of $(\$0.60)(4) = \2.4 million per day.
- (d) The buyers pay \$1.20 for each pound of oranges.
- (e) Before paying the \$0.60 tax the sellers receive \$1.20 per pound, so after paying the tax the sellers receive \$0.60 per pound.
- (f) Without the tax, buyers paid \$0.80 per pound; they now pay \$1.20 per pound, so they are worse off by $T_B = \$0.40$ per pound. Similarly, the sellers received \$0.80 per pound without the tax, but now they

only receive \$0.60 per pound, so they are worse off by $T_S = \$0.20$ per pound. (As a check here, note that the \$0.40 per pound impact on the buyers plus the \$0.20 per pound impact on the sellers equals the \$0.60 per pound tax.) The ratio of the tax burdens is $\frac{T_B}{T_S} = \frac{.40}{.20} = 2$.

- (g) The tax burden ratio is the same magnitude as the ratio of the slopes calculated previously! Intuitively, this is because the ratio of the slopes measures the relative responsiveness of buyers and sellers to price changes. The side that is most responsive to price changes (in this case, the sellers) can push the lion's share of the tax burden onto the other side.

12.4 See the answer in Figure 12.8.

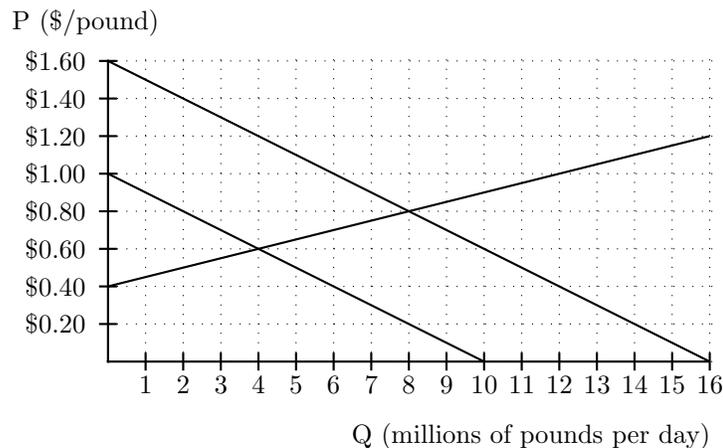


Figure 12.8: The hypothetical market for oranges

- (a) The demand curve shifts down by \$0.60.
 (b) The new equilibrium features a price of \$0.60 per pound and a quantity of 4 million pounds per day.
 (c) A tax of \$0.60 per pound levied on 4 million pounds per day yields revenue of $(\$0.60)(4) = \2.4 million per day.
 (d) The buyers pay \$0.60 to the sellers for each pound, plus the tax of \$0.60 to the government, so they pay a total of \$1.20 per pound.
 (e) The sellers receive \$0.60 per pound of oranges.
 (f) Without the tax, buyers paid \$0.80 per pound; they now pay \$1.20 per pound, so they are worse off by $T_B = \$0.40$ per pound. Similarly, the sellers received \$0.80 per pound without the tax, but now they only receive \$0.60 per pound, so they are worse off by $T_S = \$0.20$ per pound. (As a check here, note that the \$0.40 per pound impact on

the buyers plus the \$0.20 per pound impact on the sellers equals the \$0.60 per pound tax.) The ratio of the tax burdens is $\frac{T_B}{T_S} = \frac{.40}{.20} = 2$.

- (g) The tax burden ratio is the same magnitude as the ratio of the slopes calculated previously!

12.5 The answers are essentially identical—regardless of the *legal* incidence of the tax (i.e., whether it's levied on the buyers or the sellers), the *economic* incidence of the tax comes out the same, i.e., the buyers always end up bearing \$0.40 of the tax burden and the sellers end up bearing \$0.20 of the tax burden. This is an example of the **tax incidence** result.

12.6 The supply curve rotates as shown in Figure 12.9. The remaining answers are identical to those above; note that at the equilibrium price of \$1.20 per pound, the 50% tax on the sellers amounts to \$0.60 per pound.

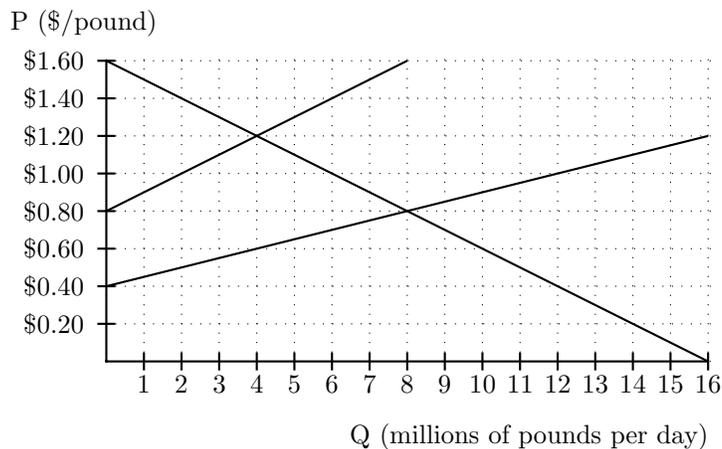


Figure 12.9: The hypothetical market for oranges

12.7 The demand curve rotates as shown in Figure 12.10. The remaining answers are identical to those above; note that at the equilibrium price of \$0.60 per pound, the 100% tax on the buyers amounts to \$0.60 per pound.

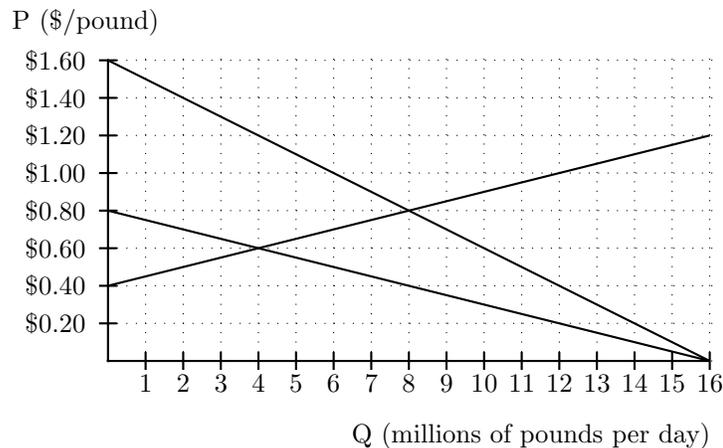


Figure 12.10: The hypothetical market for oranges

12.8 **The answers are the same!** Again, this is an example of tax equivalence, but with sales taxes the comparison is a little bit less obvious than with per-unit taxes: in the case of sales taxes, a 50% tax on the sellers is equivalent to a 100% tax on the buyers. This is because the market price is the price the buyer pays the seller; since the market price is twice as high when the tax is on the seller (\$1.20 versus \$0.60), the sales tax rate on the sellers needs to be only half as high as the sales tax rate on the buyers in order to yield an equivalent result.

12.9 [“Consider a market...”]

- Simultaneously solving the demand and supply equations gives us a market equilibrium of $p = \$0.80$ per pound and $q = 8$ million pounds per day. Total revenue is therefore \$6.4 million per day.
- The slope of the demand curve is -10 , so the price elasticity of demand at the market equilibrium is $-10 \cdot \frac{.8}{8} = -1$. The slope of the supply curve is 20 , so the price elasticity of supply at the market equilibrium is $20 \cdot \frac{.8}{8} = 2$.
- A tax of \$0.60 on the buyers will change the demand curve to $q = 16 - 10(p + .6)$, i.e., $q = 10 - 10p$. The supply curve is still $q = -8 + 20p$, so the new market equilibrium is at $p = \$0.60$ per pound and $q = 4$ million pounds per day. The buyers end up paying $0.60 + 0.60 = \$1.20$ per pound, and the sellers get \$0.60 per pound. The original equilibrium was at a price of \$0.80 per pound, so the buyers end up paying \$0.40 more and the sellers end up getting \$0.20 less. The ratio of these tax burdens is $\frac{.40}{.20} = 2$. This is the negative inverse of the ratio of the elasticities.

- (d) A tax of \$0.60 on the sellers will change the supply curve to $q = -8 + 20(p - .6)$, i.e., $q = -20 + 20p$. The demand curve is still $q = 16 - 10p$, so the new market equilibrium is at $p = \$1.20$ per pound and $q = 4$ million pounds per day. The buyers end up paying \$1.20 per pound, and the sellers get $1.20 - .60 = \$0.60$ per pound. This is the same result as above, demonstrating the tax equivalence result.
- (e) A tax of 50% on the sellers will change the supply curve to $q = -8 + 20(.5p)$, i.e., $q = -8 + 10p$. The demand curve is still $q = 16 - 10p$, so the new market equilibrium is at $p = \$1.20$ per pound and $q = 4$ million pounds per day. The buyers end up paying \$1.20 per pound, and the sellers get $.5(1.20) = \$0.60$ per pound.
- (f) A tax of 100% on the buyers will change the demand curve to $q = 16 - 10(2p)$, i.e., $q = 16 - 20p$. The supply curve is still $q = -8 + 20p$, so the new market equilibrium is at $p = \$0.60$ per pound and $q = 4$ million pounds per day. The buyers end up paying $2(.60) = \$1.20$ per pound, and the sellers get \$0.60 per pound. This is the sales tax version of the tax equivalence result.

12.10 [“Consider a market...”]

- (a) Solving the two equations simultaneously we have $500 - 20p = 50 + 25p$, which simplifies to $45p = 450$ or $p = 10$. Plugging this back in to either of the two original equations yields $q = 300$.
- (b) The supply equation is unchanged and the demand equation becomes $q = 500 - 20(p + 2)$, i.e., $q = 460 - 20p$. Solve this equation and the supply equation simultaneously to get the new equilibrium price and quantity.
- (c) The demand equation is unchanged and the supply equation becomes $q = 50 + 25(p - 2)$, i.e., $q = 25p$. Solve this equation and the demand equation simultaneously to get the new equilibrium price and quantity.
- (d) The supply equation is unchanged and the demand equation becomes $q = 500 - 20(1.2p)$, i.e., $q = 500 - 24p$. Solve this equation and the supply equation simultaneously to get the new equilibrium price and quantity.
- (e) The demand equation is unchanged and the supply equation becomes $q = 50 + 25(.8p)$, i.e., $q = 50 + 20p$. Solve this equation and the demand equation simultaneously to get the new equilibrium price and quantity.

12.11 [“Consider a world...”]

- (a) Solving simultaneously we get $50 - 6p = 20p - 28$, which yields $p = 3$. Plugging this into either the market demand curve or the market supply curve yields $q = 32$.

(b) The supply curve is unchanged. The demand curve becomes

$$q = 50 - 6(p + .5).$$

(c) The demand curve is unchanged. The supply curve becomes

$$q = 20(.8p) - 28.$$

Chapter 13

Margins

When we draw supply and demand graphs, we put the quantity Q on the x -axis and the price P on the y -axis. This is odd because we spend all our time talking about Q as a function of P , i.e., asking questions like, “At a price of \$5, how many units will buyers want to buy?” It would be more natural to put P on the x -axis and Q on the y -axis, so why don’t economists do it that way?

One answer is that it was a mistake, a typographical error in the first economics textbook ever written—Alfred Marshall’s *Principles of Economics* from 1890—that led to our being “locked in” to this weird way of looking at supply and demand in the same way that we’re locked into using the particular layout of a computer keyboard. But another answer—and they could both be true, but this second one is definitely true—is that we can reinterpret supply and demand curves to look at P as a function of Q by asking about **marginal benefit** and **marginal cost**.

13.1 Reinterpreting the supply curve

Consider a firm whose supply curve indicates that at any price above \$5 per unit the firm wants to sell at least 10 units and that at any price below \$5 per unit the firm wants to sell at most 9 units. We can now ask the following question: What is the firm’s *marginal cost* of producing the 10th unit? (**Marginal cost** is the cost of producing an additional unit. The marginal cost of the 10th unit is the cost differential between producing 10 units and producing only 9 units.)

The answer follows from profit maximization: the firm’s marginal cost of producing the 10th unit must be \$5. Here’s why:

- Since the firm wants to sell at least 10 units at any price above \$5 per unit (say, \$5.10), the marginal cost of the 10th unit cannot be more than \$5. If the marginal cost of the 10th unit was, say, \$6, the firm wouldn’t be maximizing profits by selling 10 units at a market price of \$5.10 per unit; it would make \$.90 more profit by selling only 9 units instead of 10.

- Since the firm doesn't want to sell 10 units at any price below \$5 per unit (say, \$4.90), the marginal cost of the 10th unit cannot be less than \$5. If the marginal cost of the 10th unit was, say, \$4, the firm wouldn't be maximizing profits by selling only 9 units at a market price of \$4.90 per unit; it would make \$.90 more profit by selling 10 units instead of 9.
- Since the firm's marginal cost of producing the 10th unit cannot be more than \$5 and cannot be less than \$5, it must be equal to \$5.

It follows that the point on the graph corresponding to a price of \$5 and a quantity of 10 units can be interpreted in two different ways. From the perspective of the **supply curve**, it indicates that at a price of \$5 per unit the firm wants to supply 10 units. From the perspective of the **marginal cost curve**, it indicates that the marginal cost of the 10th unit of production is \$5. This is an example of a more general result: *every point on an individual supply curve is also a point on that individual's marginal cost curve.*¹

The same result applies to the market as a whole: every point on the **market supply curve** is also a point on the **social marginal cost curve**. If the market wants to sell at least 1 million units at any price over \$5 per unit but fewer than 1 million units at any price under \$5 per unit, the marginal cost to society of producing the 1 millionth unit must be \$5.

The close connection between supply curves and marginal cost curves gives us two ways to think about the graph in Figure 13.1. If you start with a price p , an individual supply curve will tell you how many units q that individual seller would want to sell at that price; if you start with a quantity q , an individual marginal cost curve will tell you the marginal cost to that individual seller of producing that q th unit. Similarly, if you start with a price p , a market supply

¹If you take more economics classes, you'll learn that the *converse* of this statement is not true. It turns out that the supply curve is only part of the upward-sloping section of the marginal cost curve, so there are some points on the firm's marginal cost curve that are not on the firm's supply curve.

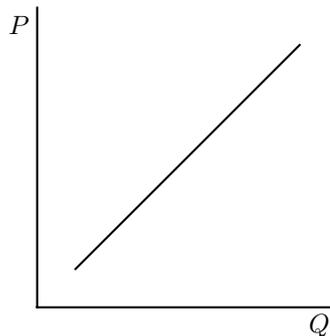


Figure 13.1: Every point on an individual (or market) supply curve is also a point on an individual (or social) marginal cost curve

curve will tell you how many units all the sellers together would want to sell at that price; if you start with a quantity q , the social marginal cost curve will tell you the marginal cost to society of producing that q th unit.

Example: Taxes revisited

Let's take another look at taxes. If the government imposes a tax on the sellers of \$1 per unit, we know from Chapter 12 how to shift the supply curve, e.g., at a price of \$4 per unit, the firms in the market are only going to get to keep \$3 per unit, so at a price of \$4 with a \$1 tax the sellers should be willing to sell exactly what they were willing to sell at \$3 without the tax. But we also saw in Chapter 12 that it looked like the supply curve shifted *up* by \$1. We can't explain this from the supply curve perspective, but if we remember that the supply curve is also the marginal cost curve then the explanation is simple: the marginal cost of producing and selling each additional unit has increased by \$1 because of the tax. So the marginal cost curve shifts up by \$1. Same result, different story!

13.2 Reinterpreting the demand curve

Just as we've refashioned the supply curve as a marginal cost curve, we can refashion the demand curve as a **marginal benefit curve**. Unfortunately, as in the analysis Chapter 20 concerning downward sloping demand curves, it turns out that there are some theoretical complications.² A more advanced course can address those complications; our approach will be to gloss over the theoretical difficulties in order to describe the basic ideas.

Consider an individual whose demand curve indicates that (1) at any price below \$5 per unit she wants to buy at least 10 units; and (2) at any price above \$5 per unit she wants to buy at most 9 units. We can now ask the following question: What is this person's *marginal benefit* from obtaining the 10th unit? (**Marginal benefit** is the benefit from obtaining an additional unit. The marginal benefit of the 10th unit is the difference in benefits between having 10 units and having only 9 units.)

The answer follows from individual optimization: the individual's marginal benefit from obtaining the 10th unit must be \$5. Here's why:

- Since she wants to buy at least 10 units at any price below \$5 per unit (say, \$4.90 per unit), the marginal benefit of the 10th unit cannot be less than \$5. If the marginal benefit of the 10th unit was, say, \$4, she wouldn't be optimizing by buying 10 units at a market price of \$4.90; she would be better off by buying only 9 units instead of 10.
- Since she wants to buy at most 9 units at any price above \$5 per unit (say, \$5.10), the marginal benefit of the 10th unit cannot be more than

²Again, the complications stem from the income effect.

\$5. If the marginal benefit of the 10th unit was, say, \$6, she wouldn't be optimizing if she bought only 9 units at a market price of \$5.10; she would be better off buying 10 units instead of 9.

- Since her marginal benefit from the 10th unit cannot be less than \$5 and cannot be more than \$5, it must be equal to \$5.

It follows that the point on the graph corresponding to a price of \$5 and a quantity of 10 units can be interpreted in two different ways. From the perspective of the **demand curve**, it indicates that at a price of \$5 per unit this individual wants to buy 10 units. From the perspective of the **marginal benefit curve**, it indicates that the marginal benefit of obtaining the 10th unit of production is \$5. This is an example of a more general result: *every point on an individual demand curve is also a point on that individual's marginal benefit curve.*

The same result applies to the market as a whole: every point on the **market demand curve** is also a point on the **social marginal benefit curve**. If the market wants to buy 1 million units at any price below \$5 per unit but fewer than 1 million units at any price above \$5 per unit, the marginal benefit to society of obtaining the 1 millionth unit must be \$5.

The close connection between demand curves and marginal benefit curves gives us two ways to think about the graph in Figure 13.2. If you start with a given p , an individual demand curve will tell you how many units that individual buyer would want to buy at that price; if you start with a quantity q , an individual marginal benefit curve will tell you the marginal benefit to that individual seller of obtaining that q th unit. Similarly, if you start with a price p , a market demand curve will tell you how many units all the buyers together would want to buy at that price; if you start with a quantity q , the social marginal benefit curve will tell you the marginal benefit to society of obtaining that q th unit.

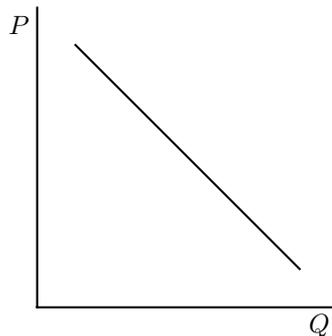


Figure 13.2: Every point on an individual (or market) demand curve is also a point on an individual (or social) marginal benefit curve

Example: Taxes revisited

Let's take another look at taxes. If the government imposes a tax on the buyers of \$1 per unit, we know from Chapter 12 how to shift the demand curve, e.g., at a price of \$4 per unit, the buyers in the market are going to end up paying \$5, so at a price of \$4 with a \$1 tax the buyers should be willing to buy exactly what they were willing to buy at \$5 without the tax. But we also saw in Chapter 12 that it looked like the demand curve shifted *down* by \$1. We can't explain this from the demand curve perspective, but if we remember that the demand curve is also the marginal benefit curve then the explanation is simple: the marginal benefit of obtaining each additional unit has decreased by \$1 because of the tax. So the marginal benefit curve shifts down by \$1. Same result, different story!

13.3 Conclusion: Carrots and sticks

Another way to think about marginal benefit and marginal cost curves is in terms of carrots and sticks. Incentives generally take one of two forms: rewards (i.e., carrots) or punishments (sticks).³ You can think of the marginal benefit curve as the carrot, and the marginal cost curve as the stick. The marginal benefit curve (i.e., the demand curve) provides an incentive for sellers to produce *more* of good X, and the marginal cost curve (i.e., the supply curve) provides an incentive for buyers to consume *less* of good X.

Note that in competitive markets the market price, the marginal cost at the equilibrium quantity, and the marginal benefit at the equilibrium quantity are equal. This fact underlies many of the efficiency results that we will discuss in the next chapter.

³To get a donkey to move forward you can either tie a carrot to a pole and dangle it in front of the donkey, or you can hit the donkey with a stick. To analyze this to death even more, we can use a decision tree to analyze the donkey's choices: go forward or do nothing. To get the donkey to go forward you can either make "Go forward" look attractive—e.g., with a carrot—or you can make "Do nothing" look unattractive—e.g., with a stick. Both are—at least in theory—equally good as incentive mechanisms.

Chapter 14

Elasticity

Let's go back to the demand curve and ask this question: How sensitive are buyers to changes in price? In other words, how much less (or more) would buyers want to buy if the price went up (or down) by a little bit? This question is relevant, for example, in addressing global warming. Increasing the price of gasoline and other fossil fuels linked to global warming is one way to reduce emissions, so many environmentalists—and even more economists—support a **carbon tax** on fossil fuels. But what tax rate is needed to cut gasoline use by 5%, or by 10%? The answer depends in part on how sensitive buyers are to price changes.

14.1 The price elasticity of demand

One way to measure the sensitivity of demand is to just eyeball the slope of the demand curve to see if it's "steep" or "flat". But it turns out that a better way to analyze sensitivity is to put everything in percentages.

The **price elasticity of demand at point A** measures the percentage change in quantity demanded (relative to the quantity demanded at point A) resulting from a 1% increase in the price (relative to the price at point A). For example, if a 1% increase in price (relative to the price at point A) results in a 5% reduction in quantity demanded (relative to the quantity demanded at point A), we say that the price elasticity of demand at point A is -5 .

Demand at point A is **unit elastic** if the percentage change in quantity demanded is *equal* in magnitude to the percentage change in price. In English, this means that a change in price causes a *proportional* change in the quantity demanded. In math, this means that the elasticity of demand is -1 : a 1% increase in price results in a 1% decrease in quantity demanded.

Demand at point A is **elastic**—think of a stretchy rubber band that extends a lot when you pull on it—if the percentage change in quantity demanded is *greater* in magnitude than the percentage change in price. In English, this means that a small change in price causes the quantity demanded to increase

more than proportionally: demand is sensitive to price changes. In math, this means that the elasticity of demand is more negative than -1 , e.g., -2 , in which case a 1% increase in price results in a 2% decrease in quantity demanded. Of occasional theoretic interest is the case of infinite sensitivity to price changes: demand is **perfectly elastic** at point A when the elasticity is $-\infty$.

Demand at point A is **inelastic**—think of a tight rubber band that doesn't move much when you pull on it—if the percentage change in quantity demanded is *smaller* in magnitude than the percentage change in price. In English, this means that a small change in price causes the quantity demanded to increase *less than proportionally*; demand is insensitive to price changes. In math, this means that the elasticity of demand is more positive than -1 , e.g., $-\frac{1}{2}$, in which case a 1% increase in price results in only a $\frac{1}{2}$ % decrease in quantity demanded. Of occasional theoretic interest is the case of complete insensitivity to price changes: demand is **perfectly inelastic** at point A when the elasticity is 0.

Note that elasticities are always measured at a point: even a straight line demand curve will have different price elasticities of demand at different points! Also note that demand elasticities are never greater than zero because demand curves are downward-sloping: a 1% increase in price never causes buyers to want to buy more. Because demand elasticities are always negative, economists often talk about elasticities in terms of magnitude or absolute value. Saying that demand is more elastic at point A than at point B means that the elasticity is more negative at A than at B, e.g., -5 compared to -1 .

To sum up in English If, at point A, a small change in price causes the quantity demanded to increase by a lot, demand at point A is elastic. If quantity demanded only changes by a little, demand at point A is inelastic. And if quantity demanded changes by a proportional amount, demand at point A has unit elasticity.

To sum up in math If, at point A, the price elasticity of demand is greater in magnitude than -1 (e.g., -2), demand at point A is elastic. If the elasticity is lesser in magnitude than -1 (e.g., $-\frac{1}{2}$), demand at point A is inelastic. And if the elasticity is equal to -1 , demand at point A has unit elasticity.

To sum up in pictures See Figure 14.1.

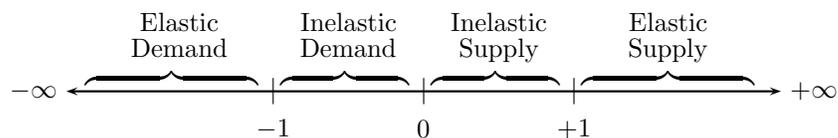


Figure 14.1: Elastic and inelastic demand and supply; “unit elasticity” refers to an elasticity of -1 for demand or $+1$ for supply.

Calculating the price elasticity of demand

To measure the price elasticity of demand at point A, we find some other convenient point B that is near point A and calculate the percentage changes in price and quantity demanded between them¹:

$$\varepsilon(A) = \frac{\% \text{ change in } q}{\% \text{ change in } p} = \frac{\frac{\Delta q}{q_A}}{\frac{\Delta p}{p_A}} = \frac{\Delta q}{\Delta p} \cdot \frac{p_A}{q_A} = \frac{q_B - q_A}{p_B - p_A} \cdot \frac{p_A}{q_A}.$$

14.2 Beyond the price elasticity of demand

The sensitivity of *demand* to changes in *price* is not the only item of interest. For example, we might want to know the sensitivity of *consumption* to changes in *wealth*. In general, the **X elasticity of Y at point A** measures the percentage change in Y resulting from a 1% increase in X. (As always, these percentage changes are calculated relative to point A.) So if someone tells you that the wealth elasticity of consumption at present levels is 3, she means that a 1% increase in wealth from today's levels would increase consumption by 3%, or, equivalently, that a 1% decrease in wealth would decrease consumption by 3%.

As another example, consider the sensitivity of *supply* to changes in *price*, i.e., the **price elasticity of supply**. Like demand elasticities, supply elasticities are always measured at a point. Unlike demand elasticities, supply curves are positive: a 1% increase in price results in an increase in the quantity supplied. As shown in Figure 14.1, supply at point A is **elastic** if the price elasticity of supply at point A is greater than +1; **inelastic** if the price elasticity of supply at point A is less than +1; and **unit elastic** if the price elasticity of supply at point A is equal to +1.

Perfectly elastic and inelastic supply

As with demand elasticities, supply elasticities have the extremes of **perfectly elastic supply** (a price elasticity of supply of $+\infty$) and **perfectly inelastic supply** (a price elasticity of supply of 0). Unlike their analogues on the demand side, however, these extremes values on the supply side are of great interest.

Short-run supply curves can be perfectly inelastic at every point. As in Figure 14.2a, such a supply curve is completely insensitive to price changes: no matter what the price, suppliers want to sell the same amount. This makes sense for certain short-run supply curves, such as the supply of apartments in Seattle. Suppliers cannot instantly respond to higher rental prices by building more apartment buildings; it takes time to get permits, do the construction, etc. Similarly, suppliers are unable to instantly respond to lower rental prices by taking apartment buildings off the market; it takes time to convert apartments

¹If you do calculus, what we're getting at is $\frac{dq}{dp} \cdot \frac{p_A}{q_A}$.

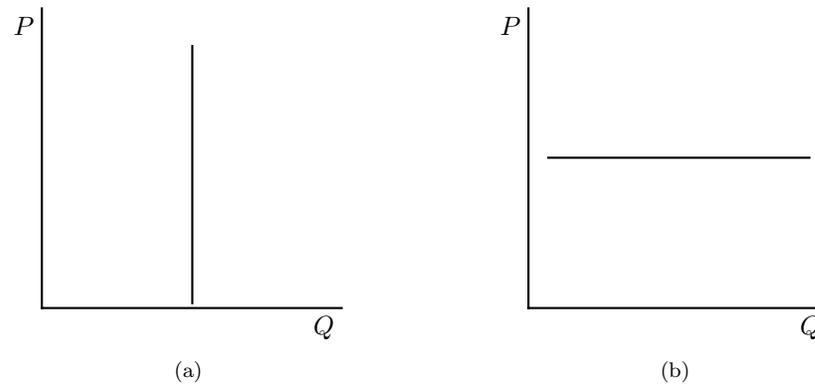


Figure 14.2: (a) A supply curve that is **perfectly inelastic** at every point; (b) a supply curve that is **perfectly elastic** at every point

into condominiums or to tear down apartment buildings and build something else. In the short run, then, the supply of apartments is fixed: if there are 50,000 apartments in Seattle today, there will be 50,000 apartments in Seattle next week, regardless of the rental price. In the short run, the supply curve for apartments is perfectly inelastic.

On the other extreme, *long-run supply curves can be perfectly elastic at every point*. As in Figure 14.2b, such a supply curve is infinitely sensitive to price changes: at any price higher than p , suppliers want to sell an infinite amount; at any price lower than p , suppliers want to sell zero. This makes sense for many long-run supply curves such as the supply of apple cider. In the long run, there is some price p at which the apple cider business generates the same rate of return as comparable investments. (Recall here the ideas from Chapter 5.) At a price of p , suppliers *in the long run* are indifferent between investing in the apple cider business and investing elsewhere. As a result, they are indifferent concerning how much apple cider they want to sell at price p : they would be willing to sell 100 gallons, 1 million gallons, or any other amount.

At any price other than p , however, the rate of return in the apple cider business will be different than the rate of return from comparable investments. Just as drivers shift out of slow-moving lanes and into fast-moving ones—the principle of arbitrage—investors will respond by entering or exiting the apple cider industry. At any price higher than p , the apple cider business generates higher rates of return than comparable investments, so *in the long run* sellers would want to sell an infinite amount of apple cider. At any price lower than p , the apple cider business generates lower rates of return than comparable investments, so *in the long run* nobody would want to sell apple cider.

The end result is a perfectly elastic long-run supply curve: suppliers want to sell zero units at any price less than p , they want to sell an infinite number of units at any price above p , and they are indifferent concerning the number of units they want to sell at a price exactly equal to p .

14.3 Applications

One application of elasticities involves how price changes affect total revenue. Consider a shift in the supply curve that leads to a small increase in the market price. Since the demand curve remains unchanged, the new equilibrium will feature a higher price and a lower quantity than the old equilibrium. So the impact on total revenue, pq , is unclear: p goes up but q goes down.

For simplicity, imagine that the increase in the market price amounts to 1%. *If the quantity remained unchanged*, then, total revenue would increase by 1%. But quantity does not remain unchanged. In fact, we have a name for the percentage change in quantity demanded resulting from a 1% increase in price: the price elasticity of demand!

This means we can use the price elasticity of demand to determine the impact of a supply-induced price change on total revenue. If demand is elastic at our original equilibrium, a 1% increase in p will lower q by more than 1%; the combined effect on pq will be to lower total revenue. If demand is inelastic at our original equilibrium, a 1% increase in p will lower q by less than 1%; the combined effect on pq will be to raise total revenue.

Elasticities also provide some insight into the behavior of a monopoly that engages in uniform pricing, i.e., that charges everyone the same price. Monopolists maximize profit, which is total revenue minus total costs, and although we need calculus to figure out exactly what price a monopolist will charge, we can show that *monopolists will never choose a price at which demand is inelastic*. To see why, imagine that the monopolist does choose a price at which demand is inelastic, and then consider increasing the price by 1%. This price increase will reduce the quantity demanded, so the monopolist doesn't have to produce as much; this lowers the total costs of production. And total revenue increases because we're on the inelastic portion of the demand curve. Both the higher total revenue and the lower total costs increase the firm's profits, so a profit-maximizing monopolist will always increase its price until it reaches the elastic portion of the demand curve.

A final application of elasticities is tax incidence: the ratio of the elasticities of supply and demand matches the ratio of the tax burdens. Since elasticities measure sensitivity to price changes, we can see that the burden of taxes falls most heavily on the party that is least able to avoid them, i.e., the party that is least sensitive to price changes. See problem 14.3e for an example.

Problems

Answers are in the next section. The first few words of each question (and each answer) are hyperlinked to allow online users to quickly move back and forth.

- 14.1 Go over previous problems (and/or make up some of your own) and calculate the elasticities of demand and supply at various points.
- 14.2 Do you think supply and demand curves are more elastic (i.e., more responsive to price changes) in the short run or in the long run? Why? How does your result translate graphically in terms of the steepness or flatness of the various curves?
- 14.3 Consider the hypothetical demand curve for oranges in Figure 14.3.

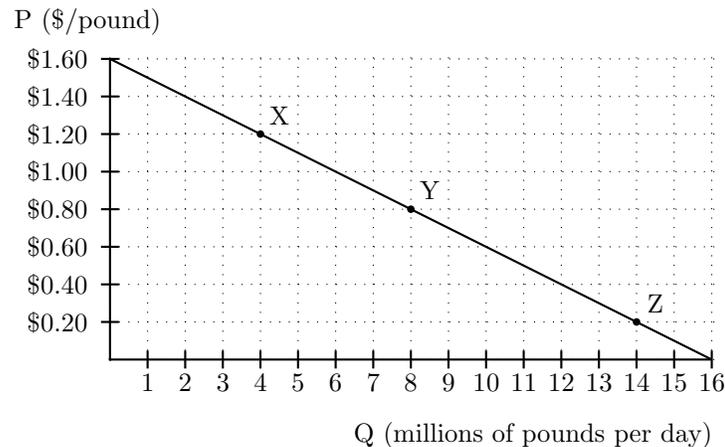


Figure 14.3: A hypothetical market for oranges

- Calculate the price elasticity of demand at point Y.
- During normal years, the supply curve is such that point Y is the equilibrium. Of the other two points, one is the equilibrium during “bad” years (when frost damages the orange crop), and one is the equilibrium during “good” years (when the orange crop thrives). Which one is point X?
- What is the total revenue at point X? At point Y? At point Z? (Use correct units!)
- The orange growers’ profit is total revenue minus total costs. If total costs are the same in all years, do the growers have higher profits in “bad” years or “good” years? Can you explain what’s going on here?
- This demand curve is the same as in problems 12.2–12.8 in Chapter 12. Go back to those problems, calculate the price elasticity of supply at the original equilibrium, and combine your answer there with your answer from problem 14.3a above to calculate the ratio of the elasticities $\left(\frac{\epsilon_S}{\epsilon_D}\right)$. Compare the results with the tax burden ratios and the slope ratios you calculated in those previous problems.

- 14.4 Show mathematically that the ratio of the elasticities of supply and demand is the inverse of the ratio of the slopes of the supply and demand curves, i.e., $\left(\frac{\epsilon_S}{\epsilon_D}\right) = \left(\frac{S_D}{S_S}\right)$.
- 14.5 Very long-run supply curves are often assumed to be perfectly elastic.
- Explain the intuition behind perfectly elastic long run supply curves. (Hint: Recall that comparable assets should have comparable rates of return.)
 - Draw a graph of a perfectly elastic supply curve.
 - Add a normal downward sloping demand curve to your graph. Then use the usual analysis to determine the incidence of a per-unit tax in this market. How much of the tax is paid by the buyers, and how much is paid by the sellers?

Answers

- 14.1 Check your answers with a classmate.
- 14.2 Supply and demand curves are more elastic in the long run than in the short run because in the long run buyers and sellers can respond in all the way they can respond in the short run, plus some others. In the long run a buyer of gasoline can move their home or their job to reduce their commute, can buy a new car, etc. Similarly, in the long run a seller of gasoline can open new oil fields or close existing ones. These choices indicate that long-run curves will be more elastic than short-run curves. Graphically this means that these curves will be *flatter*: a small change in price has a greater impact on the amount that buyers or sellers want to trade.
- 14.3 [“Consider the hypothetical. . .”]
- The point Y corresponds to point A in the elasticity formula, so we have $p_A = \$.80$ and $q_A = 8$. For point B we can take any other point, e.g., the convenient point with $p_B = \$.60$ and $q_B = 10$. Then

$$\epsilon = \frac{q_B - q_A}{p_B - p_A} \cdot \frac{p_A}{q_A} = \frac{2}{-.20} \cdot \frac{.80}{8} = -10 \cdot \frac{1}{10} = -1.$$
 - Point X is during bad years, when frost reduces supply.
 - Total revenue at the three points are pq , i.e., $(1.2)(4) = \$4.8$ million per day at point X, $(.8)(8) = \$6.4$ million per day at point Y, and $(.2)(14) = \$2.8$ million per day at point Z.

- (d) Growers make higher profits during “bad” years: their revenue is higher and their costs are assumed to be identical. This is basically a Prisoner’s Dilemma situation for the growers: they would all be better off if they could restrict supply during “good” years, but the individual incentives lead them to flood the market.
- (e) The point Y corresponds to point A in the elasticity formula, so we have $p_A = \$.80$ and $q_A = 8$. For point B we can take any other point on the supply curve, e.g., the convenient point with $p_B = \$.60$ and $q_B = 4$. Then

$$\varepsilon = \frac{q_B - q_A}{p_B - p_A} \cdot \frac{p_A}{q_A} = \frac{-4}{-.20} \cdot \frac{.80}{8} = 20 \cdot \frac{1}{10} = 2.$$

So the ratio of the elasticities is $\frac{\varepsilon_S}{\varepsilon_D} = \frac{2}{-1} = -2$. This is the same as the ratio of the slopes calculated previously! (This result follows from problem 14.4.)

14.4 We have

$$\frac{\varepsilon_S}{\varepsilon_D} = \frac{\frac{\Delta q_S}{\Delta p_S} \frac{p}{q}}{\frac{\Delta q_D}{\Delta p_D} \frac{p}{q}} = \frac{\frac{\Delta p_D}{\Delta q_D}}{\frac{\Delta p_S}{\Delta q_S}} = \frac{S_D}{S_S}.$$

14.5 [“Very long-run supply curves...”]

- (a) At some market price p , firms making (say) widgets earn the market rate of return; in the long run, then, firms are indifferent between making widgets and making other things, so they are willing to produce any number of widgets at price p . At any price less than p , firms would earn less than the market rate of return; in the long run, then, no firms would be willing to produce widgets, meaning that quantity supplied would be zero at any price less than p . Similarly, at any price greater than p , firms would earn more than the market rate of return; in the long run, then, everybody would rush into the widget-making business, meaning that the quantity supplied would be infinite at any price greater than p .
- (b) A perfectly elastic supply curve is a horizontal line at some price p .
- (c) A tax on the seller would shift the supply curve up by the amount of the tax. Since the supply curve is horizontal, the equilibrium price would increase by the full amount of the tax, meaning that buyers would pay the entire tax burden. (Similarly, a tax on the buyer would shift the demand curve down, but the equilibrium price would not change, meaning that the buyers bear the full burden of the tax.) This makes sense because of the analysis above: if sellers bear part of the tax burden then they would be earning less than the market rate of return. So in the long run buyers bear the entire burden of taxes in a competitive market.

Chapter 15

Transition: Welfare economics

How many economists does it take to screw in a lightbulb?

Two: they take turns explaining to each other how the invisible hand will take care of it without any effort on their parts.

The economic theory concerning the role of government focuses on the potential for governments to increase (or reduce) social welfare. This branch of economics, called **welfare economics**, is dominated by two theoretical results. First, as a mechanism for improving social welfare, free markets (i.e., letting people do whatever they want) work pretty well. The fact is summed up in one of the most important theorems in economics:

First Welfare Theorem: Complete and competitive markets yield Pareto efficient outcomes.

In other words, competition exhausts all possible gains from trade. If you're put in charge of the world and told to bring about an efficient allocation of resources, the advice from the First Welfare Theorem is to ensure that markets exist for all goods—i.e., that markets are **complete**—and that these markets are competitive. Once you've done that, all you need to do to get a Pareto efficient outcome is sit back and allow people to trade freely. This amazing result is sometimes called the Invisible Hand Theorem, a reference to the following quote from Adam Smith's *The Wealth of Nations*, originally published in 1776:

[People are] at all times in need of the co-operation and assistance of great multitudes. . . . The woolen coat, for example. . . is the produce of the joint labor of. . . [t]he shepherd, the sorter of the wool, the wool-comber or carder, the dyer, the scribbler, the spinner, the weaver, the fuller, the dresser, with many others. . . . [But man's]

whole life is scarce sufficient to gain the friendship of a few persons...and it is in vain for him to expect [help from] benevolence only....

It is not from the benevolence of the butcher, the brewer, or the baker that we expect our dinner, but from their regard to their own interest. We address ourselves, not to their humanity, but to their self-love, and never talk to them of our own necessity but of their advantages. [Man is] led by an invisible hand to promote an end which was no part of his intention. Nor is it always the worse for society that it was no part of it. By pursuing his own interest he frequently promotes that of society more effectually than when he really intends to promote it.

As long as markets are complete and competitive, then, the First Welfare Theorem says that competition will result in a Pareto efficient outcome. Recall from Chapter 6, however, that there are usually *many* Pareto efficient allocations of resources; the First Welfare Theorem tells us how to get *one* of them, but it doesn't indicate which one we will get, much less tell us how to get *a* particular one. This is an important issue, because Chapter 6 also shows that there's more to life than Pareto efficiency: a world where Bill Gates owned everything would be Pareto efficient, but it doesn't seem equitable, and it appears to leave something to be desired from the perspective of social welfare.

Fortunately, there is another theoretical result that provides a more well-rounded perspective on social welfare, most notably by addressing concerns about equity as well as efficiency:

Second Welfare Theorem: *Any* Pareto efficient outcome can be reached via complete and competitive markets, provided that trade is preceded by an appropriate reallocation of resources (also called a **lump sum transfer**).

To see how this works, pretend again that you're in charge of the world, but now your task is to achieve not just an *efficient* allocation of resources, but one that it *equitable* as well. (For the sake of argument, assume that "equitable" involves reducing the wealth disparities between, say, rich Americans and poor Americans, or between America and Africa.) To follow the prescription of the Second Welfare Theorem, then, you should begin by reallocating resources, which in this case means taking a certain amount of money and other assets away from rich Americans and giving it to poor Americans, or—in the second example—taking a certain amount of money and other assets away from Americans in general and giving it to Africans. Once you've done this, follow the instructions of the First Welfare Theorem: ensure that markets are complete and competitive, and then just sit back and let people trade freely. Provided that you reallocated resources in an appropriate way, you should end up with an outcome that is both Pareto efficient and equitable.

The Second Welfare Theorem has two key messages. The first is that *it is possible—at least in theory—to achieve both efficiency and equity*. There does

not have to be a trade-off between equity—however this is defined—and Pareto efficiency. The second key message is that *it is possible—at least in theory—to separate these two issues*. You can get efficiency by ensuring that there are complete and competitive markets; in doing so, you don't need to worry about equity because equity can be dealt with separately. And you can get equity by redistributing resources via lump sum transfers; in doing so, you don't need to worry about efficiency because efficiency can be dealt with separately.

15.1 From theory to reality

The First and Second Welfare Theorems indicate that the role of government—at least in theory—can be limited to only two tasks: one is ensuring that there are complete and competitive markets, and the other is reallocating resources in order to address equity concerns. Of course, the fact that such-and-such is possible in theory does not mean that it is possible in reality. What, then, are the real-world implications of these results? That is, what do they suggest about the role of government *in practice*?

Two particular implications are worth mentioning as a conclusion to this chapter and as food for thought:

Competitive markets deserve respect. Economists have a great deal of respect for competitive markets; they cannot take care of everything—for example, they don't address important equity concerns—but they can make an important contribution to social welfare by promoting Pareto efficiency. This suggests that there should be a presumption in favor of competitive markets: given a choice between a philosophy *supporting* government intervention in competitive markets unless conditions A, B, or C are met and a philosophy *opposing* government intervention in competitive markets unless conditions A, B, or C are met, a strong case can be made for the latter.

Lump sum transfers are key. The Second Welfare Theorem suggests that the ability to reallocate resources via lump sum transfers can make life easy. Unfortunately, the obverse is also true: the inability to reallocate resources via lump sum transfers can make life difficult. In particular, the inability to carry out lump sum transfers can make trade-offs between efficiency and equity unavoidable.

Part IV

Supplemental chapters

Chapter 16

Inflation

Present value calculations are complicated by the existence of **inflation**, a general increase in prices over time. In the last chapter we assumed that there was no inflation—i.e., that prices were constant over time—so it made sense to refer to *the* interest rate. But the presence of inflation requires us to distinguish between *nominal* interest rates and *real* interest rates.

16.1 Nominal and real interest rates

The **nominal interest rate** is what the bank pays. If you have \$100 and you put it in a bank paying 5% interest, in one year you'll have 5% more money. With inflation, however, having 5% more *money* next year doesn't mean you'll be able to buy 5% more *stuff* next year. Prices next year will be higher than prices today, so even though your bank account balance will be 5% larger, your ability to buy stuff—the **purchasing power** of your money—will not be 5% greater. In fact, if the inflation rate is higher than 5%, your purchasing power will be *lower* next year even though you have more money!

The **real interest rate** measures changes in purchasing power. For example, imagine that the nominal interest rate in Colombia is 13% and the inflation rate is 9%. If you put 1,000 Colombian pesos in the bank today at 13% interest, in one year you'll have 13% more, i.e., 1,130 pesos. But you won't be able to buy 13% more stuff, because prices will have risen by 9%: the same candy you can buy for 10 pesos today will cost 10.9 pesos next year. To figure out the actual increase in your purchasing power, we need to compare your purchasing power today with your purchasing power in one year. Today candy costs 10 pesos, so with your 1,000 pesos you could buy 100 candies; next year candy will cost 10.9 pesos, so with your 1,130 pesos you'd be able to buy 103.7 candies. Since putting your money in the bank for a year will enable you to buy 3.7% more candy, the real interest rate is 3.7%.

16.2 Inflation

To measure inflation economists look at how prices change for a representative “market basket” of goods and services: food, housing, transportation, education, haircuts, etc.¹ If that market basket had a price of \$10,000 in 1990 and \$12,900 in 2000, then we would say that the price level in the United States increased 29% during the decade. (This was in fact the price increase over that decade; it is equivalent to an annual inflation rate of 2.6%.²)

The most commonly used measure of inflation in the United States is the Consumer Price Index (CPI), which is shown in Figure 16.1a. According to the CPI, a market basket of goods and services that cost \$10 in 1920 would have cost about \$100 in 2005; in other words, the purchasing power of \$10 in 1920 was about the same as \$100 in 2005.

Taking inflation into account can produce significant shifts in perspective. As an example, consider Figures 16.1b and 16.1c, both of which show U.S. gas prices since 1920. Figure 16.1b features *nominal* prices: a gallon of gasoline actually sold for about \$0.30 in 1920 and \$4.10 in mid-2008. This perspective shows gas prices increasing over time, with significant jumps in the late 1970s and in recent years.

The perspective in Figure 16.1c—which adjusts for inflation by putting everything in **constant year 2005 dollars**, i.e., in terms of year 2005 purchasing power—is quite different. It shows gas prices mostly *falling* over time, interrupted by temporary price shocks in the 1930s and the 1970s and by a more recent price shock that began after oil prices reached an *all-time low* in 1998. Whether the recent price shock is also temporary or is a more permanent price increase caused by “peak oil” is a matter of debate.

For a more explicit comparison between Figures 16.1b and 16.1c, consider the price of a gallon of gasoline in 1920. Figure 16.1b asks this question: “What was the actual price of gasoline in 1920?” The answer is about \$0.30 per gallon. In contrast, the question posed in Figure 16.1c is this: “What would the average price of gasoline have been in 1920 *if the general price level in that year had been the same as the general price level in the year 2005?*” The answer is about \$3 per gallon because, according to the CPI, the general price level in the year 2005 was about 10 times higher than in 1920.

Because we have a good grasp of the purchasing power of today’s money, it often makes sense to use the current year, or a recent year, as the **base year** for analyzing real prices. It is, however, possible to use any year as the base year; Figure 16.1d shows gas prices in constant 1920 dollars, i.e., in terms of 1920 purchasing power. Note that the graph is identical to that in Figure 16.1c except for the labels on the vertical axis.

¹This would be as simple as it sounds if people bought the same stuff every year, but it isn’t because they don’t. Take an advanced microeconomics course to learn more.

²See problem 3.10 to understand why a 29% increase over a decade works out to only 2.6% per year over the course of that decade.

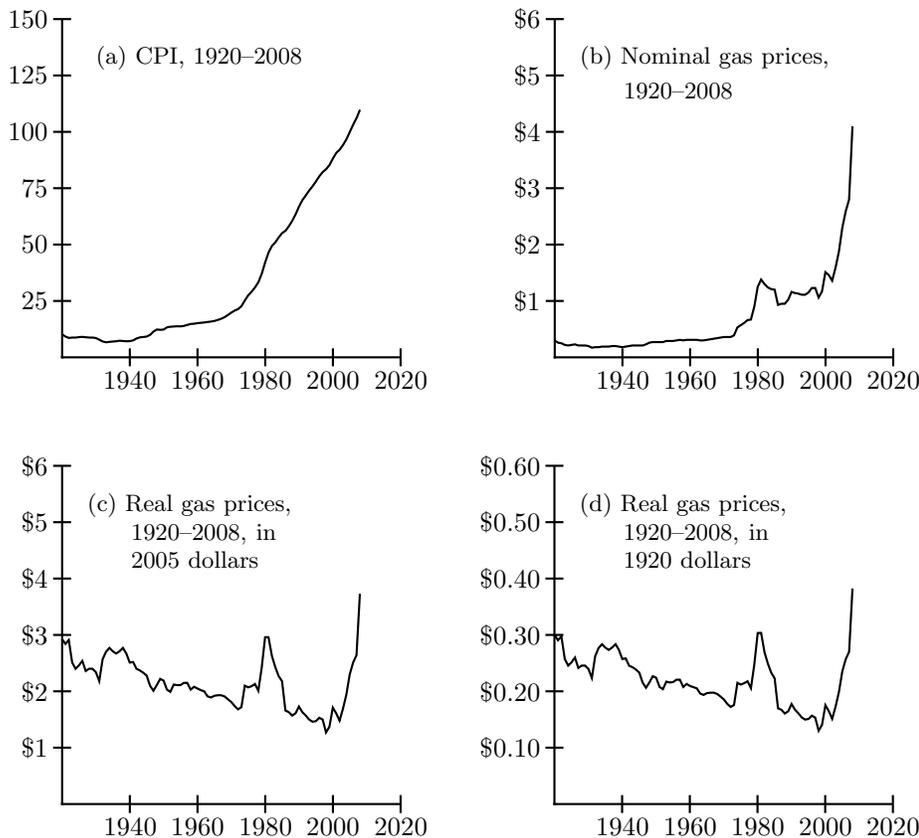


Figure 16.1: Inflation and U.S. gas prices, 1920–2008. Figure (a) shows the Consumer Price Index (CPI) measure of inflation; a representative “market basket” of goods and services that cost \$10 in 1920 would have cost about \$100 in 2005. (Note that there was a period of **deflation**—a general *decrease* in prices over time—during the Great Depression that began in 1929.) Figure (b) shows the average price for a gallon of gasoline using *nominal* prices: a gallon of gasoline actually sold for an average of \$0.30 in 1920 and \$4.10 in mid-2008. Figure (c) shows average gas prices using *real* year 2005 dollars, meaning that it adjusts for inflation by putting everything in terms of year 2005 purchasing power. Since a “market basket” that cost \$10 in 1920 would have cost about \$100 in 2005, the \$0.30 price tag on a gallon of gasoline in 1920 is equivalent to about \$3 in year 2005 dollars. Figure (d) shows the average gas price using *real* 1920 dollars. Figures (c) and (d) show that the \$4.10 price tag on a gallon of gasoline in mid-2008 is equivalent to about \$3.73 in 2005 dollars and \$0.37 in 1920 dollars. (Sources: American Petroleum Institute for pre-1949 gas prices, U.S. Energy Information Administration, U.S. Bureau of Labor Statistics.)

16.3 Mathematics

We can get a mathematical formula relating the nominal interest rate, the real interest rate, and the inflation rate by generalizing the approach from Section 16.1's discussion of Colombia. If the nominal interest rate is r_N and the inflation rate is i (e.g., $r_N = 0.13$ and $i = 0.09$), the real interest rate r_R is given by

$$1 + r_R = \frac{1 + r_N}{1 + i}, \quad \text{i.e.,} \quad r_R = \frac{1 + r_N}{1 + i} - 1.$$

Intuitively, the numerator $(1 + r_N)$ tells you how much more money you'll have in one year; the denominator $(1 + i)$ tells you how much prices will have risen in one year; and dividing one by the other tells you how much more purchasing power you'll have in one year. For the Colombia situation described above, the formula gives a real interest rate of 3.7%:

$$\frac{1 + 0.13}{1 + 0.09} - 1 = 0.037.$$

There is also a handy rule of thumb that works well when the inflation rate is small (say, less than 10%, so that $1 + i \approx 1$). In this case,

$$r_R = \frac{1 + r_N}{1 + i} - 1 = \frac{1 + r_N}{1 + i} - \frac{1 + i}{1 + i} = \frac{r_N - i}{1 + i} \approx r_N - i.$$

In short, $r_R \approx r_N - i$. In English, this says that the real interest rate is approximately equal to the nominal interest rate minus the inflation rate.

Remember that the rule of thumb is an approximation that works well only when the inflation rate is small. In the Colombia example from above, the rule of thumb estimates the real interest rate at $13\% - 9\% = 4\%$, which is pretty close to the actual rate of 3.7%. In contrast, consider a nominal interest rate of 113% and an inflation rate of 109% ($r_N = 1.13$, $i = 1.09$). The rule of thumb estimates the real interest rate at $113\% - 109\% = 4\%$, but this time the real interest rate is actually only 1.9%:

$$\frac{1 + 1.13}{1 + 1.09} - 1 = 0.019.$$

When to use which

It can be difficult to figure out when to use the nominal interest rate and when to use the real interest rate when computing present values. Two rules of thumb are described in Problem 16.4, but the only sure-fire strategy is to think about the goal, which is to figure out how much to put in the bank today to be able to afford a certain stream of income and expenses.

Problems

Answers are in the next section. The first few words of each question (and each answer) are hyperlinked to allow online users to quickly move back and forth.

- 16.1 If a bank is paying 14.4% and inflation is 8%, calculate the real interest rate. Round to the nearest .1% (Hint: Think about purchasing power relative to a good whose price increases at the rate of inflation.) Use both the true formula and the approximation and compare them.
- 16.2 Explain (as if to a non-economist) why the formula relating real and nominal interest rates makes sense. (Hint: Recall that the real interest rate measures increases in purchasing power, and think about how much more of some good you'll be able to purchase in one year if your bank account pays the nominal interest rate and the good's prices increases with inflation.)
- 16.3 Assume that the nominal interest rate is 10% per year and that the rate of inflation is 5% per year. Round all your answers as appropriate.
- You put \$100 in the bank today. How much will be in your account after 10 years?
 - You can buy an apple fritter (a type of donut) for \$1 today. The price of donuts goes up at the rate of inflation. How much will an apple fritter cost after 10 years?
 - Calculate x , the number of apple fritters you could buy for \$100 today. Then calculate y , the number of apple fritters you could buy after ten years if you put that \$100 in the bank. Finally, calculate $z = 100 \cdot \frac{y - x}{x}$. (The deal with z is that you can say, "If I put my money in the bank, then after ten years I will be able to buy $z\%$ more apple fritters.")
 - Given the nominal interest rate and inflation rate above, calculate the real interest rate to two significant digits (e.g., "3.81%"). Check your answer with the "rule of thumb" approximation.
 - Calculate how much money you'd have after 10 years if you put \$100 in the bank today *at the real interest rate you calculated in the previous question (16.3d)*. Compare your answer here with the result from question 16.3c.
- 16.4 Here are a couple of rules of thumb concerning the use of real (rather than nominal) interest rates in present value calculations.

Use real when your payments are inflation-adjusted. Somebody offers to sell you a lemon tree that will bear 100 lemons at the end of each year. The price of lemons is \$1.00/lemon right now, and will rise

at the rate of inflation, which is 4% per year; the nominal interest rate is 6%.

- (a) What is the present value of the lemon tree if it will bear fruit for 5 years and then die?
- (b) What if it will bear fruit forever?

Use real to calculate future purchasing power. You and a buddy win the \$20 million grand prize in a lottery, and you choose to accept a lump sum payment of \$10 million, which you divide equally. Your buddy immediately quits school and moves to Trinidad and Tobago. Being more cautious, you put your \$5 million in a 40-year CD paying 6%, figuring that after 40 years your wealth will have increased 10-fold and so you'll be able to buy 10 times more stuff. Are you figuring correctly if the inflation rate is 4%?

- 16.5 *Fun.* Recall the question from Chapter 3 concerning the Louisiana Purchase. That problem (#3.9 on page 34) asked you to calculate the present value of the \$15 million President Jefferson spent in 1803 to buy the Louisiana Territory from France, using interest rates of 2% and 8%. Assume now that 2% was the real interest rate over that time period and that 8% was the nominal interest rate over that time period. Which is the correct interest rate to use?
- 16.6 **In choosing between** a lottery annuity and a lump sum, should you use the real interest rate or the nominal interest rate?
- 16.7 **In 2009 the EPA** (Environmental Protection Agency) and the CBO (Congressional Budget Office) both did analyses of the Waxman-Markey cap-and-trade bill (HR 2454) that estimate permit prices for the next few decades. The EPA estimates that permit prices in 2020 will be about \$18 per metric tonne of CO₂ *in 2005 dollars*. The CBO estimates that permit prices in 2020 will be about \$28 per metric tonne of CO₂ *in 2020 dollars*. Assume that inflation in the 15 years between 2005 and 2020 is expected to be 2% per year.
- (a) Translate both the EPA and the CBO estimates into 2009 dollars so we can see which one is larger.
 - (b) Explain (as if to a non-economist) what it means that the CBO estimate for permit prices in 2020 is \$28 *in 2020 dollars*.
 - (c) Explain (as if to a non-economist) what it means that the CBO estimate for permit prices in 2020 is \$X *in 2009 dollars*.
 - (d) One gallon of gasoline produces about 20 pounds of CO₂, and there are about 2,205 pounds in a metric tonne. How much does a tax of (say) \$20 per metric tonne of CO₂ translate into in terms of dollars per gallon of gasoline?

16.8 Imagine that you are a profit-maximizing forester. You currently own trees containing 100 board-feet of timber.

- (a) With probability 2%, a fire will destroy your trees, and you'll have no harvestable timber. With probability 98%, your trees will grow and in one year you'll have 5% more board-feet of timber. What is the expected number of board-feet of timber you'll have next year?
- (b) Explain (as if to a non-economist) why the interest rate at the bank matters in deciding to cut the trees down now or to cut them down in year. (Hint: "trees are capital.")
- (c) Continuing with the story from part (a) above, assume that the price of lumber grows at the rate of inflation and that you're a risk-neutral forester. In order for cutting the trees down *now* to be a better choice than waiting until next year, the (pick one: nominal real) interest rate at the bank has to be (pick one: higher lower) than _____%.

16.9 You win a \$100 lump sum payment in the lottery! You decide to put your money in a 40-year Certificate of Deposit (CD) paying 6% annually. The inflation rate is 4% annually.

- (a) How much money will be in your bank account at the end of 40 years?
- (b) Assume that after 40 years you'll have 10 times more money (i.e., \$1000). Does this mean you'll be able to buy 10 times more stuff? *Briefly* explain.
- (c) Assume that "It's It" ice cream bars cost \$1 today, and that their price increases at the rate of inflation. How much will an It's It bar cost in 40 years? How many will you be able to buy with the money you'll have in 40 years? (Note: If you didn't get an answer to question 16.9a, use \$1000 for the amount of money you'll have in 40 years.)
- (d) Calculate the real interest rate using *both* the "rule of thumb" and the true formula.
- (e) Assume that the real interest rate is 1.92%. Use this interest rate to calculate the future value of your \$100 lump sum if you let it gain interest for 40 years. How does your answer compare with your answer from question 16.9c?

16.10 Social Security benefits are adjusted for inflation, meaning that payments to retirees increase at the rate of inflation.

- (a) How much will Grammy be paid in one year, in two years, and in three years if her current benefit is \$1000 and inflation is 3%?

- (b) One way to calculate the present value of these three payments is to use brute force: determine the present value of each payment separately and then add them together. Go ahead and do this when the nominal interest rate is 5%.
- (c) Calculate the present value of receiving \$1000 at the end of each year for 3 years if the relevant interest rate is 2%. Compare with your answer from 16.10b and explain.

Answers

16.1 The approximation is $14.4\% - 8\% = 6.4\%$. The actual answer comes from the formula, and gives a result of 5.9%.

16.2 This is explained in the text.

16.3 [“Assume that the nominal interest rate. . .”]

- (a) Use the nominal interest rate and the future value formula to get a bank account balance of about \$259.37.
- (b) Use the inflation rate and the future value formula to get an apple fritter price of about \$1.63.
- (c) Today you have \$100 and fritters cost \$1, so you can buy $x = 100$ of them. In ten years you’ll have \$259.37 and fritters will cost \$1.63, so you’ll be able to buy about $y = 159$ of them. So we can calculate $z \approx 59$.
- (d) The rule of thumb approximation says that the real interest rate should be about $10\% - 5\% = 5\%$. The actual value is $\frac{1+.1}{1+.05} - 1 \approx .0476$, i.e., 4.76%.
- (e) If you put \$100 in the bank at this interest rate, after 10 years you’d have about \$159. So you get $z = 59$ as your gain in purchasing power.

16.4 [“Here are a couple of rules. . .”]

- Use real when your payments are inflation-adjusted.**
- (a) Use the annuity formula and the real interest rate (about $6 - 4 = 2\%$) to get a present value of about \$470.
 - (b) Use the perpetuity formula and the real interest rate (about $6 - 4 = 2\%$) to get a present value of about \$5,000.

Use real to calculate future purchasing power. You’re not figuring correctly because you’re forgetting that prices are going to rise. Yes, you’ll have 10 times more money, but you won’t be able to buy 10 times more stuff. Using the real interest rate and the future value formula, we get a future value of \$11 million, or about 2.2 times more purchasing power.

16.5 Since we're trying to figure out what the current bank account balance is and the bank pays the nominal interest rate, we should use the nominal interest rate to determine if the Louisiana Purchase was really a great deal.

Note that an estimate of 2% for the real interest rate actually does make sense; in fact, it's called the Fischer Hypothesis, which you might have (or may yet) come across in macroeconomics. The 8% figure for the nominal interest rate, in contrast, is entirely fictitious; you'll have to study some economic history if you want a real approximation for the nominal interest rate over the last 200 years.

16.6 We're trying to figure out how much money we need to put in the bank today in order to finance cash payments in the future. Since the bank pays the nominal interest rate, that's the rate we should use.

16.7 ["In 2009 the EPA..."]

- (a) In 2009 dollars, the EPA estimate is $\$18(1.02)^4 \approx \19.48 and the CBO estimate is $\frac{\$28}{(1.02)^{11}} \approx \22.52 . So if we assume 2% inflation then the CBO estimate is slightly higher.
- (b) The CBO estimates that when 2020 rolls around, you'll be able to buy or sell a permit for \$28.
- (c) This means that if the general price level in 2020 were the same as it is in 2009, permits would cost \$X. In other words, we're adjusting for inflation by putting 2020 prices into terms that we understand today (in 2009).
- (d) We have

$$\frac{\$20}{1 \text{ tonne CO}_2} \cdot \frac{1 \text{ tonne CO}_2}{2,205 \text{ pounds CO}_2} \cdot \frac{20 \text{ pounds CO}_2}{1 \text{ gallon}} \approx \$0.18/\text{gallon}.$$

16.8 ["Imagine that you are..."]

- (a) $(0.1)(0) + (.98)(105) = 102.9$.
- (b) To maximize your present value you need to compare the return you'll get from "investing in the trees" with the return you'll get from investing in the bank. Investing in the bank means cutting down the trees and putting the proceeds in the bank. Investing in the trees means letting the trees grow so there will be more lumber next year.
- (c) The real interest rate has to be higher than 2.9%.

16.9 ["You win a \$100 lump sum..."]

- (a) Put \$100 and 6% in the future value formula to get about \$1028.57.

- (b) No: inflation means that you'll have 10 times more money, but not 10 times more purchasing power.
- (c) Plug \$1 and 4% into the future value formula to get a price of about \$4.80. With \$1028.57, you'll be able to buy about 214 ice cream bars.
- (d) The rule of thumb says that the real interest rate is approximately $6 - 4 = 2\%$. The true formula gives us $r = \frac{1+n}{1+i} - 1 = \frac{1.06}{1.04} - 1 \approx .019$, i.e., about 1.9%.
- (e) Plug \$100 and 1.92% into the future value formula to get a future value of about \$214. This equals the answer from question 16.9c.

16.10 ["Social Security benefits. . ."]

- (a) Plug the inflation rate (.03) into the future value formula to get a payment of \$1030 in one year, \$1060.90 in two years, and 1092.73 in three years.
- (b) Use the nominal interest rate and the lump sum formula to get a present value of approximately

$$\$980.95 + \$962.27 + \$943.94 = \$2887.16.$$

- (c) Plug \$1000, .02, and 3 years into the annuity formula to get a present value of \$2883.88. This is very close to the answer from above! (If you use 1.94% as a better estimate of the real interest rate, the annuity formula gives you a present value of \$2887.26, which is very close indeed to the \$2887.16 figure above.) The punch line here is that you can use the real interest rate to determine the present value of inflation-adjusted annuities.

Chapter 17

Fisheries

What is owned by many is taken least care of, for all men regard more what is their own than what others share with them.

—Aristotle, *A Treatise on Government*, Book 2¹

One place where *laissez faire* policies have definitely failed is in the water. Consider, for example, the history of the sea otter, an adorable marine mammal whose original range stretched from coastal areas in northern Japan to Alaska and down the west coast of North America to Baja California. Sea otters have more hairs on each square centimeter of their bodies than the average human being has on his or her entire head—giving them the unfortunate distinction of having the most luxurious fur of any animal.

Although sea otters had peacefully co-existed with native populations, they did not fare well after their discovery by Russian explorers in 1741. In that year there were perhaps 500,000 sea otters in the world. By the early 1900s there were less than 2,000. Motivated by prices of up to \$1,000 per pelt, hunters from Russia, Japan, and the United States drove the sea otter to the brink of extinction in less than 200 years.

Thanks to government intervention, the story of the sea otter did not end there. An international agreement ended the hunting of sea otters in 1911, and populations have rebounded. There are now an estimated 100,000 sea otters in Alaska, plus small populations elsewhere in North America (most famously in Monterey Bay, California).

The story of the sea otter is not unique. Another example is the Southern Atlantic wreckfish, which live at such great depths that they went undiscovered for most of the twentieth century. Then, in the mid-1980s, one was accidentally

¹This chapter was written in conjunction with graduate students in the UW School of Marine Affairs during a course on the microeconomics of marine affairs. Thanks to Melissa Andersen, Heather Brandon, Katie Chamberlin, Ruth Christiansen, Stacy Fawell, Julie Fields, Jason Gasper, Kevin Grant, Andy Herdon, Rus Higley, Heather Ludemann, Elizabeth Petras, Amy Seward, Pete Stauffer, Linda Sturgis, Brie van Cleve, and Jay Watson!

caught by a fisherman trying to recover lost equipment off the coast of Georgia. The fishery rapidly expanded, with catch totaling 30,000 pounds in 1987 and 4,000,000 pounds in 1990. Fearing the decimation of the fishery, regulators established a catch limit of 2,000,000 pounds for the next year; that limit was reached after only two months.

The same story is mirrored in the history of cod, rockfish, dogfish, and many other fish species. *Laissez faire* policies led to overfishing, the collapse of fish stocks, and government intervention.

17.1 An economic perspective

Instead of leading to Pareto efficiency, *laissez faire* fishery policies lead to disaster. Is this a failure of the “invisible hand”, a case in which free markets don’t work? An economist’s answer might be yes and no, respectively. Yes, the “invisible hand” fails: letting everyone do whatever they want leads to a Pareto inefficient outcome. But no, the problems in fisheries are not caused by the presence of markets. Indeed, economists argue that the problems in fisheries are caused by the *absence* of markets.

To see why, recall from Chapter 3 that **fish are capital**. In other words, fish can be seen as an investment: like money left in a savings account, fish left in the ocean will grow and multiply, yielding more fish to catch later. A profit-maximizing fisheries owner would compare the interest rate at the “Bank of Fish” with the interest rate at the Bank of America in order to determine how many fish to catch this year and how many to leave for later.

The problem with fisheries, from the perspective of economics, is that the theoretical “profit-maximizing fisheries owner” does not exist. Nobody owns the ocean, and nobody owns the fish in the ocean. In **open access** situations, anybody who wants to can go out and fish. An analogous investment situation would be a savings account for which everybody has an ATM card.

The owners of such a savings account would have little incentive to invest money in it, and the account would quickly become empty. Similarly, individuals in an open access fishery have little incentive to “invest in the fish” because someone else is likely to come along and catch their investment, leaving them empty-handed. The result is a **race for fish** and the subsequent collapse of the fishery.

Economists therefore attribute fisheries problems to a lack of **property rights**, i.e., of ownership. Government intervention is therefore necessary in order to prevent the Pareto inefficient outcome of overfishing.

17.2 A brief look at government intervention

Governments responded to fisheries problems in a variety of ways. Many established a goal of Maximum Sustainable Yield, i.e., management with the goal of getting the maximum possible catch that can be sustained year after year

forever. (See Figure 3.3 on page 29.) Although economists would quibble with the establishment of MSY as the desired objective (see Section 3.3 for details), it is interesting to see how governments attempted to achieve this goal.

One tactic was the creation and extension of Exclusive Economic Zones: by 1980 many countries had given themselves sole rights to all fishing grounds within 200 nautical miles of their shores.² While kicking out foreigners provided some temporary relief for fish stocks, it was not a permanent solution. To see why, consider the bank account analogy: investment in a savings account is unlikely to increase by changing from a situation in which everybody in the world has an ATM card to one in which only Americans (to take one example) have such a card.

Another tactic was the determination of Total Allowable Catch (TAC) limits: by limiting the number of fish that could be caught each year, regulators hoped to allow fish stocks to recover and eventually reach an equilibrium at the desired goal of MSY. Fisheries managers instituted a number of policies designed to keep total catch within the desired limits. Such policies included banning new vessels from entering a fishery; restricting fishing to certain types of equipment and/or certain types of vessels; restricting the days during which fishing was allowed; and, ultimately, closing the fishery after the TAC limit was met.

To the extent that such policies succeeded in limiting total catch, they succeeded in the long-term goal of promoting sustainable fisheries. Unfortunately, they failed to lead to Pareto efficient management of the fisheries in the short-term. *Each fishing season now featured its own race for fish*, with vessels desperately trying to catch as many fish as possible before the TAC limit triggered the close of the fishery for that season.

The results were a mixture of tragedy and farce. Despite the TAC limits, new vessels continued to crowd into the fisheries. Where new vessels were banned, existing vessels were rebuilt and equipped with high-tech electronics in order to expand capacity. As a consequence of the expansion of fish-catching capacity, the TAC limits were reached in ever-shorter periods of time; shorter fishing seasons and other additional restrictions in turn led to even more capacity expansion, creating a vicious cycle. In the New England surf clam fishery, the amount of time that vessels were allowed to fish fell from 96 hours per week in 1978 to under 3 hours per week in 1988. In British Columbia, the halibut season shrank from 65 days in 1980 to 6 days in 1990. At the start of the 1994 halibut season in Alaska, 2,400 fishing boats pulled up 24 million pounds of halibut within 24 hours; this was more than half of the TAC limit, and the fishery was closed shortly thereafter.

The race for fish—now evident each year rather than over time—was on with a vengeance. The result was Pareto inefficient for a number of reasons.

Danger The race for fish provided a huge incentive for fishing boats to head out

²For obvious reasons, the establishment of Exclusive Economic Zones greatly displeased foreigners who were used to fishing in these areas. Iceland's efforts led to three "Cod Wars" with Great Britain, the most serious of which (the Third Cod War, in 1975-76) featured a number of boat ramming and the deployment of part of the British Navy.

as soon as the season began, regardless of the weather. This threatened the safety of vessels and the crews aboard them.

Wasted fish Fishing in stormy weather also led to lost gear: in 1990 lost gear killed an estimated 2 million pounds of halibut in the Alaska fishery. A Pareto improvement would be to end the race for fish, or at least ground all vessels during bad weather. In fisheries that used nets, another problem was compaction: so many fish were being caught in each haul that the fish in the bottom of the nets would get squashed. Again, a Pareto improvement would be to end the race for fish; fishing boats could then afford to pull their nets up more often, reducing compaction.

Overcapitalization in the fishing industry There was a tremendous overinvestment in fishing capital (e.g., fishing boats, automatic hook baiters and other fancy fishing gear, etc.). After the TAC limit triggered the close of the fishing season, the boats would return to dock, often to sit dormant until the next season. A Pareto improvement would be to end the race for fish; given a longer fishing season, the same number of fish could be caught with less capital.

Overcapitalization in the processing industry With so many fish being caught in such short time frames, the race for fish led to overinvestment in fish processing. As with fishing boats, factories for processing and freezing fish would sit unused for much of the year, springing to life only during the short fishing season. Again, ending the race for fish would lead to a Pareto improvement: given a longer fishing season, the same number of fish could be processed with less capital.

Frozen fish Consumers prefer fresh fish, but short fishing seasons meant that they had to settle for frozen fish during much of the year.

Put it all together and the result was a **dissipation of resource rents**. In English, this means that human beings were not receiving the maximum benefit from this natural resource. Since consumers are willing to pay a lot of money for fish that nature produces “for free”, the potential exists for somebody to get a great deal: consumers should get highly valued fish for low prices, or fishermen should make large profits, or both. In other words, the potential exists for somebody to capture the resource rents associated with the fishery. But the race for fish squandered much of this potential bounty. Even though consumers placed much higher values on fresh fish, they had to settle for frozen fish for much of the year. Fishermen weren’t getting ahead either: they had to spend so much money on capital (bigger fishing boats with fancy technology) that they didn’t make much profit. In short, everybody lost.

It is important to note that the race for fish is not inevitable. The Coase Theorem says that there are always incentives for individuals to bargain their way out of Pareto inefficient situations, and in some fisheries this happened. Most notably, factory trawlers in the Pacific whiting fishery avoided the race for fish by forming a co-operative that imposed limits on each vessel and put

observers on each boat to make sure the limits were obeyed. Although participation in the co-operative was voluntary, each owner knew that breaking the rules would result in an end of the co-operative and a return of the (unprofitable) race for fish the following season. In a repeated game, this **trigger strategy** gives each firm an incentive to participate in the co-operative.

In short, the factory trawlers avoided the race for fish by colluding to fix harvest amounts. Its success, however, depended on the unique situation in the whiting fishery: with only three firms in the market and a government-mandated moratorium on entry by other firms, the conditions for successful collusion were perfect. As in normal markets, however, collusion between fishermen becomes more difficult as the number of fishermen increases, and is extremely difficult when there is free entry into the fishery. These impediments to bargaining help explain why fishermen in most fisheries were collectively unable to stop the race for fish.

17.3 ITQs to the rescue?

To solve the short-term race for fish, economists advocate the use of Individual Transferable Quotas (ITQs) in conjunction with limits on total allowable catch (TAC). Since the economics explanation for the race for fish is a lack of property rights, the economics solution is for the government to create property rights: give specific individuals the rights to a certain number of fish (a **quota**) during each season. ITQ systems even allow individuals to trade “their” fish, i.e., to buy or sell these fishing rights. (Some fisheries use quotas but do not allow these trades; these policies are sometimes known as Individual Fishing Quotas rather than Individual Transferable Quotas.)

In theory, ITQs hold a great deal of promise. First, they end the race for fish: quota owners don’t have to worry about other fishermen stealing “their” fish. With their property rights secure, those individuals can then fish whenever they want to during the season. Second, the transferability of quotas provides an opportunity for additional Pareto improvements. Fishing vessels with low costs can purchase or lease quotas from vessels with high fishing costs; both parties benefit from this trade, and this Pareto improvement pushes the industry toward an efficient system of least-cost fishing. There is also some possibility (discussed below) that ITQs can promote the long-term sustainability of fisheries.

The economic theory of ITQs has been put to the test in Iceland and New Zealand, which rely heavily on ITQs. (They are also used in four fisheries in the United States: surf clams/ocean quahog in New England, wreckfish in the South Atlantic, and halibut and sablefish in Alaska.) These experience suggest three main conclusions about ITQs.

First, **ITQs successfully eliminate the short-term race for fish.** Less danger (in some cases), less waste, less overcapitalization in fishing and processing, more fresh fish for consumers, less dissipation of rents.

Second, **adopting ITQs may not make everyone better off.** Although the race for fish is Pareto inefficient, it is not easy to make ITQs into a Pareto

improvement. Most obviously, fishermen who don't get ITQs might be hurt: they now have to pay for permits in order to do something that they were previously able to do for free. But there are other concerns as well. Rationalization of the fishing industry might result in lower employment levels and other negative impacts on fishing communities. Ending the race for fish might also hurt people in the boat-building industry: they benefit from overcapitalization in the fishing industry, and lose when overcapitalization ends. As a final example, fish processors might suffer short-term losses from the transition to ITQs. Dealing with the sudden glut of fish that occurs with a race for fish requires lots of processing capacity; ITQs spread the fishing season out over more time, so less processing capacity is required. Processors who have built factories in anticipation of a continued race for fish will therefore be hurt by the introduction of ITQs.

Finally, **ITQs are not perfect**. One concern is **highgrading**: fishermen may attempt to fill their quota with the most highly valued specimens. This may lead to wasteful practices such as discarding fish of lower value. A second concern is **concentration** of ITQs in the hands of only a few individuals. Some fisheries now have rules preventing excessive consolidation. And there is the perennial concern about **equity**: who should get the ITQs? Most fisheries allocate quotas on the basis of historical catch, but this approach—like all others—doesn't please everybody. This issue is difficult for economists to analyze because economics is largely silent when it comes to equity. Like the allocation of cake in the cake-cutting games discussed in Chapter 6, the initial allocation of ITQs is unimportant as far as economics is concerned. What is important is that the ITQs, once allocated, can be freely traded so as to produce a Pareto efficient outcome.

ITQs and sustainability

The role of ITQs in promoting sustainable fisheries is unclear. At first glance, ITQs are irrelevant to the long run sustainability of a fishery. After all, ITQs are designed to solve the race for fish in the short run, i.e., within a particular fishing season. Solving the race for fish in the long run requires an appropriate determination of the Total Allowable Catch (TAC) limit. If fisheries managers set the TAC limit too high, ITQs will do nothing to stop the collapse of a fishery.

Upon further reflection, however, ITQs may have a role after all. Fishery quotas are often calculated as a percentage of the (TAC) limit. Since TAC levels will go down in the future if overharvesting occurs in the present, quota owners have some incentive to preserve stocks at sustainable levels. These incentives might lead quota owners to pressure fisheries managers to maintain a truly sustainable TAC limit. In contrast, the pressure in non-ITQ fisheries is often to increase TAC for short-term gain at the expense of long-term sustainability.

This argument has its limits. Although all the ITQs together may benefit from sustainable management, each ITQ owner individually may not have much of an incentive to push for sustainability. The **Prisoner's Dilemma** rears its ugly head yet again...

Chapter 18

Sequential-move games

A **sequential-move game**—such as chess or poker—is a game in which the players take turns moving. We can analyze these games with **game trees**, which are multi-person extensions of the decision trees from Chapter 1. Figure 18.1 shows an example.

The label at each node of a game tree indicates which player moves at that node, e.g., in this game player 1 moves first. The various branches coming out of that node identify the different options for that player, e.g., in this game player 1 begins by choosing Up, Middle, or Down. When the player chooses, we move along the chosen branch; sometimes this leads to another node, where it's another player's turn to move, e.g., player 2's turn if player 1 chooses Up or Down; and sometimes—e.g., if player 1 chooses Middle—this leads to the end of the game, where we find the payoffs for all of the players. In this game, if player 1 chooses Up and then player 2 chooses Up then the outcome is $(-3, 2)$, meaning that player 1's payoff is -3 and player 2's payoff is 2 . (Payoffs are listed in order of the players, with player 1's payoff first.)

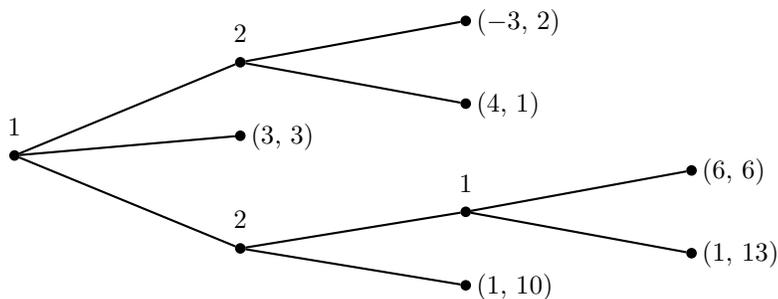


Figure 18.1: An example of a game tree

18.1 Backward induction

In sequential-move games players need to *anticipate* what's going to happen later on (“If I do X, she’ll do...”) and so the way to approach these games is to work backwards from the end, a process called **backward induction**:

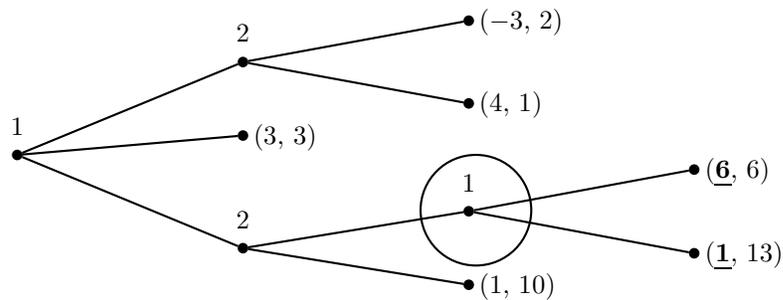


Figure 18.2: We begin at one of the nodes at the end of the tree—like the one circled here—and ask what player 1 would do *if* the game got to this node. Since 6 is a better payoff than 1, we can anticipate that player 1 will choose Up *if* the game gets to this circled node, draw an arrow (as in the next figure) to indicate this choice, and continue to move *backwards* through the tree.

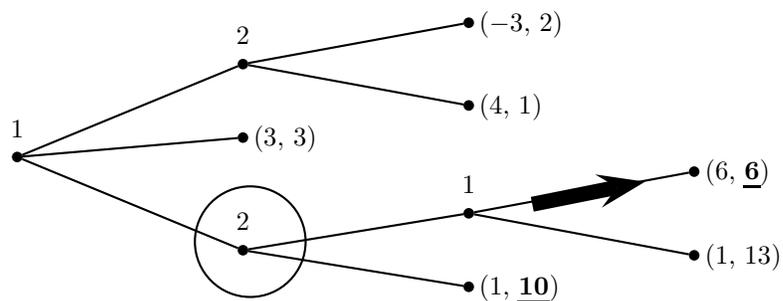


Figure 18.3: Analyzing this circled node is easy because the arrow shows how player 1 would respond *if* player 2 chooses Up: player 1 would choose Up, and both players would receive a payoff of 6. Since player 2 can get a payoff of 10 by choosing Down at the circled node, we can anticipate that player 2 will choose Down *if* the game gets to the circled node, draw an arrow (as in the next figure) to indicate this choice, and continue to move *backwards* from the end of the tree.

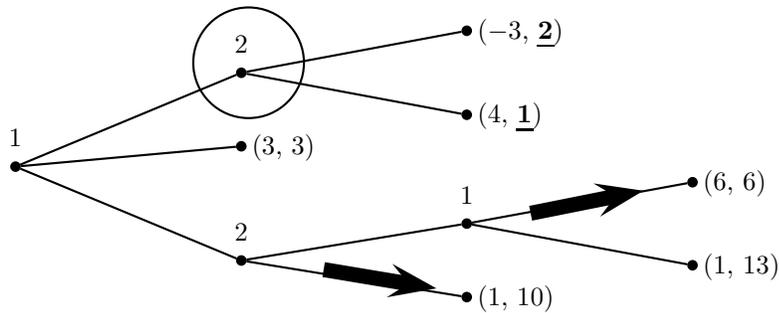


Figure 18.4: At this point it is tempting to move backwards to the opening node of the game and ask what player 1 will choose, but we cannot do that until we have worked backwards through the rest of the tree. So we must now analyze the circled node by asking what player 2 would do *if* the game got to this node. Since 2 is a better payoff than 1, we can anticipate that player 2 will choose Up *if* the game gets to this circled node, draw an arrow (as in the next figure) to indicate this choice, and continue to move backwards to the one node remaining on the tree: the opening node.

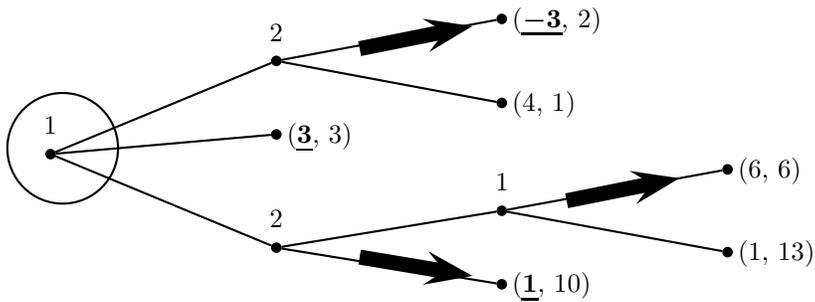


Figure 18.5: Analyzing this circled node—the opening move of the game—is now easy because the arrows show how the game will play out *if* player 1 chooses Up, Middle, or Down. If player 1 chooses Up, we see that player 2 would choose Up, giving player 1 a payoff of -3 ; if player 1 chooses Middle, we see that the game would end immediately, giving player 1 a payoff of 3 ; and if player 1 chooses Down, we see that player 2 would choose Down, giving player 1 a payoff of 1 . Since 3 is bigger than -3 or 1 , we can anticipate that player 1 will begin the game by choosing Middle and draw an arrow (as in the next figure) to show the complete analysis of the game.

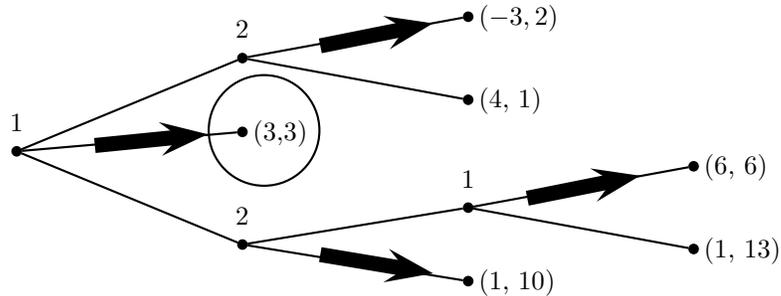


Figure 18.6: After using backward induction to figure out how the players would choose *if* the game reached different nodes, we can now simply follow the arrows to see how the game will play out.

Problems

Answers are in the next section. The first few words of each question (and each answer) are hyperlinked to allow online users to quickly move back and forth.

18.1 *Challenge.* Explain (as if to a non-economist) why backward induction makes sense.

18.2 Analyze the sequential-move game in Figure 18.7 using backward induction.

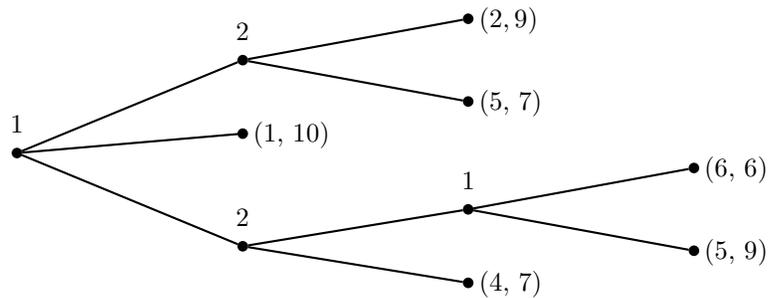


Figure 18.7: A random game tree

(a) Identify (e.g., by circling) the likely outcome of this game.

- (b) Is this outcome Pareto efficient? If it is not Pareto efficient, identify a Pareto improvement.

18.3 Analyze the sequential-move game in Figure 18.8 using backward induction.

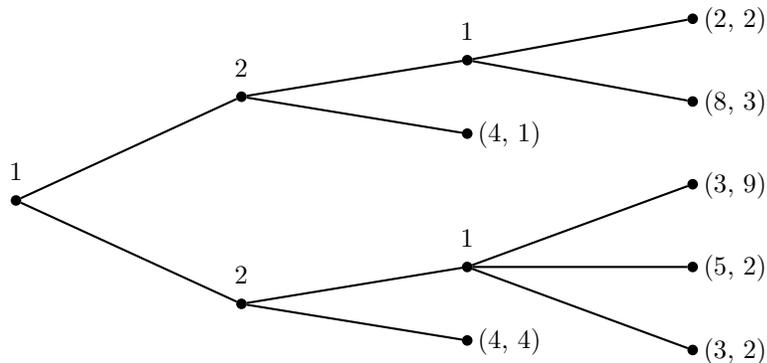


Figure 18.8: Another random game tree

- (a) Identify (e.g., by circling) the likely outcome of this game.
- (b) Is this outcome Pareto efficient? If it is not Pareto efficient, identify a Pareto improvement.
- 18.4 Consider tipping, a social phenomenon observed in some (but not all!) countries in which restaurant patrons leave some extra money behind for their waiter or waitress. Would tipping provide much of an incentive for good service if the tips were handed over at the beginning of the meal rather than at the end? Are there any difficulties in the incentive structure when tips are left at the end of meals? Write down game trees to support your arguments.
- 18.5 (Overinvestment as a barrier to entry) Consider the following sequential-move games of complete information. The games are between an incumbent monopolist (M) and a potential entrant (PE). You can answer these questions without looking at the stories, but the stories do provide some context and motivation.

Story #1 (See Figure 18.9): Firm M is an incumbent monopolist. Firm PE is considering spending \$30 to build a factory and enter the market. If firm PE stays out, firm M gets the whole market. If firm PE enters the market, firm M can either build another factory and engage in a price war or peacefully share the market with firm PE.

- (a) Identify (e.g., by circling) the likely outcome of this game.
- (b) Is this outcome Pareto efficient? Yes No (Circle one. If it is not Pareto efficient, identify, e.g., with a star, a Pareto improvement.)

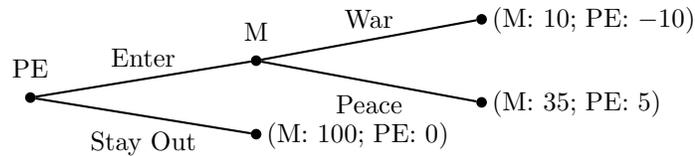


Figure 18.9: Story #1

Story #2 (See Figure 18.10): The monopolist (firm M) chooses whether or not to overinvest by building a second factory for \$30 even though one factory is more than enough. Firm PE (the potential entrant) sees what firm M has done and decides whether to enter or stay out, and if PE enters then M decides whether or not to engage in a price war.

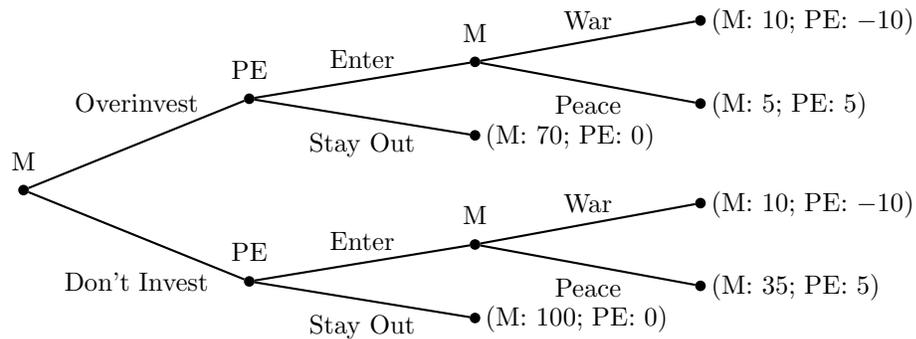


Figure 18.10: Story #2

- (c) Identify (e.g., by circling) the likely outcome of this game.
 (d) Is this outcome Pareto efficient? Yes No (Circle one. If it is not Pareto efficient, identify, e.g., with a star, a Pareto improvement.)

18.6 (The Sticks Game) The sticks game works as follows: We put n sticks on the table. Beginning with Player 1, the two players take turns removing either one or two sticks. The player who removes the last stick must pay the other player \$1.

- (a) If there are 10 sticks on the table, which player would you rather be, and what strategy will you employ?

- (b) *Challenge* If there are n sticks on the table, which player would you rather be? Can you describe a general strategy?

18.7 *Challenge (The Ice Cream Pie Game, from Dixit and Nalebuff)* Two players take turns making take-it-or-leave-it offers about the division of an ice cream pie. In the first round, the whole pie is available; if Player 2 accepts Player 1's proposal then the two players share the entire pie; if Player 2 rejects Player 1's proposal, half of the pie melts away, and we go to round two (in which Player 2 makes a take-it-or-leave-it offer about the division of the remaining pie). The game ends when an offer is accepted, or after the end of the n th period (at which point Mom eats the remaining pie, meaning that the players get nothing). Assume that each player just wants to get as much pie as possible, even if it's just a sliver.

- (a) Predict the outcome of the game when there are 1, 2, and 3 periods.
 (b) Now assume that 1/3rd of the pie (rather than 1/2) melts away each period. Predict the outcome when there are 1, 2, and 3 periods.
 (c) Hopefully your prediction is that the first offer made is always accepted. Try to understand and explain (as if to a non-economist) why this happens.

18.8 *Make up some game trees* (players, options, payoffs, etc.) and solve them using backward induction.

18.9 *(The Draft Game, from Brams and Taylor)* Three football teams (X, Y, Z) are involved in a draft for new players. There are six players to choose from (Center, Guard, Tailback, Quarterback, Halfback, Fullback), and the draft works as follows: X chooses a player, then Y chooses one of the remaining five players, then Z chooses one of the remaining four players (this constitutes the first round of the draft); the same procedure is repeated in the second round, at the end of which all six players are taken.

The teams' preferences are as follows:

	Top choice	Second	Third	Fourth	Fifth	Sixth
X	C	G	T	Q	H	F
Y	H	F	G	C	Q	T
Z	T	F	H	Q	C	G

Assume that the teams all know each others' preferences. Then we can model the draft as a game tree, with team X choosing first &etc. The complete game tree for this draft is quite involved, but *trust me, it all boils down to the game tree shown in Figure 18.11.*

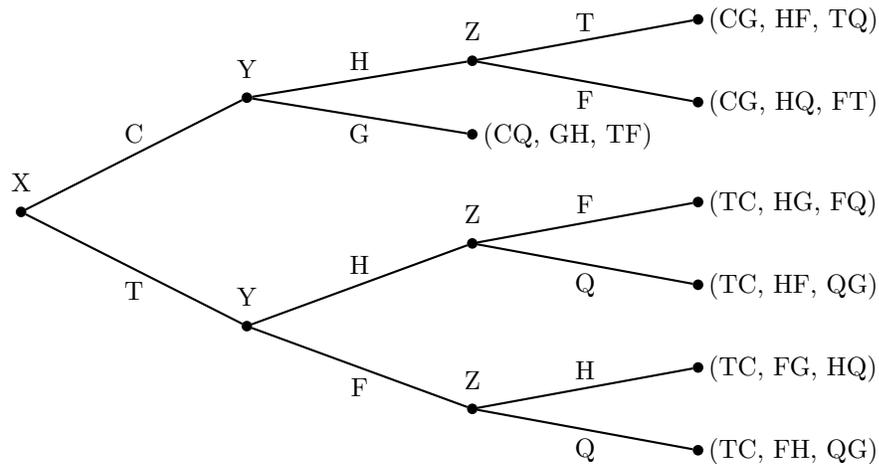


Figure 18.11: The draft game

The payoffs for this game are the players each team gets. For example, (CG, HQ, TF) indicates that team X gets the Center and the Guard (its #1 and #2 choices), team Y gets the Halfback and the Quarterback (#1 and #5), and team Z gets the Tailback and the Fullback (#1 and #2). Clearly each team would prefer to get the players it likes the most, e.g., team X prefers CT (or TC) to CQ or GQ.

- The “naive” strategy is for each team to choose its top choice among the available players every time it gets to pick. What is the outcome of this strategy?
- If teams X and Y pursue this naive strategy by picking C and H in the first round, should team Z also pursue this strategy, i.e., pick T? Briefly explain why or why not.
- What outcome do you expect from this game using backward induction?
- Is the expected outcome you identified Pareto efficient? If so, explain. If not, identify a Pareto improvement.
- Statement 1: “In the first round, the optimal move for each team is to pick the best available player.” Statement 2: “In the second round, the optimal move for each team is to pick the best available player.” Explain why Statement 1 is false but Statement 2 is true.
- Challenge* Prove that the game tree really does boil down to what’s shown on the previous page.

18.10 *Fun.* (The Hold-Up Problem) In the movie *Butch Cassidy and the Sundance Kid* (1969), Paul Newman and Robert Redford play Wild West bank robbers who are particularly fond of robbing the Union Pacific Railroad. The CEO of the railroad, Mr. E. H. Harriman, hires a “superposse” of gunslingers to bring the two robbers to justice, dead or alive. After a long (and rather boring) chase scene, Butch and Sundance manage to escape. Afterwards, Butch reads about the superposse in the newspaper and has this to say:

A set-up like that costs more than we ever took...That crazy Harriman. That’s bad business. How long do you think I’d stay in operation if every time I pulled a job, it cost me money? If he’d just pay me what he’s spending to make me stop robbin’ him, I’d stop robbin’ him. [Screaming out the door at E. H. Harriman:] Probably inherited every penny you got! Those inherited guys—what the hell do they know?

- (a) Is what Harriman is doing bad business? Explain why or why not. Your answer may depend on the assumptions you make, so explicitly state any and all assumptions. You might also want to draw a game tree, make up appropriate payoffs, and solve the game using backwards induction.
 - (b) “If he’d just pay me what he’s spending to make me stop robbin’ him, I’d stop robbin’ him.” Assume this statement is true. What does it say about the efficiency or inefficiency of the situation?
 - (c) What do you think about the argument contained in the previous quote? Can you see why this is called the “hold-up problem”?
 - (d) The hold-up problem also applies to students working jointly on projects and to firms engaged in joint ventures: after one member makes an irreversible investment, the other member may try to renegotiate the terms of the deal. Explain how contracts might help in preventing this difficulty, and why contracts wouldn’t work in the case of Butch Cassidy.
- 18.11 *Fun.* The IgoUgo travel guide *Andiamo* provides the following advice for crossing the street in Rome: “First, stand towards a coffee bar and watch a local or two. See how they boldly walk out into traffic? Now it’s your turn! Choose your moment but don’t hesitate for too long. Waiting for traffic to clear will not happen. When you appear to have the most lead-time, step boldly off the curb and walk swiftly and confidently for the opposite side of the street. Do not look at the traffic careening towards you—believe it or not, they will stop for you! But do not look at them—do not make eye contact—this is an invitation for sport. Just walk briskly with your head up and your eyes on the prize- the opposite sidewalk.”

- (a) “[D]o not make eye contact—this is an invitation for sport.” Explain.

- (b) Set up a game tree for this “pedestrian in Rome” game and solve it.
- 18.12 (The Centipede Game) There are 5 \$1 bills on a table. Players 1 and 2 take turns moving. Each turn the player moving takes either \$1 (in which case it becomes the other player’s turn) or \$2 (in which case the game ends). Each player wants to get as much money as possible.
- (a) Draw a game tree for this game.
 - (b) Predict the outcome of this game.
 - (c) Is this outcome Pareto efficient? If so, explain briefly. If not, identify a Pareto improvement.
 - (d) *Challenge.* Can you generalize this result when there are n bills on the table? (Hint: Try induction.)
 - (e) *Super Challenge.* Can you reconcile the previous answer with your intuition about how this game might actually get played in real life?
- 18.13 *Fun.* (The surprise exam paradox) Your class meets 5 days a week, and on Friday your teacher tells you that there will be a surprise exam next week, meaning (1) that there will be an exam, and (2) that it will be a surprise (i.e., you won’t be able to anticipate the night before the exam that the exam will be the next day). What can you conclude about the exam? Relate this problem to the Centipede Game discussed previously.
- 18.14 The game in Figure 18.12 is just one example of how game theory might be applied to international climate negotiations. (Repeat, it’s just one example.) The game involves player 1 (the U.S.) choosing whether to unilaterally cut emissions (“cut”) or whether to delay (“delay”) in the hopes of reaching agreement on an international climate treaty. In either case, player 2 (China) then chooses whether to agree to an international climate treaty (“agree”) or refuse to be part of an international treaty (“not agree”). The outcomes could be considered to be measures of societal well-being in the two countries (U.S., China). Analyze the following sequential move game using backward induction.
- (a) Identify (e.g., by circling) the likely outcome of this game.
 - (b) Is this outcome Pareto efficient? If it is not Pareto efficient, identify, e.g., with a star, a Pareto improvement.)
- 18.15 Consider the following version of a game called the Ultimatum Game: Player 1 begins by proposing a take-it-or-leave-it division of ten \$1 bills between himself and Player 2. (For the sake of simplicity, assume that he has only three options: he can keep \$9 himself and offer \$1 to Player 2, or he can keep \$5 himself and offer \$5 to Player 2, or he can keep \$1 himself and offer \$9 to Player 2.) Player 2 then either accepts or rejects the offer. If she accepts the offer, the players divide up the money and the game ends; if she rejects the offer, both players get nothing.

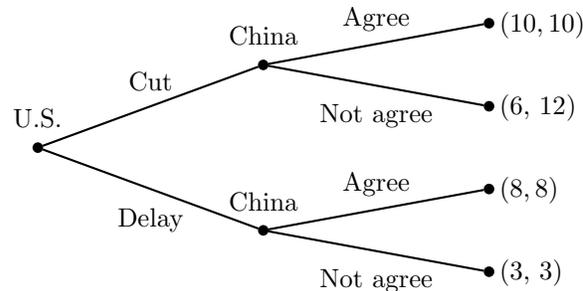


Figure 18.12: A climate negotiations game.

- (a) Draw a game tree for this game.
- (b) *Assuming that each player's sole motivation is to get as much money as possible*, backward induction predicts that the outcome of this game will be for Player 1 to choose the first option (keeping \$9 for himself and offering \$1 to Player 2) and for Player 2 to accept his offer. Explain—as if to a non-economist—the underlying logic here, either in words or using the game tree. (Note: if you think backward induction predicts a different solution, well, explain that one.)
- (c) Daniel Kahneman and Vernon Smith won the 2002 Nobel prize in economics for exploring how real people actually make decisions in games like these. (Amos Tversky would have won the prize, too, but he died in 1996.) In classroom experiments, Player 1 sometimes offered Player 2 more than \$1, and—in situations where Player 1 did offer Player 2 only \$1—Player 2 sometimes rejected the offer. For each of the following statements, indicate whether it is true or false *on the basis of these experimental results* and provide a brief sentence of explanation.
- i. These experiments showed that some of the Player 2's were not motivated solely by getting as much money as possible.
 - ii. These experiments showed that some of the Player 1's were not motivated solely by getting as much money as possible.
 - iii. These experiments showed that the basic assumption of economics (that decisions are made by optimizing individuals) is wrong.

18.16 Consider the following 2-period cake-cutting game between Jack and Jill, each of whom has as his or her sole objective the desire for as much cake as possible. In round 1 there are three ounces of cake, and Jack makes a take-it-or-leave-it offer to Jill. If Jill accepts, the game ends and the

players divide and eat the cake; if Jill rejects, Mom eats one ounce and the game moves to round 2. In round 2 there are two ounces of cake, and Jill makes a take-it-or-leave-it offer to Jack. If Jack accepts, the game ends and the players divide and eat the cake; if Jack rejects, the game ends and both players get nothing.

- (a) Backward induction predicts that Jack will offer two ounces of cake to Jill in round 1, leaving one ounce for himself, and that Jill will accept. Explain the reasoning behind this prediction.
- (b) This cake-cutting game has something in common with such real-world phenomena as labor disputes or lawsuits in that delay hurts both sides: the longer the strike or lawsuit drags on, the worse off the various players are. As in the game above, settlement in round 1 is the only way to reach an outcome that is Pareto (circle one: efficient inefficient). What does the Coase theorem have to say about *when* such conflicts are likely to be resolved? Give an example of what the Coase theorem predicts using labor disputes or lawsuits. (Recall the essence of the Coase theorem: “If there is nothing to stop people from trading, nothing will stop people from trading.”)

Answers

- 18.1 This is explained (to the best of my abilities) in the text. The basic idea is that you need to anticipate your rival’s response.
- 18.2 [“Analyze the sequential-move game. . .”]
- (a) Backward induction predicts an outcome of (4, 7).
 - (b) No, it is not Pareto efficient; both (5, 7) and (5, 9) are Pareto improvements.
- 18.3 [“Analyze the sequential-move game. . .”]
- (a) Backward induction predicts an outcome of (8, 3).
 - (b) Yes, it is Pareto efficient.
- 18.4 Tipping at the beginning of the meal is problematic because then the waitress has no incentive to provide good service. (The tip is already sunk.) Tipping at the end of the meal is problematic because then the customer has no incentive to provide the tip. (The service is already sunk.)
- 18.5 [“Overinvestment as a barrier to entry. . .”]
- (a) Backward induction predicts an outcome of (M: 35, PE: 5).
 - (b) Yes.

- (c) Backward induction predicts an outcome of (M: 70, PE: 0).
- (d) No; a Pareto improvement is (M: 100, PE: 0).

18.6 [“The Sticks Game...”]

- (a) If there are 10 sticks on the table, you should be player 2. Whenever your opponent takes 1 stick, you take 2; when he takes 2 sticks, you take 1. So you can force your opponent to move with 7 sticks, then 4 sticks, then 1 stick—so you win!
- (b) Hint: The above answer suggests a general strategy to follow.

18.7 [“The Ice Cream Pie Game...”]

- (a) With one period, Player 1 offers Player 2 a sliver, and Player 2 accepts. With two periods, Player 1 offers Player 2 half the cake, and Player 2 accepts. (Both know that if Player 2 refuses, half the cake melts, Player 2 will offer Player 1 a sliver of the remaining half, and Player 1 will accept.) With three periods, Player 1 offers Player 2 one-quarter of the cake, and Player 2 accepts. (Both know that if Player 2 refuses, she’ll have to offer Player 1 at least half of the remaining half, meaning that she’ll get at most one-quarter.)
- (b) With one period, Player 1 offers Player 2 a sliver, and Player 2 accepts. With two periods, Player 1 offers Player 2 two-thirds of the cake, and Player 2 accepts. (Both know that if Player 2 refuses, one-third of the cake melts, Player 2 will offer Player 1 a sliver of the remaining two-thirds, and Player 1 will accept.) With three periods, Player 1 offers Player 2 two-ninths of the cake, and Player 2 accepts. (Both know that if Player 2 refuses, she’ll have to offer Player 1 at least two-thirds of the remaining two-thirds, meaning that she’ll get at most two-ninths.)
- (c) This is the magic of the Coase Theorem. It is in neither player’s interest to let the cake melt away, so they have a strong incentive to figure things out at the beginning and bring about a Pareto efficient outcome. You can see the same phenomenon at work in labor disputes and lawsuits, many of which get settled before the parties really begin to hurt each other.

18.8 Share them with a classmate.

18.9 [“The Draft Game...”]

- (a) The naive outcome is for X to choose C, Y to choose H, and Z to choose T, producing the “naive outcome” at the top of the game tree.
- (b) No. If X and Y choose C and H, Z will choose F because this produces a better outcome for Z: FT is better than TQ! (But now backward induction kicks in: Y anticipates this, and so Y will choose G instead

of H—GH is better than HQ. But X anticipates this, and so knows that a choice of C will result in CQ. X then uses backward induction to solve the bottom half of the tree—Z will choose F in the top part and H in the lower part, so Y will choose H because HG is better than FG—and determine that a choice of T will result in TC. Since X prefers TC to CQ, X chooses T in the first round, leading Y to choose H and Z to choose F.

- (c) Backward induction leads to a result of (TC, HG, FQ).
- (d) This is not Pareto efficient: the “naive” strategies produce better outcomes for all three teams!
- (e) Statement #1 is false because each team’s choice in the first round will have strategic implications for its options in the second round. Statement #2 is true because each team’s choice in the second round has no further ramifications; since there are no more rounds, in the second round each team faces a simple decision tree.
- (f) This is a time-consuming problem. Thanks to Kieran Barr for finding two strategies that yield this same outcome!

18.10 [“The Hold-Up Problem...”]

- (a) The answer here depends on your assumptions. See below for one take on it.
- (b) The situation is Pareto inefficient.
- (c) The key issue here is that Butch Cassidy is a bank robber, and hence cannot be bound to contracts or other agreements. Sure, Harriman could pay him the money, but what guarantee does he have that this will make Butch stop robbing his train? A more likely outcome is that Butch will take the money and continue to rob the train, and then Harriman will be out even more money. So Harriman hires the superposse instead, even though both he and Butch would be better off with an alternative outcome.
- (d) Contracts can help by forcing players to act in certain ways; then the Coase theorem allows them to negotiate an efficient outcome. The Coase Theorem doesn’t work in the case of Butch Cassidy because he’s an outlaw: there’s no way to bind an outlaw to an enforceable contract.

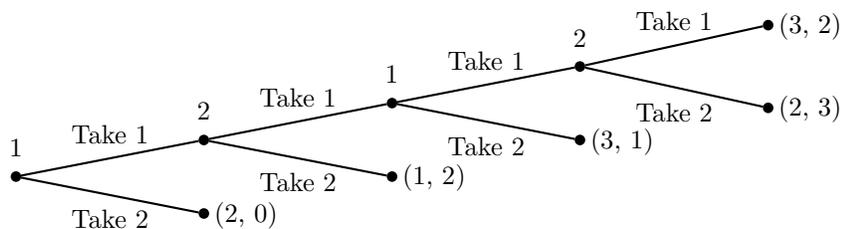
18.11 [“The IgoUgo travel guide...”]

- (a) If you make eye contact with the driver, the driver will pretend that she’s not going to stop, and then you’ll get scared and won’t go for it, so then the driver *won’t* stop.
- (b) The game tree here has you choosing to look or not look. If you choose not to look, the driver chooses to stop or not, and the payoffs

are obvious. If you choose to look, the driver chooses to stop or not, and in each of those situations you must choose whether or not to push the issue.

18.12 [“The Centipede Game...”]

- (a) Here is one way to draw the game tree; note that it looks a bit like a centipede.



- (b) As discussed in the text, backward induction predicts that Player 1 will immediately choose \$2 and end the game, yielding an outcome of $(2, 0)$.
- (c) No. There are many Pareto improvements, e.g., $(2, 1)$.
- (d) You can do this with induction; this exercise also suggests why backward induction has the name it does.
- (e) This is a truly difficult philosophical question. If you're interested, there's an interesting chapter (and a great bibliography) on this topic, in the guise of “the unexpected hanging”, in Martin Gardner's 1991 book, *The Unexpected Hanging, and Other Mathematical Diversions*.
- 18.13 Well, the exam can't be on Friday, because then on Thursday night you'd think, “Aha! The exam's got to be Friday!” So then you wouldn't be surprised; so the exam can't be on Friday. But then the exam can't be on Thursday, because then on Wednesday night you'd think, “Aha! The exam can't be on Friday, so it's got to be Thursday!” So then you wouldn't be surprised; so the exam can't be on Thursday. But then the exam can't be on Wednesday, or Tuesday, or even Monday. An apparently non-controversial statement by your teacher turns out to be quite treacherous!

18.14 [“The following game...”]

- (a) Backward induction predicts an outcome of $(8, 8)$.
- (b) No; a Pareto improvement is $(10, 10)$.

18.15 [“Consider the following version. . .”]

- (a) Player 1 moves first and has three options; in each case, Player 2 has two options (accept or reject). So there are 6 different possible outcomes.
- (b) If Player 2 is motivated solely by money, Player 1 can anticipate that Player 2 will accept any offer that he makes. If Player 1 is motivated solely by money, he will therefore offer Player 2 the minimum amount (\$1) required in order to get her to accept, thereby maximizing his financial payoff.
- (c)
 - i. This is true; otherwise, Player 2 would always have accepted Player 1’s offer.
 - ii. This is false. Player 1’s generosity (offering more than \$1) might be motivated by altruism, but it might also be motivated by a desire for money: if Player 1 thinks that Player 2 will turn down a lower offer, it’s in Player 1’s financial interest to offer more.
 - iii. The basic assumption may or may not be wrong, but this experiment didn’t show that it is wrong because optimizing individuals may not be solely motivated by money.

18.16 [“Consider the following. . .”]

- (a) With backward induction, the analysis begins at the end of the game. So: if the game reaches round 2, there are two ounces of cake left. Jill will offer Jack a tiny sliver, knowing that Jack will accept because his only alternative is to reject the offer and get nothing; so if the game reaches round 2, Jill will get a tiny bit less than two ounces of cake and Jack will get a tiny bit more than nothing. Using backward induction, we now look at round 1, where there are three ounces of cake. Jack has to offer Jill at least two ounces of cake, or Jill will reject his offer and go to round 2 (where, as we have seen, she can get two ounces). So we can predict that Jack will offer two ounces of cake to Jill, leaving one ounce for himself, and that Jill will accept the offer.
- (b) Settlement in round 1 results in a Pareto efficient outcome. The Coase theorem indicates that there is a strong incentive for both sides to settle these games in round 1 in order to reach a Pareto efficient outcome. In other words, there is a strong incentive to negotiate a labor agreement before a strike happens, or to settle a lawsuit before it goes to trial.

Chapter 19

Iterated dominance and Nash equilibrium

In Chapter 8 we examined simultaneous move games in which each player had a dominant strategy; the Prisoner's Dilemma game was one example. In many games, however, one or more players do not have dominant strategies. This chapter explores two solution concepts that we can use to analyze such games.

The first solution concept, iterated dominance, is a **refinement** of the dominant strategies approach from the previous chapter, meaning that iterated dominance is a stronger technique that builds upon (or refines) the results of the dominant strategies approach. In other words: the idea of dominant strategies often allows us to narrow down our prediction for the outcome of a game; iterated dominance allows us to narrow down our prediction at least as far, and sometimes further.

Unfortunately, this extra strength does not come for free. While dominant strategies is a reasonably simple idea, iterated dominance is (while not exactly a Nobel-prize-winning concept) one step closer to rocket science. As such, it requires more powerful assumptions about the intellectual capabilities of the optimizing individuals who are playing the games.

The second solution concept in this chapter, Nash equilibrium, is a refinement of iterated dominance: Nash equilibrium allows us to narrow down our prediction at least as far as iterated dominance, and sometimes further. Again, this extra strength does not come for free. Nonetheless, Nash equilibrium is one of the central concepts in the study of strategic behavior—a fact which helps explain why Nash equilibrium *is* a Nobel-prize-winning concept.

19.1 Iterated dominance

The transition from dominant strategies to iterated dominance involves two ideas. The first is this: even when a player doesn't have a dominant strategy (i.e., a *best* strategy, regardless of what the other players do), that player might

still have one strategy that dominates another (i.e., a strategy A that is *better* than strategy B, regardless of what the other players do). As suggested by the terms “best” and “better”, the difference here is between a superlative statement (e.g., “Jane is the best athlete in the class”) and a comparative statement (“Jane is a better athlete than Ted”); because comparatives are weaker statements, we can use them in situations where we might not be able to use superlatives.

For example, consider the game in Figure 19.1. First note that there are no strictly dominant strategies in this game: U is not the best strategy for Player 1 if Player 2 plays L or C, M is not the best strategy for Player 1 if Player 2 plays R, and D is not the best strategy for Player 1 if Player 2 plays L or C. Similarly, L is not the best strategy for Player 2 if Player 1 plays U or D, C is not the best strategy for Player 2 if Player 1 plays M, and R is not the best strategy for Player 2 if Player 1 plays U, M, or D.

Although there are no strictly dominant strategies, we can see that no matter what Player 1 does, Player 2 always gets a higher payoff from playing L than from playing R. We can therefore say that L **strictly dominates** R for Player 2, or that R is **strictly dominated** by L for Player 2. (Note that we cannot say that L is a **strictly dominant** strategy for Player 2—it does not dominate C—but we can say that R is a **strictly dominated** strategy for Player 2: an optimizing Player 2 would *never* play R.)

The second idea in the transition from dominant strategies to iterated dominance is similar to the backward induction idea of anticipating your opponents’ moves: players should recognize that other players have strictly dominated strategies, and should act accordingly. In our example, Player 1 should recognize that R is a strictly dominated strategy for Player 2, and therefore that there is no chance that Player 2 will play R. In effect, the game now looks like that shown in Figure 19.2 on the next page: the lines through the payoffs in the R column indicate that both players know that these payoffs have no chance of occurring because R is not a viable strategy for Player 2.

But now we see that Player 1 has an obvious strategy: given that Player 2 is never going to play R, Player 1 should always play M. Once R is out of the way, U and D are both dominated by M for Player 1: regardless of whether Player 2 plays L or C, Player 1 always gets his highest payoff by playing M. This is the idea of **iteration**, i.e., repetition. Combining this with the idea of dominated

		Player 2		
		L	C	R
Player 1	U	1, 10	3, 20	40, 0
	M	10, 20	50, -10	6, 0
	D	2, 20	4, 40	10, 0

Figure 19.1: A game without dominant strategies

		Player 2		
		L	C	R
Player 1	U	1, 10	3, 20	40, 10
	M	10, 20	50, -10	6, 10
	D	2, 20	4, 40	10, 10

Figure 19.2: Eliminating R, which is strictly dominated by L for Player 2

strategies gives us the process of **iterated dominance**: starting with the game in Figure 19.1, we look for a strictly dominated strategy; having found one (R), we eliminate it, giving us the game in Figure 19.2. We then repeat the process, looking for a strictly dominated strategy in that game; having found one (or, actually two: U and D), we eliminate them. A final iteration would yield (M, L) as a prediction for this game: knowing that Player 1 will always play M, Player 2 should always play L.

A complete example

Consider the game in Figure 19.3 below. There are no strictly dominant strategies, but there is a strictly dominated strategy: playing U is strictly dominated by D for Player 1. We can conclude that Player 1 will never play U, and so our game reduces to the matrix in Figure 19.4a on the next page.

But Player 2 should know that Player 1 will never play U, and if Player 1 never plays U then some of Player 2's strategies are strictly dominated! Namely, playing L and playing R are both strictly dominated by playing C as long as Player 1 never plays U. So we can eliminate those strategies for Player 2, yielding the matrix in Figure 19.4b. Finally, Player 1 should anticipate that Player 2 (anticipating that Player 1 will never play U) will never play L or R, and so Player 1 should conclude that M is strictly dominated by D (the matrix in Figure 19.4c). Using iterated strict dominance, then, we can predict that Player 1 will choose D and Player 2 will choose C.

		Player 2		
		L	C	R
Player 1	U	1,1	2,0	2,2
	M	0,3	1,5	4,4
	D	2,4	3,6	3,0

Figure 19.3: Iterated strict dominance example

	L	C	R
U	1,1	2,0	2,2
M	0,3	1,5	4,4
D	2,4	3,6	3,0

(a)

	L	C	R
U	1,1	2,0	2,2
M	0,3	1,5	4,4
D	2,4	3,6	3,0

(b)

	L	C	R
U	1,1	2,0	2,2
M	0,3	1,5	4,4
D	2,4	3,6	3,0

(c)

Figure 19.4: Solution to iterated strict dominance example

Question: Does the order of elimination matter?

Answer: Although it is not obvious, the end result of iterated strict dominance is always the same regardless of the sequence of eliminations. In other words, if in some game you can either eliminate U for Player 1 or L for Player 2, you don't need to worry about which one to "do first": either way you'll end up at the same answer.

A side note here is that this result only holds under iterated **strict** dominance, according to which we eliminate a strategy only if there is some other strategy that yields payoffs that are *strictly higher* no matter what the other players do. If you eliminate a strategy when there is some other strategy that yields payoffs that are *higher or equal* no matter what the other players do, you are doing iterated **weak** dominance, and in this case you will *not* always get the same answer regardless of the sequence of eliminations. (For an example see problem 19.10.) This is a serious problem, and helps explain why we focus on iterated strict dominance.

19.2 Nash equilibrium

Tenuous as it may seem, iterated strict dominance is not a very strong solution concept, meaning that it does not yield predictions in many games. An example is the game in Figure 19.5: there are no strictly dominant strategies and no strictly dominated strategies.

So game theorists have come up with other solution concepts. The most important one is called **Nash equilibrium** (abbreviated NE). A Nash equi-

		Player 2		
		L	C	R
Player 1	U	5,1	2,0	2,2
	M	0,4	1,5	4,5
	D	2,4	3,6	1,0

Figure 19.5: Nash equilibrium example

librium occurs when the strategies of the various players are best responses to each other. Equivalently but in other words: given the strategies of the other players, each player is acting optimally. Equivalently again: No player can gain by deviating alone, i.e., by changing his or her strategy single-handedly.

In the game in Figure 19.5, the strategies (D, C) form a Nash equilibrium: if Player 1 plays D, Player 2 gets her best payoff by playing C; and if Player 2 plays C, Player 1 gets his best payoff by playing D. So the players' strategies are best responses to each other; equivalently, no player can gain by deviating alone. (Question: Are there any other Nash equilibria in this game?)

Algorithms for finding Nash equilibria

The best way to identify the Nash equilibria of a game is to first identify all of the outcomes that are *not* Nash equilibria; anything left must be a Nash equilibrium. For example, consider the game in Figure 19.5. The strategy pair (U, L) is not a Nash equilibrium because Player 2 can gain by deviating alone to R; (U, C) is not a NE because Player 1 can gain by deviating alone to D (and Player 2 can gain by deviating alone to L or R); etc. If you go through the options one by one and cross out those that are *not* Nash equilibria, the remaining options *will* be Nash equilibria (See Figure 19.6a).

A shortcut (but one you should use carefully!) is to underline each player's best responses.¹ To apply this to the game in Figure 19.5, first assume that Player 2 plays L; Player 1's best response is to play U, so underline the "5" in the box corresponding to (U, L). Next assume that Player 2 plays C; Player 1's best response is to play D, so underline the "3" in the box corresponding to (D, C). Finally, assume that Player 2 plays R; Player 1's best response is to play M, so underline the "4" in the box corresponding to (M, R). Now do the same thing for Player 2: go through all of Player 1's options and underline the best response for Player 2. (Note that C and R are both best responses when Player 1 plays M!) We end up with Figure 19.6b: the only boxes with both payoffs underlined are (D, C) and (M, R), the Nash equilibria of the game.

	L	C	R
U	5 / 1	2 / 0	2 / 2
M	0 / 4	1 / 5	4,5
D	2 / 4	3,6	1 / 0

(a)

	L	C	R
U	<u>5</u> ,1	2,0	2, <u>2</u>
M	0,4	1, <u>5</u>	<u>4</u> , <u>5</u>
D	2,4	<u>3</u> , <u>6</u>	1,0

(b)

Figure 19.6: Finding Nash equilibria: (a) with strike-outs; (b) with underlinings

¹It is easy to confuse the rows and columns and end up underlining the wrong things. Always double-check your answers by confirming that no player can gain by deviating alone.

Some history

Nash equilibrium is one of the fundamental concepts of game theory. It is named after John Nash, a mathematician born in the early part of this century. He came up with his equilibrium concept while getting his Ph.D. in mathematics at Princeton, then got a professorship at MIT, then became mentally ill—e.g., claimed that aliens were sending him coded messages on the front page of the *New York Times*—then spent many years in and out of various mental institutions, then slowly got on the road to recovery, then won the Nobel Prize in Economics in 1994, and now spends his time at Princeton playing with computers. You can read more about him in a fun book called *A Beautiful Mind* by Sylvia Nasar.²

19.3 Infinitely repeated games

We saw in the last chapter that there’s no potential for cooperation (at least in theory) if we play the Prisoner’s Dilemma game twice, or 50 times, or 50 million times. What about infinitely many times? To examine this possibility, we will look at the Prisoners’ Dilemma game in Figure 19.7.

		Player 2	
		D	C
Player 1	D	0,0	10,-5
	C	-5,10	1,1

Figure 19.7: Another version of the Prisoners’ Dilemma

We must first figure out exactly what it means to win (or lose) this game infinitely many times. Here it helps to use the present value concepts from Chapter 1: with an interest rate of 5%, winning \$1 in each round does *not* give you infinite winnings. Rather, the present value of your winnings (using the perpetuity formula, assuming you get paid at the end of each round) is $\frac{\$1}{.05} = \20 .

So: with an interest rate of r we can ask meaningful questions about the potential for cooperation. One point that is immediately clear is that there is still plenty of potential for *non-cooperation*: the strategies of playing (D, D) forever continue to constitute a Nash equilibrium of this game.

But perhaps there are other strategies that are also Nash equilibria. Because the game is played infinitely many times, we cannot use backward induction to

²There is also a movie of the same name, starring Russell Crowe. Unfortunately, it takes some liberties with the truth; it also does a lousy job of describing the Nash equilibrium concept.

solve this game. Instead, we need to hunt around and look for strategies that might yield a cooperative Nash equilibrium.

One potentially attractive idea is to use a **trigger strategy**: begin by cooperating and assuming that the other player will cooperate (i.e., that both players will play C), and enforce cooperation by threatening to return to the (D, D) equilibrium. Formally, the trigger strategy for each player is as follows: In the first stage, play C. Thereafter, if (C, C) has been the result in all previous stages, play C; otherwise, play D.

We can see that the cooperative outcome (C, C) will be the outcome in each stage game if both players adopt such a trigger strategy. But do these strategies constitute a Nash equilibrium? To check this, we have to see if the strategies are best responses to each other. In other words, given that Player 1 adopts the trigger strategy above, is it optimal for Player 2 to adopt a similar trigger strategy, or does Player 2 have an incentive to take advantage of Player 1?

To find out, let's examine Player 2's payoffs from cooperating and from deviating:

If Player 2 cooperates, she can expect to gain \$1 at the end of each round, yielding a present value payoff of $\frac{\$1}{r}$. (If $r = .05$ this turns out to be \$20.)

If Player 2 tries to cheat Player 1 (e.g., by playing D in the first round), Player 2 can anticipate that Player 1 will play D thereafter, so the best response for Player 2 is to play D thereafter as well. So the best deviation strategy for Player 2 is to play D in the first round (yielding a payoff of \$10 since Player 1 plays C) and D thereafter (yielding a payoff of \$0 each round since Player 1 plays D also). The present value of all this is simply \$10.

We can now compare these two payoffs, and we can see that cooperating is a best response for Player 2 as long as $\frac{\$1}{r} \geq 10$. Since the game is symmetric, cooperating is a best response for Player 1 under same condition, so we have a Nash equilibrium (i.e., mutual best responses) as long as $\frac{\$1}{r} \geq 10$. Solving this yields a critical value of $r = .1$. When r is below this value (i.e., the interest rate is less than 10%), cooperation is possible. When r is above this value (i.e., the interest rate is greater than 10%), cheating is too tempting and the trigger strategies do not form a Nash equilibrium. The intuition here is quite nice: By cooperating instead of deviating, Player 2 accepts lower payoffs now (1 instead of 10) in order to benefit from higher payoffs later (1 instead of 0). Higher interest rates make the future less important, meaning that Player 2 benefits less by incurring losses today in exchange for gains tomorrow. With sufficiently high interest rates, Player 2 will take the money and run; but so will Player 1!

19.4 Mixed strategies

Figure 19.8 shows another game, called the Battle of the Sexes. In this game, Player 1 prefers the opera, and Player 2 prefers wrestling, but what both players really want above all is to be with each other. They both choose simultaneously, though, and so cannot guarantee that they'll end up together. (Imagine, for example, that they are at different work places and can't reach each other and must simply head to one of the two events after work and wait for the other person at will-call.)

The Nash equilibriums of this game are (Opera, Opera) and (WWF, WWF). But there is another Nash equilibrium that is perhaps a little better at predicting reality: that equilibrium is for both players to play a **mixed strategy**, i.e., to choose different strategies with various probabilities. (In this case, the mixed strategy equilibrium is for Player 1 to choose opera with probability $2/3$ and WWF with probability $1/3$, and for Player 2 to choose opera with probability $1/3$ and WWF with probability $2/3$. You should be able to use what you've learned about expected value to show that these are mutual best responses.) One of the main results from game theory is that every finite game has at least one Nash equilibrium. That Nash equilibrium may only exist in mixed strategies, as in the following example.

Example: (Matching pennies, Figure 19.9). Players 1 and 2 each have a penny, and they put their pennies on a table simultaneously. If both show the same face (both heads or both tails), Player 2 must pay \$1 to Player 1; if one is heads and the other is tails, Player 1 must pay \$1 to Player 2.

Not surprisingly, the only NE in this game is for each player to play heads with probably $1/2$ and tails with probability $1/2$.

Problems

Answers are in the next section. The first few words of each question (and each answer) are hyperlinked to allow online users to quickly move back and forth.

- 19.1 *Challenge.* Explain (as if to a non-economist) why iterated dominance make sense.

		Player 2	
		Opera	WWF
Player 1	Opera	2,1	0,0
	WWF	0,0	1,2

Figure 19.8: The battle of the sexes

- 19.2 *Super Challenge.* Explain (as if to a non-economist) why Nash equilibrium makes sense.
- 19.3 Show that there are no strictly dominant strategies in the game in Figure 19.3 on page 191.
- 19.4 The game “Rock, Paper, Scissors” works as follows: You and your opponent simultaneously choose rock, paper, or scissors. If you pick the same one (e.g., if you both pick rock), you both get zero. Otherwise, rock beats scissors, scissors beats paper, and paper beats rock, and the loser must pay the winner \$1.
- Write down the payoff matrix for this game.
 - Does iterated dominance help you solve this game?
 - Calculus/Challenge.* Can you find any *mixed strategy* Nash equilibria?
- 19.5 Analyze games (a) through (e) on the following page(s). First see how far you can get using iterated dominance. Then find the Nash equilibrium(s). If you can identify a unique outcome, determine whether it is Pareto efficient. If it is not, identify a Pareto improvement.
- Game a.
 - Game b.
 - Game c.
 - Game d.
 - Game e.
- 19.6 *Challenge.* Prove that the pure strategy Nash equilibrium solutions are a subset of the iterated dominance solutions, i.e., that iterated dominance never eliminates any pure strategy Nash equilibrium solutions.
- 19.7 Rewrite Story #1 from the Overinvestment Game (problem 18.5 on page 177) as a simultaneous move game and identify the (pure strategy) Nash equilibria. Does your answer suggest anything about the relationship between backward induction and Nash equilibrium?

		Player 2	
		Heads	Tails
Player 1	Heads	1,-1	-1,1
	Tails	-1,1	1,-1

Figure 19.9: Matching pennies

		Player 2		
		L	C	R
Player 1	U	0,3	2,1	5,0
	M	4,8	3,2	8,3
	D	3,7	6,3	6,8

		Player 2		
		L	C	R
Player 1	U	-1,4	7,3	5,2
	M	2,0	5,-1	6,2
	D	1,2	1,0	1,0

		Player 2		
		L	C	R
Player 1	U	1,0	7,3	2,1
	M	1,0	1,2	6,2
	D	1,2	1,-3	1,0

		Player 2		
		L	C	R
Player 1	U	3,-1	5,4	3,2
	M	-2,5	1,3	2,1
	D	3,3	3,6	3,0

		Player 2		
		L	C	R
Player 1	U	3,-1	1,0	-1,-1
	M	1,-5	6,3	-7,-5
	D	-8,-10	-1,-3	-1,-1

- 19.8 *Challenge.* Prove that backward induction solutions are a subset of Nash equilibrium solutions, i.e., that any backward induction solution is also a Nash equilibrium solution. (Note: Backward induction is in fact a refinement of Nash equilibrium called “subgame perfect Nash equilibrium”.)
- 19.9 *Fun/Challenge.* Section 19.3 describes a trigger strategy for yielding cooperating in the infinitely repeated Prisoner’s Dilemma game shown in Figure 19.7. Can you think of another strategy that yields even higher payoffs for the players? Can you show that it’s a Nash equilibrium?
- 19.10 *Challenge.* The end of the section on iterated dominance mentioned the dangers of iterated weak dominance, namely that different sequences of elimination can yield different predictions for the outcome of a game. Show this using the game in Figure 19.10. (Hint: Note that U is weakly dominated by M for Player 1 and that M is weakly dominated by D for Player 1.)

		Player 2	
		L	R
Player 1	U	50, 10	6, 20
	M	50, 10	8, 9
	D	60, 15	8, 15

Figure 19.10: The dangers of iterated weak dominance

Answers

- 19.1 This is explained to the best of my abilities in the text. The key idea is to anticipate your opponent’s behavior.
- 19.2 This is a hard philosophical problem.
- 19.3 U is not the best strategy for Player 1 if Player 2 plays R, M is not the best strategy for Player 1 if Player 2 plays C, and D is not the best strategy for Player 1 if Player 2 plays R. Similarly, there are no strictly dominant strategies for Player 2: L is not the best strategy for Player 2 if Player 1 plays U, C is not the best strategy for Player 2 if Player 1 plays U, and R is not the best strategy for Player 2 if Player 1 plays D.
- 19.4 [“The game “Rock, Paper, Scissors”...”]
- (a) The payoff matrix is shown in Table 19.1.
 - (b) Iterated dominance does not help you solve this game because there are no dominated strategies.

		Player 2		
		R	P	S
Player 1	R	0,0	-1,1	1,-1
	P	1,-1	0,0	-1,1
	S	-1,1	1,-1	0,0

Table 19.1: The payoff matrix for the game “Rock, Paper, Scissors”

- (c) In accordance with intuition, the NE is for both players to choose randomly among the three strategies.

19.5 [“Analyze games (a) through (e)...”]

- (a) Here U is dominated by M for player 1, then C is dominated by L (or R) for player 2, then D is dominated by M for player 1, then R is dominated by L for player 2. The result: (M, L), with a payoff of (4,8). This is also the unique Nash equilibrium of the game; it is not a Pareto efficient outcome because of (D, R).
- (b) Here D is dominated by M for player 1, then C is dominated by L for player 2, then U is dominated by M for player 1, then L is dominated by R for player 2. The result: (M, R), with a payoff of (6,2). This is also the unique Nash equilibrium of the game; it is not a Pareto efficient outcome because of (U, C).
- (c) There are no strictly dominated strategies. The NE are (U, C), (D, L) and (M, R).
- (d) Here M is dominated by U for player 1, then L and R are dominated by C for player 2, then D is dominated by U for player 1. The result: (U, C), with a payoff of (5,4). This is also the unique Nash equilibrium of the game; it is a Pareto efficient outcome.
- (e) Here L is dominated by C for player 2, and that is as far as iterated dominance can take us. We do not get a unique prediction for the outcome of this game. (All that we can say is that a rational player 2 would never play L.) With Nash, the NE are (M, C) and (D, R). Note that these Nash equilibria are a subset of the iterated dominance solutions; see the next problem for details.

19.6 This is a tough problem.

- 19.7 The Nash equilibria are (War, Stay Out) and (Peace, Enter). This suggests that backward induction is in fact a refinement or strengthening of Nash equilibrium (which it is, namely **subgame perfect Nash equilibrium**).

		Entrant	
		Enter	Stay Out
Monopolist	War	10, -10	100, 0
	Peace	35, 5	100, 0

19.8 This is a hard problem.

19.9 Yes! As long as the interest rate is sufficiently low, the players can also cooperate by taking turns: (C, D), (D, C), (C, D),... Instead of gaining \$1 every stage (which is the result with the trigger strategies), each player now gains \$5 every two stages. As an exercise, you can formally define this strategy (what happens if the other player doesn't cooperate?) and determine the interest rates that allow this strategy as a Nash equilibrium and those that make this strategy a Pareto improvement over the trigger strategy. There are also plenty of other strategies, e.g., tit-for-tat, that you can play around with if you wish.

19.10 One possibility is to proceed as follows: U is weakly dominated by M for Player 1, and then R is weakly dominated by L for Player 2, and then M is (strictly) dominated by D for Player 1, yielding a prediction of (D, L). But another possibility is to proceed like this: M is weakly dominated by D for Player 1, then L is weakly dominated by R for Player 2, then U is (strictly) dominated by D for Player 1, yielding a prediction of (D, R). Conclusion: the order of elimination matters for iterated weak dominance!

Chapter 20

Supply and demand details

The great thing about supply and demand is that the basic ideas are relatively easy and extremely powerful. The previous chapters analyze a wide range of situations using only three concepts: market supply curves, market demand curves, and market equilibrium. This chapter looks at some details of supply and demand, with a focus on these questions:

- Where do market supply and demand curves come from, and how do they relate to the assumption in economics that decisions are made by individuals and not by groups?
- Why do market supply and demand curves have the shape they do? In particular, why do economists usually assume that market demand curves are downward sloping, and that market supply curves are upward sloping?

20.1 Deconstructing supply and demand

The answer to the first question is this: *market* supply and demand curves come from *individual* supply and demand curves, and individual supply and demand curves come from individual optimization.

Recall that a **market supply curve** is a graph that answers the following question: *If* the market price were p , how many units of this good would sellers want to sell? An **individual supply curve** (say, for a particular firm) is a graph that answers the same question for that particular firm: *If* the market price were p , how many units of this good would that firm want to sell? For each price p , we can imagine the firm using individual optimization and a decision tree to consider all the options (sell 4, sell 5, sell 6, . . .) and then picking the best one. The answer that the firm comes up with at any particular price (e.g., “At a price of \$6, I can maximize my profit by selling 4”) gives us *one point* on that firm’s individual supply curve.

The same ideas apply to the demand side. Recall that a **market demand curve** is a graph that answers the following question: *If* the market price were

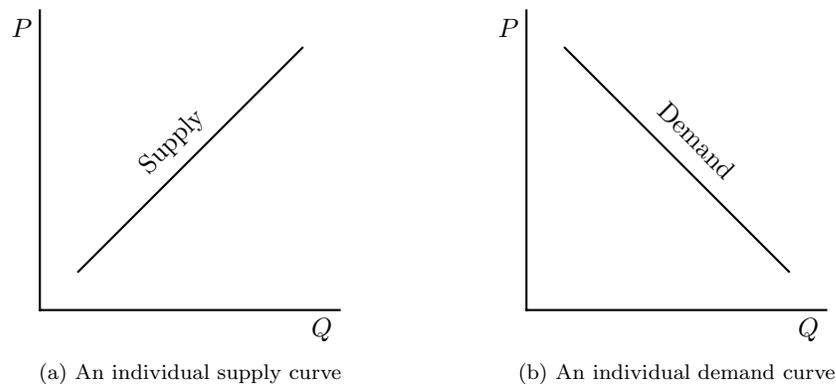


Figure 20.1: Individual supply and demand curves

p , how many units of this good would buyers want to buy? An **individual demand curve** (say, for a particular person) is a graph that answers the same question for that particular person: *If* the market price were p , how many units of this good would that person want to buy? Again, the link to individual optimization is that we can imagine that person answering this question by using a decision tree to write down all the different options (buy 4, buy 5, buy 6, ...) and picking the best one. The answer that that person comes up with for any particular price (e.g., “At a price of \$8, I would want to buy 6”) gives us *one point* on that person’s individual demand curve.

Note that individual supply and demand curves (see Figure 20.1) look just like market supply and demand curves, the only difference being that individual curves are scaled down. All the sellers together may want to sell 100,000 units at a price of \$5 per unit, but one individual seller may only want to sell 100 units at that price. Similarly, all the buyers together may want to buy 100,000 units at a price of \$5 per unit, but one individual buyer may only want to buy 10 units at that price.

Reconstructing supply and demand

Beginning with optimizing individuals, each with an individual supply or demand curve, we can aggregate over all the individual sellers to get the market supply curve, as in Figure 20.2, and over all the individual buyers to get the market demand curve, as in Figure 20.3.

20.2 *Math*: The algebra of markets

Algebra allows us to easily aggregate individual demand or supply curves into market demand or supply curves. If, for example, there are 500 consumers, each of whom wants to buy 3 units at a price of \$10, then the buyers as a whole want to buy $500 \cdot 3 = 1500$ units at a price of \$10.

If the price were	0	1	2	3	4	5
Firm 1 would want to sell	0	1	2	3	4	5
Firm 2 would want to sell	0	2	4	6	8	10
Together they would want to sell	0	3	6	9	12	15

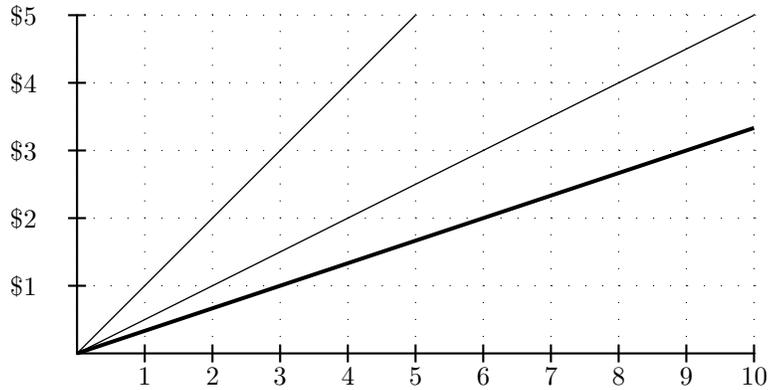


Figure 20.2: Aggregating supply

If the price were	0	1	2	3	4	5
Person 1 would want to buy	10	8	6	4	2	0
Person 2 would want to buy	5	4	3	2	1	0
Together they would want to buy	15	12	9	6	3	0

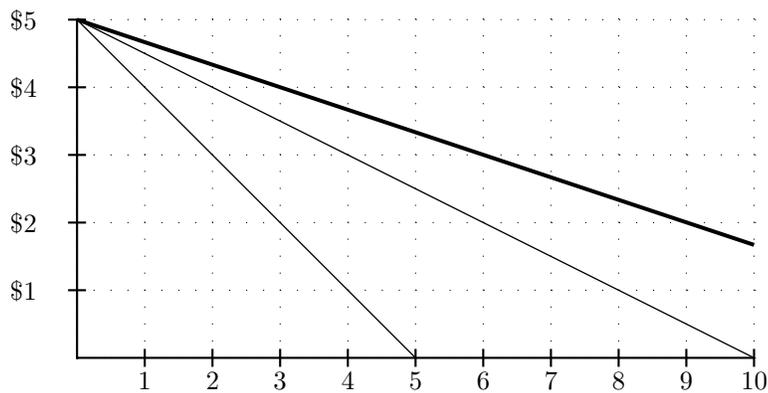


Figure 20.3: Aggregating demand

Mathematically, the market demand curve is simply the summation of all the individual demand curves. For example, if there are 500 consumers in an economy, each with an individual demand curve of $q_i = 15 - p$, then the total demand from the market q_M is

$$q_M = q_1 + q_2 + \dots + q_{500} = 500(15 - p) \implies q_M = 7500 - 500p.$$

The same approach works for market supply curves, which are simply summations of individual supply curves. If there are 400 suppliers with individual supply curves of $q_i = 15 + 2p$, then the market supply curve is given by

$$q_S = q_1 + q_2 + \dots + q_{400} = 400(15 + 2p) = 6000 + 800p.$$

The same process works if there are multiple types of buyers or sellers. For example, if there are 500 consumers with individual demand curves of $q_i = 15 - p$ and 300 consumers with individual demand curves of $q_i = 30 - 2p$, then the total demand from the market is

$$q_M = 500(15 - p) + 300(30 - 2p) = 16500 - 1100p.$$

At a price of \$10, the buyers want to buy $16500 - 1100 \cdot 10 = 5500$ units. Each of the 500 buyers with an individual demand curves of $q_i = 15 - p$ wants to buy $15 - 10 = 5$ units, for a total of 2500. And each of the 300 buyers with individual demand curves of $q_i = 30 - 2p$ wants to buy $30 - 2 \cdot 10 = 10$ units, for a total of 3000.

20.3 On the shape of the demand curve

We are now better able to answer the second question posed at the beginning of this chapter: Why do market supply and demand curves have the shape they do? We begin with the demand side, where a common assumption is that market demand curves are downward sloping, i.e., that as prices fall buyers want to buy more and that as prices rise buyers want to buy less. A correct (if unsatisfying) explanation is that we assume that *market* demand curves are downward sloping because we assume that *individual* demand curves are downward sloping. Aggregating a bunch of downward-sloping individual demand curves produces a downward-sloping market demand curve.

Of course, this simply leads to the question of why we assume *individual* demand curves are downward sloping? There are two possible explanations:

The substitution effect suggests that a drop in price makes that good look more attractive relative to other goods. If Coke is on sale, I'm going to buy less Pepsi and more Coke, i.e., I'm going to **substitute out of** Pepsi and **substitute into** Coke.

The income effect suggests that a drop in price is in some ways similar to my having more income. Because the price of Coke has fallen, I can now afford more of everything, and in particular I can afford more Coke.

The substitution effect is usually more important than the income effect, especially for low-budget items like soda. The dominance of the substitution effect is a good thing, because it is the substitution effect that underlies our assumption about downward-sloping demand curves: as the price of some good goes up *relative to its substitutes*, the substitution effect causes people to buy less; as its relative price goes down, the substitution effect causes people to buy more.

Unfortunately, that is not all there is to it. Although the income effect is unlikely to be important in practice, there is no such restriction in theory. And this leads to a terrible complication. It turns out that the impact of the income effect is theoretically unclear. For **normal goods**, the income effect reinforces the substitution effect: as in the Coke example above, a reduction in the price of Coke effectively boosts my income, which leads me to buy more Coke. But for **inferior goods** the income effect works in the opposite direction: by effectively boosting my income, lower prices for inferior goods lead me to buy *less* of them. (An example here might be Ramen noodles: as incomes rise, many individuals buy less Ramen, not more.) And since it is theoretically possible for the income effect to be more important than the substitution effect, *it is theoretically possible for demand curves to be upward sloping*: as the price of such a **Giffen good** goes down, the impact on your income—i.e., the income effect—is so strong that you end up buying *less* of that good. Equivalently, as the price of a Giffen good goes up, the income effect is so strong that you end up buying *more* of that good.

Although all of this has significant implications for economic theory, the practical implications are pretty minor. Economists argue about whether there has ever been *even a single real-life situation* featuring an upward sloping demand curve. (The focus of many such arguments is the Irish potato famine of 1845–49: potatoes were such a large portion of household expenditures that when potato prices rose, households arguably had to cut back on other food expenditures, leading them to buy even more potatoes.) For practical purposes, then, it is perfectly reasonable to assume—as we will throughout this book—that demand curves are downward sloping.

20.4 On the shape of the supply curve

We now turn to the supply side, where our task is to analyze the assumption that market supply curves are upward sloping. More precisely, our task is to analyze the assumption that market supply curves are not downward sloping. (Recall that we sometimes assume that supply curves are vertical or horizontal lines—perfectly inelastic or perfectly elastic—as in Figure 14.2.) As on the demand side, we begin with a correct (if unsatisfying) answer: we assume that *market* supply curves are not downward sloping because we assume that *individual* supply curves are not downward sloping. Aggregating a bunch of individual supply curves that are not downward sloping produces a market supply curve that is not downward sloping.

The obvious follow-up question is: why do we assume that individual supply

curves are not downward sloping? The reason is that *it is theoretically impossible for a profit-maximizing firm to have a downward sloping supply curve*. A downward sloping supply curve would mean, for example, that a firm would want to sell 1,000 units at a price of \$1 per unit, but only 500 units at a price of \$2 per unit. This sort of behavior is incompatible with profit maximization: if the firm maximizes profit by producing 1,000 units at a price of \$1 per unit, it must produce *at least* that many units in order to maximize profit at a price of \$2 per unit. (For a mathematical proof of this, see problem 20.2.)

20.5 Comparing supply and demand

A close reading of the preceding material in Part III suggests that supply curves and demand curves have much in common. Indeed, many of the analyses of demand curves have followed essentially the same path as the analyses of supply curves. In some cases the only difference is that the words “supply”, “sell”, and “sellers” were replaced with “demand”, “buy”, and “buyers”. What, then, are the differences (if any) between supply and demand?

One apparent difference is that we usually think of demand curves as pertaining to people, and of supply curves as pertaining to firms. It turns out, however, that the important differences that do exist between optimizing firms and optimizing individuals—notably, that firms have more flexibility to enter or exit markets in the long run—do not usually smoothly translate into differences between supply and demand.¹

The reason is that the connections between demand curves and people—and between supply curves and firms—are not as clear cut as they appear. Consider, for example, the labor market: here the sellers are individual people, and the buyers are individual firms. Another example on the supply side is the housing market, where the sellers are individual people rather than firms. On the demand side, many market demand curves—such as those for electricity, shipping, and paper—reflect demand from individual firms instead of (or in addition to) demand from individual people. Firms’ demand curves for labor, raw materials, and other factors of production are called **factor demand curves**.

Close scrutiny suggests that even the “obvious” difference—that supply curves are about selling and demand curves are about buying—is not such a big deal. In fact, this distinction disappears entirely in the context of a **barter economy**. Instead of using money as a medium for exchange, such an economy relies on the direct exchange of goods. (In the famous example from the story *Jack and the Beanstalk*, Jack trades the family cow for some beans.) Barter economies do not differentiate between buyers and sellers: Jack is both a seller of cows and a buyer of beans. This perspective is hidden in money-based economies—we usually think about buying a car for \$3,000, not about selling \$3,000 for a car—but it is still valuable. The value is in the idea that *there is*

¹One exception is the treatment of supply and demand in extremely short-run or long-run situations. Long-run supply curves are often assumed to be perfectly elastic, as in Figure 14.2b. Demand curves, in contrast, do not exhibit perfect elasticity in the long run.

really no difference between supply and demand. They are two sides of the same coin: the same terminology (e.g., elasticities) applies to both, and the analysis on one side (e.g., taxes on the sellers) mirrors the analysis on the other side. In educational terms, this means that you get two for the price of one: master how either supply or demand works and mastery of the other should follow close behind.

Problems

Answers are in the next section. The first few words of each question (and each answer) are hyperlinked to allow online users to quickly move back and forth.

- 20.1 Consider a world with 1,000 buyers: 500 of them have an individual demand curve of $q = 20 - 2p$, and 500 of them have an individual demand curve of $q = 10 - 5p$. There are also 500 sellers, each with an individual supply curve of $q = 6p - 10$. Determine the market demand curve and the market supply curve.
- 20.2 *Challenge* Show that it is theoretically impossible for a profit-maximizing firm to have a downward-sloping supply curve. To put this into mathematical terms, consider high and low prices (p^H and p^L , with $p^H > p^L$) and high and low quantities (q^H and q^L , with $q^H > q^L$). With a downward-sloping supply curve, q^H would be the profit-maximizing quantity at p^L and q^L would be the profit-maximizing quantity at p^H . Your job is to show that this is not possible. (*Hint*: Let $C(q)$ be the cost of producing q units of output, so that profit is $pq - C(q)$. To show that downward sloping supply curves are impossible, assume that q^H maximizes profit at market price p^L and then show that a firm facing market price p^H will make more profit by producing q^H instead of q^L .)
- 20.3 Consider a world with 300 consumers, each with demand curve $q = 25 - 2p$, and 500 suppliers, each with supply curve $q = 5 + 3p$.
- Calculate the equations for the market demand curve and the market supply curve.
 - Determine the market equilibrium price and quantity and the total revenue in this market.
 - Calculate the price elasticity of demand and the price elasticity of supply at the market equilibrium. Then calculate their ratio.
 - Imagine that the government imposes a \$1 per-unit tax on the buyers. Write down the new market supply and demand curves, and find the new market equilibrium price and quantity. How much of the

tax burden is borne by the buyers, and how much by the sellers? Calculate the ratio of these tax burdens, and compare with the ratio of the elasticities calculated above.

- (e) Now imagine that the government instead decides to impose a \$1 per-unit tax on the sellers. How will this change things? Write down the new market supply and demand curves, find the new market equilibrium price and quantity, and compare with your answer from above (where the tax is on the buyer).
- (f) Now imagine that the government instead decides to impose a 50% sales tax on the sellers. Find the new market equilibrium price and quantity.
- (g) Finally, imagine that the government instead decides to impose a 100% sales tax on the buyers. Find the new market equilibrium price and quantity, and compare with your answer from above (where the tax is on the seller).

Answers

20.1 The market demand curve is

$$q = 500(20 - 2p) + 500(10 - 5p) = 15000 - 3500p.$$

The market supply curve is

$$q = 500(6p - 10) = 3000p - 5000.$$

20.2 Let the relevant variables be $p^H > p^L$ and $q^H > q^L$. A downward-sloping supply curve means q^L is optimal at the higher price (so that $p^H q^L - C(q^L)$ maximizes profit at price p^H) but that q^H is optimal at the lower price (so that $p^L q^H - C(q^H)$ maximizes profit at price p^L). To proceed by contradiction, note that profit maximization at the lower market price yields

$$p^L q^H - C(q^H) \geq p^L q^L - C(q^L).$$

It follows that q^L is not profit-maximizing at the higher price:

$$\begin{aligned} p^H q^H - C(q^H) &\geq (p^H - p^L)q^H + p^L q^L - C(q^L) \\ &= (p^H - p^L)(q^H - q^L) + p^H q^L - C(q^L) \\ &> p^H q^L - C(q^L). \end{aligned}$$

20.3 [“Consider a world...”]

- (a) The market demand curve is $q = 300(25 - 2p)$, i.e., $q = 7500 - 600p$. The market supply curve is $q = 500(5 + 3p)$, i.e., $q = 2500 + 1500p$.

- (b) Solving simultaneously we get $p = \frac{50}{21} \approx 2.38$ and $q \approx 6071$. Total revenue is therefore $pq \approx (2.38)(6071) \approx 14450$.
- (c) To calculate the price elasticity of demand at the market equilibrium, we take the market equilibrium as point A ($p_A = 2.38, q_A = 6071$) and any other point on the demand curve as point B (e.g., $p_B = 10$, from which the market demand curve gives $q_B = 7500 - 6000 = 1500$) and then plug them into the formula:

$$\varepsilon_D = \frac{q_B - q_A}{p_B - p_A} \frac{p_A}{q_A} = \frac{1500 - 6071}{10 - 2.38} \frac{2.38}{6071} \approx -0.235.$$

To calculate the price elasticity of supply at the market equilibrium, we take the same point A and any other point on the supply curve as point B (e.g., $p_B = 1$, from which the market supply curve gives $q_B = 2500 + 1500 = 4000$) and then plug them into the formula:

$$\varepsilon_S = \frac{q_B - q_A}{p_B - p_A} \frac{p_A}{q_A} = \frac{4000 - 6071}{1 - 2.38} \frac{2.38}{6071} \approx 0.588.$$

The ratio of the elasticities is $\frac{\varepsilon_D}{\varepsilon_S} = \frac{-0.235}{0.588} = -0.4$.

- (d) With a \$1 per-unit tax on the buyers, the market demand curve becomes $q = 7500 - 600(p + 1)$, i.e., $q = 6900 - 600p$. The market supply curve is still $q = 2500 + 1500p$, so the new equilibrium is at $p = \frac{44}{21} \approx 2.10$ and $q \approx 5643$. The seller therefore receives \$2.10 for each unit, and the buyer pays a total of $2.10 + 1.00 = \$3.10$ for each unit. Compared with the original equilibrium price of \$2.38, the seller is worse off by $2.38 - 2.10 = .28$, and the buyer is worse off by $3.10 - 2.38 = .72$. The ratio of these two is $\frac{.72}{.28} \approx 2.5$. Since $(0.4)^{-1} = 2.5$, the tax burden ratio is the negative inverse of the ratio of the elasticities.
- (e) With a \$1 per-unit tax on the sellers, the market supply curve becomes $q = 2500 + 1500(p - 1)$, i.e., $q = 1000 + 1500p$. The market demand curve is, as originally, $q = 7500 - 600p$. So the new equilibrium is at $p = \frac{65}{21} \approx \3.10 and $q \approx 5643$. This shows the tax equivalence result: the buyers and sellers end up in the same spot regardless of whether the tax is placed on the buyer or the seller.
- (f) With a 50% sales tax on the sellers, the market supply curve becomes $q = 2500 + 1500(.5p)$, i.e., $q = 2500 + 750p$. The demand curve is, as originally, $q = 7500 - 600p$. So the new equilibrium is at $p = \frac{500}{135} \approx \3.70 and $q \approx 5278$.
- (g) With a 100% sales tax on the buyers, the demand curve becomes $q = 7500 - 600(2p)$, i.e., $q = 7500 - 1200p$. The supply curve is, as originally, $q = 2500 + 1500p$. So the new equilibrium is at $p = \frac{50}{27} \approx \1.85 and $q \approx 5278$. This shows the tax equivalence result for sales taxes: the buyers and sellers end up in the same spot regardless of

whether there's a 50% sales tax on the sellers or a 100% sales tax on the buyers.

Glossary

ad valorem tax A tax based on value or sale amount, as in a per-dollar **sales tax**. (Compare with **per-unit tax**)

annuity A stream of annual payments for a finite number of years, as in a 20-year lottery jackpot payment. (Compare with **perpetuity**)

arbitrage [An attempt] to profit by exploiting price differences of identical or similar financial instruments, on different markets or in different forms. The ideal version is *riskless arbitrage* (investorwords.com).

ascending price auction See Chapter 9 for descriptions of different kinds of auctions.

auction A method of selling an item that involves pitting potential buyers against each other. See Chapter 9 for descriptions of the different kinds of auctions.

backward induction A solution concept for **sequential move games** that involves reasoning backward from the end of the **game tree** to the beginning.

barrier to entry A legal or economic barrier that protects a monopoly by preventing other firms from entering the market.

capital theory The branch of microeconomics dealing with investment decisions.

collective action problem A game, such as the **prisoners' dilemma**, in which decisions that are individually optimal lead to results that are collectively sub-optimal (and, in particular, to results that are **Pareto inefficient**).

competitive market A market with many buyers, each small in relation to all the buyers together, and many sellers, each small in relation to all the sellers together. Also called (somewhat redundantly) a **perfectly competitive market**.

complement Intuitively, a good that is used in combination with another good: bread is a complement to butter. (Compare with **substitute**)

- consumer surplus** The gains from trade accruing to the buyer. (Compare with **producer surplus**)
- decision tree** A method of visualizing how individuals make decisions. The branches coming out of each node represent the range of possible choices for the individual at that point. (Compare with **game tree**)
- demand curve** A curve relating the **market price** to desired purchase amounts of some good. An *individual demand curve* relates how much of that good some individual wants to buy at any given price; a *market demand curve* relates how much of that good all of the buyers together want to buy at any given price.
- descending price auction** See Chapter 9 for descriptions of different kinds of auctions.
- dominant strategy** See **strictly dominant strategy** and **weakly dominant strategy**.
- dominated strategy** See **strictly dominated strategy** and **weakly dominated strategy**.
- duopoly** A market with only two sellers. (Compare with **monopoly** and **oligopoly**)
- Dutch auction** See Chapter 9 for descriptions of different kinds of auctions.
- economics** The study of the actions and interactions of optimizing individuals. See also **microeconomics** and **macroeconomics**
- efficient** See **Pareto efficient**
- elastic** Indicating an elasticity bigger than +1 or less than -1; used to describe a relatively large degree of responsiveness, e.g., to price changes. (Compare with **unit elastic** and **inelastic**)
- elasticity** A measure of responsiveness, as in the *price elasticity of demand*, which measures the percentage change in quantity demanded resulting from a one percent increase in price. In general, the *x elasticity of y* measures the percentage change in *y* resulting from a one percent increase in *x*.
- English auction** See Chapter 9 for descriptions of different kinds of auctions.
- expected value** A probabilistic measure of the *average* outcome of a situation involving uncertainty.
- experimental economics** A branch of microeconomics that uses real-life experiments to test the predictions of economic theory, especially game theory.

fair bet A bet with an **expected value** of zero.

first price auction See Chapter 9 for descriptions of different kinds of auctions.

First Welfare Theorem See **welfare theorems**.

fiscal policy Government activities dealing with taxing and spending; in macroeconomics, compare with *monetary policy*, which refers to government activity dealing with interest rates and the money market.

free-rider An individual who makes an individually optimal decision that is detrimental to some larger group; for example, a student who does no work in a group project.

future value An economic measure of the value at some point in the future of resources in existence today. (Compare with **present value**)

game theory The branch of microeconomics dealing with strategic interactions between a small number of individuals, as in bargaining or auctions. (Compare with **price theory**)

game tree A method of visualizing sequential move games between individuals. The branches coming out of each node represent the range of possible choices for the relevant player at that point. Note: A one-player game tree is called a **decision tree**.

income effect Together with **substitution effect**, the components of how individuals respond to price changes. The income effect focuses on the effect of price changes on *income* or *purchasing power* : a price decrease effectively increases one's income (making it possible to buy more of everything), while a price increase effectively reduces one's income. Note that the income effect moves in opposite directions for **normal goods** and **inferior goods**.

inefficient See **Pareto inefficient**

inelastic Indicating an elasticity between -1 and $+1$; used to describe a relatively small degree of responsiveness, e.g., to price changes. (Compare with **unit elastic** and **elastic**)

inferior good A good (such as Ramen noodles) that you buy less of as your income increases. (Compare with **normal good**)

inflation A general increase in prices over time.

input One good that is used to make another good, as with grapes and labor in wine-making.

iterated dominance A solution concept in game theory that involves the successive elimination of **strictly dominated strategies**.

- lump sum payment** A one-time payment, usually of cash. (Compare, e.g., with **annuity**)
- macroeconomics** The branch of economics that studies national and international issues: Gross National Product (GNP), growth, unemployment, etc. (Compare with **microeconomics**)
- marginal analysis** Comparing choices with nearby choices, i.e., by making marginal changes. This is a powerful mathematical technique for identifying optimal choices: if a given choice is really optimal, making marginal changes cannot bring improvement.
- marginal x of y** With some exceptions (such as the **marginal rate of x**), the marginal x of y is the extra amount of x resulting from one more unit of y . Examples include the **marginal benefit of oranges** (the extra benefit some individual gets from one more orange), the **marginal cost of labor** (the extra cost of hiring one more worker), and the **marginal product of capital** (the extra amount of product—i.e., output—from one more unit of capital).
- marginal revenue** The extra amount of revenue a firm receives from selling one more unit of some good.
- market-clearing price** See **market price**.
- market equilibrium** A solution concept for **competitive markets** featuring the **market price** and the quantity that buyers want to buy at that price; called an equilibrium because that quantity is the same quantity that sellers want to sell at that price.
- market price** The price of some good in a **competitive market**. Also called the **market-clearing price** because it is the price that clears the market, i.e., because the amount that sellers want to sell at that price is equal to the amount that buyers want to buy at that price.
- Maximum Sustainable Yield** A resource management policy designed to yield the maximum possible harvest (e.g., of fish or trees) that can be sustained indefinitely year after year.
- microeconomics** The branch of economics that studies individual markets, supply and demand, the impact of taxes, strategic interactions, monopoly, etc. (Compare with **macroeconomics**)
- monopoly** A market with only one seller. (Compare with **duopoly** and **oligopoly**)
- monopsony** A market with only one buyer.

Nash equilibrium An important solution concept in game theory that is related to **iterated strict dominance** and **backward induction**. A Nash equilibrium occurs when the strategies of the various players are best responses to each other. Equivalently but in other words: given the strategies of the other players, each player is acting optimally. Equivalently again: No player can gain by deviating alone, i.e., by changing his or her strategy single-handedly.

nominal interest rate The interest rate in terms of cash; for example, the interest rate at a bank is a nominal interest rate. Nominal interest rates *do not* account for **inflation**. (Compare with **real interest rate**)

normal good A good (such as fancy restaurant meals) that you buy more of as your income increases. (Compare with **inferior good**)

oligopoly A market with only a few sellers; a **duopoly** is a specific example. (Compare with **monopoly**)

open-access resource A resource (such as an ocean fishery) that is open to everyone; a commons.

Pareto efficient An allocation of resources such that it *is not* possible to make any individual better off without making someone else worse off. If allocation A is Pareto efficient, there is no other allocation that is a **Pareto improvement** over A.

Pareto improvement A method of comparing different resource allocations. Allocation B is a Pareto improvement over allocation A if nobody is worse off with B than they were with A and at least one person is better off.

Pareto inefficient An allocation of resources such that it *is* possible to make any individual better off without making someone else worse off. If allocation A is Pareto inefficient, there is at least one allocation B that is a **Pareto improvement** over A.

payoff matrix A grid used to describe simultaneous move games.

per-unit tax A tax based on quantity, as in a per-gallon tax on gasoline. (Compare with **ad valorem tax**)

perfectly competitive market See **competitive market**.

perpetuity A stream of annual payments for an infinite number of years. (Compare with **annuity**)

present value An economic measure of the value today (i.e., at present) of resources in the past or in the future. (Compare with **future value**)

price discrimination The practice of charging different customers different prices for the same good; compare with **uniform pricing**. The types of price discrimination are described in Section 7.3.

price elasticity See **elasticity**.

price-taker An individual, for example a buyer or seller in a **competitive market**, who takes the market price as given because the scope of his activities are too small to affect big-picture variables.

price theory The branch of microeconomics dealing with market interactions between a large number of individuals. (Compare with **game theory**)

prisoners' dilemma In general, a prisoners' dilemma is a **simultaneous move game** with two or more players that features (1) symmetric strategies and payoffs for all the players; (2) **strictly dominant strategies** for all the players; and (3) a predicted outcome that is Pareto inefficient. (Such problems are also called **collective action problems**.) The phrase "prisoners' dilemma" may also refer to a specific game, one featuring two prisoners who must choose whether or not to accept a plea-bargain that involves testifying against the other prisoner.

producer surplus The gains from trade accruing to the seller. (Compare with **consumer surplus**)

profit Loosely defined, money in minus money out. Slightly more accurate is to account for the changing value of money over time by defining profit as the **present value** of money in minus money out. Other definitions consider this to be *accounting profit* rather than *economic profit*, which subtracts out *opportunity costs*.

quantum *Physics.* A minimum amount of a physical quantity which can exist and by multiples of which changes in the quantity occur (Oxford English Dictionary). The "quantum" of economics is the optimizing individual.

real interest rate The interest rate in terms of *purchasing power*, i.e., in terms of your ability to buy stuff. Real interest rates account for **inflation**. (Compare with **nominal interest rate**)

rent-seeking Behavior (usually by a monopolist) intended to increase profits.

repeated game A game in which a **stage game** is played multiple times.

Revenue Equivalence Theorem A theoretical result indicating that a wide variety of auctions all generate the same **expected** revenue. (Note that this only hold under certain conditions.)

risk A **risk-averse** individual prefers to avoid risks, a **risk-loving** individual prefers to take risks, and a **risk-neutral** individual is indifferent to risk.

risk premium An added payment made to avoid risk, or accepted to take on risk.

Robinson Crusoe model A simple economic model based on Daniel Defoe's 1719 novel about life on a desert island for a shipwrecked sailor.

sales tax A tax—usually an **ad valorem tax**—on the sale of some good.

second price auction See Chapter 9 for descriptions of different kinds of auctions.

Second Welfare Theorem See **welfare theorems**.

sequential move game A game, such as chess, in which players take turns moving. These games can be analyzed with **game trees**. (Compare with **simultaneous move game**)

shading one's bid An auction strategy involving lowering one's bid.

simultaneous move game A game, such as Rock, Paper, Scissors, in which players move simultaneously; one important example is the **Prisoners' Dilemma**. These games can be analyzed with **payoff matrices**. (Compare with **sequential move game**)

stage game A game, usually a simple one, that is played multiple times to yield a **repeated game**.

strictly dominant strategy A strategy for one player that yields a payoff that is strictly greater than her payoff from any other strategy, regardless of the other players' strategies. (Compare with **weakly dominant strategy**)

strictly dominated strategy A strategy for one player that yields a payoff that is strictly less than his payoff from some other strategy, regardless of the other players' strategies. (Compare with **weakly dominated strategy**)

subgame Generally used in reference to a **game tree**, a subgame is a subset of a game, e.g., one section of a game tree.

substitute Intuitively, a good that can be used to replace another good: margarine is a substitute for butter. (Compare with **complement**)

substitution effect Together with **income effect**, the components of how individuals respond to price changes. The substitution effect focuses on the effect of price changes on *relative prices*: since a price increase in some good makes it more expensive relative to similar goods, individuals will respond by *substituting out of* the now-more-expensive good and *substituting into* the related goods.

sunk cost A cost that will be borne regardless; for example, when deciding whether or not to sell some stock, the amount you spend on buying the stock is a sunk cost. As in this example, sunk costs often are costs incurred in the past.

supply curve A curve relating the **market price** to desired sales amounts of some good. An **individual supply curve** relates how much of that good some individual wants to sell at any given price; a **market supply curve** relates how much of that good all of the sellers together want to sell at any given price.

tax See **per-unit tax** or **ad valorem tax**.

trigger strategy A strategy in **repeated games** that involves cooperating as long as the other players also cooperate, but ending cooperation as soon as any of the other players stop cooperating.

unfair bet A bet with an expected value less than zero.

uniform pricing The practice of charging all customers the same price for the same good; compare with **price discrimination**.

unit elastic Indicating an elasticity of exactly $+1$ or -1 ; used to describe a degree of responsiveness, e.g., to price changes, that is neither relatively large nor relatively small. (Compare with **elastic** and **inelastic**)

weakly dominant strategy A strategy for one player that yields a payoff that is greater than *or equal to* his payoff from any other strategy, regardless of the other players' strategies. (Compare with **strictly dominant strategy**)

weakly dominated strategy A strategy for one player that yields a payoff that is less than *or equal to* her payoff from some other strategy, regardless of the other players' strategies. (Compare with **strictly dominated strategy**)

welfare economics The branch of microeconomics dealing with social welfare.

welfare theorems Two important results from **welfare economics**. The First Welfare Theorem says that complete and competitive markets yield a Pareto efficient outcome. The Second Welfare Theorem says that *any Pareto efficient* outcome can be reached via complete and competitive markets, provided that trade is preceded by an appropriate reallocation of resources.

winner's curse Used in reference to auctions, e.g., of oil-drilling rights, in which the winning bidder overbids. Note that this pertains only to *common-value* auctions.

Index

- absolute advantage, 72
- annuity, 25, 33
- arbitrage, 55–58
 - equity premium puzzle, 58
 - explorers and sheep, 56, 57
 - financial, 57–58
 - no arbitrage result, 56
 - rush-hour, 56–57
- auction, 50, 66, 91–98
 - all-pay, 96
 - code bidding, 97
 - common values, 97
 - eBay, 95
 - equivalences, 94–96
 - FCC spectrum, 97
 - independent values, 97
 - multi-unit, 97
 - Revenue Equivalence Theorem, 96
 - shading one's bid, 93
 - types of, 92–93
 - winner's curse, 97
- backward induction, 174
- Bank
 - of America, 27, 28, 58
 - of Fish, 27, 28
 - of Gold, 28
 - of Microsoft, 28
 - of Oil, 28
- barriers to entry, 107
- Bazerman, M.H., 97
- Binmore, Ken, 66
- blue
 - and yellow makes green, 113
- Brams, Steven, 69, 179
- budget constraint, 30
 - and present value, 30
- business
 - game theory applied to, 66, 177
- cake-cutting algorithms, 65–71
 - Adjusted Winner, 69
 - books about, 68, 69
 - Divide and Choose, 67
 - divide or choose?, 68
 - Moving Knife, 68
 - objectives, 69
 - three player envy-free, 69
- calculus, 145, 147
- capital
 - fish are, 27
 - theory, 27–28
- carbon tax, 143
- Cassidy, Butch, 181
- certainty equivalent, 47
- climate change, 49, 143
- Coase Theorem, 70
- collective action problem, 85
- comparative advantage, 73
- competitive markets, 106
- complementary good, 116
- compound interest, 33
 - continuous, 31
- continuous compounding, 31
- Conway, John, 69
- cost
 - sunk, 17, 20, 184
- Crampton, Peter, 97
- Crowe, Russell, 194
- decision tree, 17, 19, 23, 43, 173, 204
- demand, 111, 203
- demand curve
 - in terms of margins, 139

- income effect, 206
- increases and decreases, 113
- individual, 204
- market, 112, 203
- substitution effect, 206
- differentiation, *see* derivative, calculus
- diversification, 44
- Dixit, Avinash K., 179
- dominance
 - iterated, 191
 - strict, 190
 - weak, 192
- duopoly, 83
- eBay, 91, 95
- economic models
 - Robinson Crusoe, 72
- economics
 - basic assumption of, 13
- elasticity, 143–147
 - X elasticity of Y , 145
 - and monopoly pricing, 147
 - and taxes, 147
 - and total revenue, 147
 - elastic, 143
 - formula, 145
 - inelastic demand, 144
 - price elasticity of demand, 143
 - price elasticity of supply, 145
 - short-run versus long-run, 148
 - summed up
 - in English, 144
 - in math, 144
 - in pictures, 144
 - unit, 143
- equilibrium
 - market, 113
 - Nash, 192
- equity premium puzzle, 58
- evolution
 - game theory applied to, 66
- expected value, 43–49, 56, 59, 100, 102, 103, 196
 - and indifference curve, 47–48
- fair bet, 44, 48
- federal government, *see* government
- First Welfare Theorem, 151
- fish
 - are capital, 27
 - Maximum Sustainable Yield, 27, 28
 - open-access fishery, 28
 - slow-growing, 28
- free-rider, 85
- Friedman, David D., 56
- future value, 24, 30, 32
- game
 - theory, 65
- game theory
 - applications of, 66–67
- game tree, 173
- games
 - Battle of the Sexes, 196
 - centipede, 182
 - draft, 179
 - duopoly, 83
 - ice cream pie, 179
 - infinitely repeated, 194
 - investment, 86
 - matching pennies, 196
 - pedestrian in Rome, 181
 - poker, 66
 - Prisoners' Dilemma, 84–85
 - Rock, Paper, Scissors, 83, 197
 - sequential move, 173–176
 - simultaneous move, 83–85
 - sticks, 178
 - surprise exam paradox, 182
- Gardner, Martin, 187
- global warming, 49, 143
- green
 - yellow and blue makes, 113
- Guy, Richard, 69
- Hall, Monty, 49
- hand
 - invisible, 14, 151, 152
- Harriman, E.H., 181
- income effect, 206

- indifference, 47
- indifference curve
 - and expected value, 47–48
- induction
 - backward, 174
- inferior good, 207
- inflation, 157–160
- interest
 - compound, 31, 33
 - nominal rate of, 157–160
 - real rate of, 157–160
 - rule of 72, 32
- invisible hand, 14, 151, 152
- jokes
 - \$10 bill on the ground, 55
 - fish, 13
 - José the bank robber, 91
 - three wishes, 65
- liquidity, 58
- local government, *see* government
- lump sum, 23, 24, 33, 152, 153, 162
- macroeconomics, 165
- marginal
 - analysis, 18
 - benefit, 18, 130, 137, 139–141
 - cost, 18, 137, 139, 141
- margins, 137–141
 - and demand curve, 139
 - and supply curve, 137
 - in terms of sticks and carrots, 141
- market
 - competitive, 106
 - demand curve, 112, 203
 - equilibrium, 113
 - market-clearing price, 113
 - perfectly competitive, 106
 - price, 111
 - quantity transacted, 111
 - supply curve, 112, 203
 - total expenditure, 111
 - total revenue, 111
- Marshall, Alfred, 137
- Maximum Sustainable Yield, 27, 28
- mechanism design, 67
- monopoly, 18, 79, 107, 147, 216
 - and elasticity, 147
 - barriers to entry, 107
 - in the long run, 106
 - regulated, 18
 - uniform pricing, 79
- Monty Hall Problem, 49
- Nalebuff, Barry J., 179
- Napster, 70
- Nasar, Sylvia, 194
- Nash equilibrium, 192
 - algorithms for finding, 193
 - mixed strategy, 196
- Nash, John, 194
- Newman, Paul, 181
- Nobel Prize, 58
 - Coase, Ronald, 70
 - King, B.B., 70
 - Nash, John, 194
- normal good, 207
- Olson, Mancur, 13
- OPEC, 85
- Pareto
 - efficient, 75–80
 - improvement, 75–80
 - inefficient, 75–80
- Pareto, Vilfrado, 75
- patent race, 97
- payoff matrix, 83
- perpetuity, 26, 28, 33, 194
- poker, 66
- present value, 24, 25, 28, 30–34, 49, 53, 157, 161, 162, 194, 195
 - and budget constraint, 30
 - of annuity, 25
 - of Louisiana Purchase, 34, 162
 - of perpetuity, 26
- price discrimination, 80
- price theory, 106
- price-taker, 106
- Prisoners' Dilemma, 84–85
- profit, 14

- accounting, 14
 - economic, 14
- quantum, 218
- Raiffa
 - Howard, 49
- Ramen noodles, 207
- Redford, Robert, 181
- repeated games
 - infinitely, 194
 - trigger strategy, 195
- risk, 14, 43, 48, 52, 58–60
 - premium, 58–60
 - equity premium puzzle, 58
 - reducing with diversification, 44
- Robertson, Jack, 68
- Robinson Crusoe model, 72
- Rock, Paper, Scissors, 83, 197
- rule of 72, 32
- Samuelson, W.F., 97
- Schwartz, Jesse A., 97
- Second Welfare Theorem, 152
- Selfridge, John, 69
- Shapiro, Carl, 80
- slope, 127
 - and elasticity, 149
 - and tax burdens, 128, 148
 - of budget constraint, 30
 - of supply/demand curve, 127, 149
- Smith
 - Adam, 151
- sports
 - game theory applied to, 66, 179
- state government, *see* government
- substitute, 116
- substitution effect, 206
- supply, 111, 203
- supply curve
 - in terms of margins, 137
 - increases and decreases, 114
 - individual, 203
 - market, 112, 203
- tax
 - carbon, 143
 - equivalence, 125–126
 - per-unit, 121–125
- taxes, 121–126
 - and elasticity, 147
 - in terms of margins, 139, 141
- Taylor, Alan, 69, 179
- tree
 - game, 173
- trigger strategy, 195
- uncertainty, 58
- unfair bet, 44
- Varian, Hal, 80
- variance, 44, 46
- von Neumann, John, 66
- war
 - game theory applied to, 66
- Webb, William, 68
- welfare economics, 151
 - First Welfare Theorem, 151
 - policy implications, 153
 - Second Welfare Theorem, 152
- winner's curse, 97
- Wolfstetter, Elmar, 91
- yellow
 - and blue makes green, 113