Quantum Microeconomics

Version 3.1

June 2004

Available with and without calculus at

http://www.smallparty.org/yoram/quantum

Email contact: yoram@smallparty.org

This work is licensed under the Creative Commons Attribution-NonCommercial License. To view a copy of this license, visit Creative Commons online at http://creativecommons.org/licenses/by-nc/2.0/ or send a letter to Creative Commons, 559 Nathan Abbott Way, Stanford, California 94305, USA.

Contents

README.TXT vii

I One: The Optimizing Individual 1

1 Decision Theory 3
- 1.1 Decision Trees 4
- 1.2 Example: Monopoly 6

2 Optimization and Risk 9
- 2.1 Reducing Risk with Diversification 11

3 Optimization over Time 17
- 3.1 Lump Sums 18
- 3.2 Annuities 18
- 3.3 Perpetuities 21
- 3.4 Capital Theory 21

4 More Optimization over Time 31
- 4.1 Nominal and Real Interest Rates 31
- 4.2 Inflation 32
- 4.3 Mathematics 34

5 Transition: Arbitrage 37
- 5.1 No Arbitrage 37
- 5.2 Rush-Hour Arbitrage 38
- 5.3 Financial Arbitrage 39

II One v. One, One v. Many 43

6 Cake-Cutting 45
- 6.1 Some Applications of Game Theory 45
- 6.2 Cake-Cutting: The Problem of Fair Division 47
- 6.3 The Importance of Trade 50

7 Economics and Social Welfare **55**
7.1 Cost-Benefit Analysis 56
7.2 Pareto 57
7.3 Examples 58

8 Sequential Move Games **63**
8.1 Backward Induction 63

9 Simultaneous Move Games **73**
9.1 Dominance 74
9.2 The Prisoner's Dilemma 74
9.3 Finitely Repeated Games 76

10 Application: Auctions **79**
10.1 Kinds of Auctions 80
10.2 Auction Equivalences 82
10.3 Auction Miscellany 84

11 Application: Marine Affairs **89**
11.1 An Economic Perspective 90
11.2 A Brief History of Government Intervention 90
11.3 ITQs to the Rescue? 93

12 Transition: Game Theory v. Price Theory **95**
12.1 Monopolies in the Long Run 96
12.2 Barriers to Entry 97

III Many v. Many **99**

13 Supply and Demand: The Basics **101**
13.1 The Story of Supply and Demand 102
13.2 *Math*: The Algebra of Markets 103
13.3 Shifts in Supply and Demand 104

14 Taxes **111**
14.1 A Per-Unit Tax Levied on the Sellers 111
14.2 A Per-Unit Tax Levied on the Buyers 112
14.3 Tax Equivalence 114
14.4 *Math*: The Algebra of Taxes 116

15 Elasticities **119**
15.1 The Price Elasticity of Demand 120
15.2 Elasticities of Supply (and Beyond) 123

16 Supply and Demand: Some Details **129**
16.1 Deconstructing Supply and Demand 129
16.2 Reconstructing Supply and Demand 130
16.3 *Math*: The Algebra of Markets 131
16.4 On the Shape of the Demand Curve 131
16.5 On the Shape of the Supply Curve 134
16.6 Comparing Supply and Demand 135

17 Margins **137**
17.1 Reinterpreting the Supply Curve 137
17.2 Total Variable Cost and Producer Surplus 139
17.3 Reinterpreting the Demand Curve 140
17.4 Total Benefit and Consumer Surplus 141
17.5 Conclusion: Carrots and Sticks 142

18 Transition: Welfare Economics **145**
18.1 From Theory to Reality 147

Appendices &etc. **149**

A Government in Practice **151**
A.1 Rules and Regulations 152

Glossary **157**

Index **164**

README.TXT

> **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. All of economics ultimately boils down to the behavior of such individuals. **Microeconomics** studies their basic actions and interactions: individual markets, supply and demand, the impact of taxes, monopoly, etc. **Macroeconomics** then lumps together these individual markets to study national and international issues: Gross National Product (GNP), growth, unemployment, etc.

In structure this book—which covers only microeconomics—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 is the only truly mandatory part of this book, in that later chapters will reference topics from this chapter. For students, this means that understanding the basics from Part I is critical. For instructors, this means that adding or removing material in later chapters should be fairly painless. In particular, the text's modular structure should provide smooth entry points for other material.

About This Book

My hope is for this book to become a successful **open source** endeavor à la the Linux operating system (though on a much smaller scale). More on the nuts and bolts of what this means in a moment, but first some philosophizing. You can find out more about open source online[1]; of particular note is Eric S.

[1] http://www.opensource.org

Raymond's essay "The Cathedral and the Bazaar" (available in bookstores and also online[2]). Two of the maxims from this essay are:

- If you treat your [users] as if they're your most valuable resource, they will respond by becoming your most valuable resource.

- ("Linus's Law") Given enough eyeballs, all bugs are shallow.

Raymond's focus was on software, but I believe that these maxims also hold true for textbooks. (In the context of textbooks, "users" are students and instructors, and "bugs" are typos, arithmetic mistakes, confusing language, substantive errors, and other shortcomings.) Regardless, Raymond's essay was one of the inspirations for this book.

Okay, back to the nuts and bolts. Legally, this book is licensed under the Attribution-NonCommercial License[3] developed by Creative Commons[4]. The details are on the websites listed above, but the basic idea is that the license allows you to use and/or modify this document for non-commercial purposes as long as you credit *Quantum Microeconomics* as the original source. Combine the legal stuff with the open source philosophy and here is what it all means...

... For students and instructors This book is freely available online[5]. (One advantage of the online edition is that all the "online" links are clickable.) Please contribute your comments, suggestions, and ideas for improvement: let me know if you find a typo or a cool website, or even if there's a section of the book that you just found confusing and/or in need of more work.[6] If you're looking for something more substantial to sink your teeth into, you can add or rewrite a section or create some new homework problems. Hopefully you will get some personal satisfaction from your contribution; instructors will hopefully offer extra credit points as well.

... For writers and publishers I am happy to share the LaTeX source code for this book if you are interested in modifying the text and/or publishing your own version. And I encourage you to submit something from your own area of expertise as a contribution to the text: the economic arguments for specialization apply nicely in the world of textbook writing, and the alternative—having one or two people write about such a broad subject—is an invitation for trouble. (For an example, see this excerpt online[7].) As mentioned above, this book is typeset in LaTeX, a free typesetting program which you can learn about online[8] and/or from many mathematicians/scientists. You can also submit material in Microsoft Word or some other non-LaTeX format.

[2] http://www.openresources.com/documents/cathedral-bazaar/
[3] http://creativecommons.org/licenses/by-nc/2.0/
[4] http://creativecommons.org/
[5] http://www.smallparty.org/yoram/quantum
[6] Another lesson from *The Cathedral and the Bazaar* is that finding bugs is often harder than fixing them.
[7] http://www.smallparty.org/yoram/humor/globalwarming.html
[8] http://www.ctan.org

Acknowledgments

The individuals listed below commented on and/or contributed to this text. Thanks to their work, the text is much improved. Please note that listings (which are alphabetical) do not imply endorsement.

Yoram Bauman
yoram@smallparty.org
Seattle, Washington
June 2004

Melissa Andersen	Aaron Finkle	Noelwah Netusil
Kea Asato	Jason Gasper	Anh Nyugen
Heather Brandon	Kevin Grant	Elizabeth Petras
Gardner Brown	Robert Halvorsen	Karl Seeley
Katie Chamberlin	Richard Hartman	Amy Seward
Rebecca Charon	Andy Herndon	Pete Stauffer
Ruth Christiansen	Rus Higley	Linda Sturgis
Stacy Fawell	Wen Mei Lin	Brie van Cleve
Julie Fields	Heather Ludemann	Jay Watson

Part I

One: The Optimizing Individual

Chapter 1

Decision Theory

> Two Zen 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?"

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, study groups, etc.) is to study the incentives and behaviors of each individual in that group. Indeed, one of the key questions in economics—arguably *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 9. For an optimistic view, see Chapter 18 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: profit-maximizing firms try to make as much money as possible, utility-maximizing individuals try to get as much **utility** (i.e., happiness) as possible, etc.

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 include the happiness of friends or family, or even of people they haven't met in distant parts of the world. 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, some people act in certain ways and other people act in other ways. In all cases, we assume that each individual is making the decision that is in his or her best interest.

Of course, it is sometimes useful to make additional assumptions about the objectives of various individuals. For example, economists usually assume that *firms maximize profit.* (For our purposes, **profit** is simply money in minus money out, also known as **cash flow**.[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 goals other than profit maximizing.

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, each with his or her own objectives—that is unlikely to function seamlessly as a cohesive optimizing unit. (Recent scandals involving Enron and other companies provide ample evidence of this fact.) 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.

1.1 Decision Trees

We can imagine a simple process that optimizing individuals follow when making decisions: list all the options, then choose the best one. We can visualize this

[1] The ideas in Chapter 3 can be used to refine this definition to account for the changing value of money over time. Also note that there are a surprising number of different definitions of profit, including "accounting profit" and "economic profit".

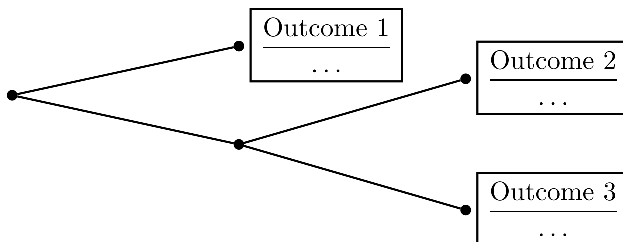

Figure 1.1: A simple decision tree

process with the help of a **decision tree**. (See Figure 1.1.) The individual starts at the left-most node on the tree, chooses between the various options, and continues moving along the branches until reaching one of the outcome boxes at the end of the tree.

Comparing the items in the various outcome boxes, we can identify two sets: those which are in all the boxes and those which are not. Items that are in all the 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 outcome 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, but the decision-making process depends crucially on items which are not in all the boxes. (These are sometimes called **opportunity costs**.)

Once you think you've found the best choice, a good way to check your work is to look at some of the nearby choices. Such **marginal analysis** basically involves asking "Are you sure you don't want one more?" or "Are you sure you don't want one 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 greater than 25 cents. Another nearby choice is "Buy six"; since our optimizing individual chose "Buy five" instead, her marginal benefit from the 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 *necessary condition*: 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 actually quite useful for checking your work. (It is not, however, infallible: although all mountaintops pass the marginal analysis test, not all of them rank as the highest spot on Earth.)

1.2 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 the freedom to engage in strategic manipulation in order to maximize their profits. A key component of that freedom is the ability to set whatever price they want for their product.[2] 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 the case of a monopolist who must charge all customers the same price. In this situation, the monopolist's profit is calculated according to

$$\begin{aligned} \text{Profit} &= \text{Total Revenue} - \text{Total Costs} \\ &= \text{Price} \cdot \text{Quantity} - \text{Total Costs} \end{aligned}$$

From this equation we can see two key ideas. First, 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.[3]

Second, we can catch a glimpse of the monopolist's fundamental problem: 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 a lot of units) The decision tree in Figure 1.2 shows some of the different price options for a hypothetical monopolist; it also shows that a higher price will drive away customers, thereby reducing the quantity that it sells. (This inverse relationship between price and the quantity that customers want to buy is formally known as the **Law of Demand**.) The monopolist 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. The decision tree shows that the monopolist's optimal choice in the example is to choose a price of $5 per unit.

[2] For this reason, monopolists are also known as **price-setting** firms, whereas firms in a competitive market are known as **price-taking** firms.

[3] More progress can be made in this direction by studying **government-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.

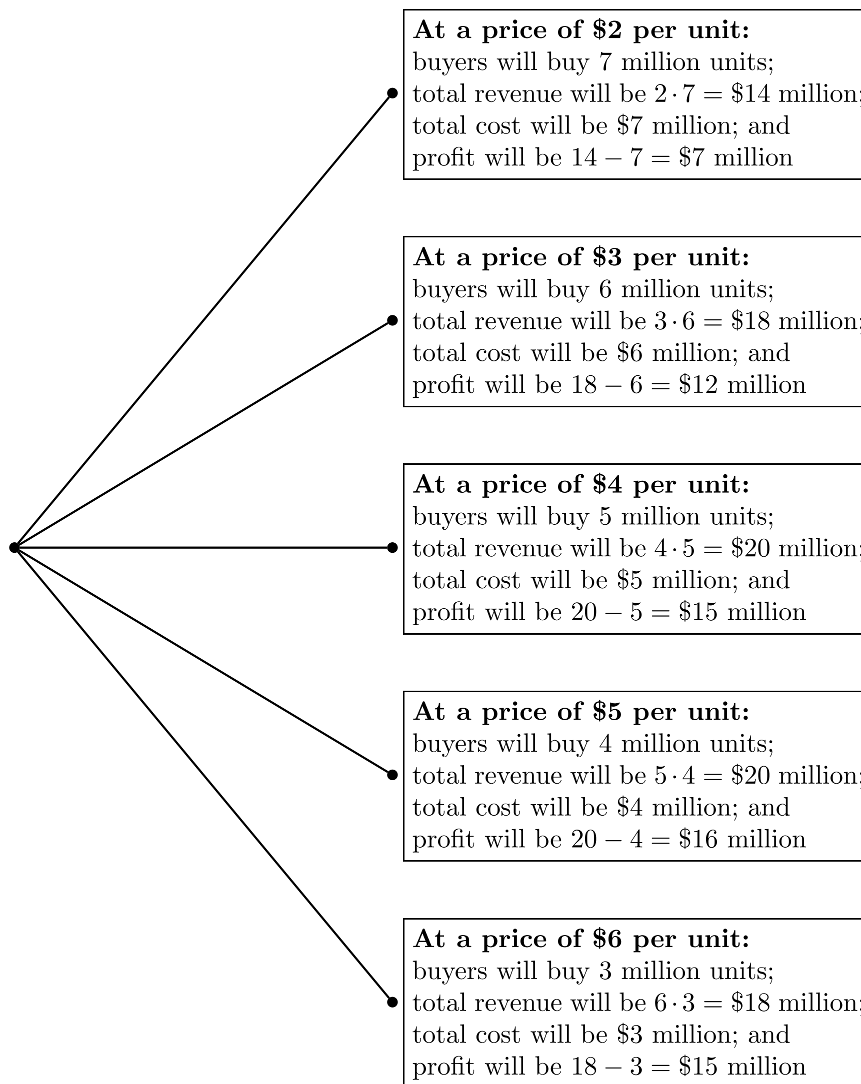

Figure 1.2: A hypothetical decision tree for a monopolist trying to find the profit-maximizing price. The second line for each pricing option shows the amount that buyers are willing to buy at the given price; the inverse relationship between the price and the amount that buyers want to buy is known as the **Law of Demand**. The third line show 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; this hypothetical situation features constant production costs of $1 per unit. The final line shows the monopolist's profit, in this case identifying a price of $5 per unit as the profit-maximizing choice.

Problems

1. A newspaper column in the summer of 2000 complained about the overwhelming number of hours being devoted to the Olympics by NBC and its affiliated cable channels. The columnist argued that NBC had established 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. 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?

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 state legislature decides to get tough on crime, and passes 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?

4. A pharmaceutical company comes out with a new pill that prevents baldness. When asked why the drug costs so much, the company spokesman replies that the company needs to recoup the $10 billion it spent on research and development.

   (a) Do you believe the spokesman's explanation?

   (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 $10 billion?

   If you said "No" above: What alternative explanation might help explain why the drug company charges so much for its pill?

Chapter 2

Optimization and Risk

Motivating question: What is the maximum amount $$x$ you would pay to play a game in which you flip a coin and get $10 if it comes up heads and $0 otherwise? (See Figure 2.1.)

The important issue in this game is your attitude toward **risk**. People who are **risk-averse** buckle their seatbelts and drive safely and 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 crazy, 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 (or other types of risk games like the lottery or Las Vegas). 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 2.1 fleshes out this idea.) The concept of expected value provides a valuable perspective by condensing this

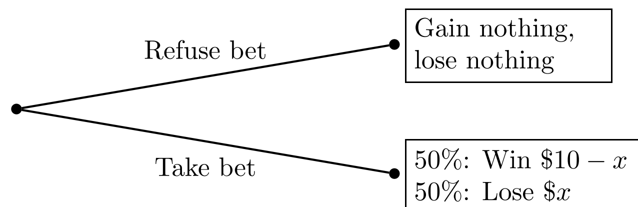

Figure 2.1: A decision tree involving risk

risky situation into a single number. Note, however, that it conveys only one perspective. In describing the average outcome, expected value fails to capture other aspects, such as the variability of the outcome.[1]

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. To calculate expected value, sum probability times value over all possible outcomes:

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

The Greek letter $\sum$ ("sigma") is the mathematical notation for summation, e.g., $\sum_{y=1,2,3} y^2 = 1^2+2^2+3^2 = 14$. In the game described above, the two possible outcomes are heads (H) and tails (T), so the expected value is

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

If it costs $\$x$ to play the game, the overall expected value becomes $\$(5 - x)$.

Example: Fair and Unfair Bets

A **fair bet** is a bet with an expected value of zero. Flipping a coin and having me pay you $\$x$ if it comes up heads and you pay me $\$x$ if it comes up tails is a fair bet (provided, of course, that the coin is not weighted). 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 a 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, you get back \$36 (the dollar you bet plus 35 more); if you guess the wrong number, you get back \$0. So if you bet \$1 on, say, number 8, then the expected value of the amount you'll get back is

$$\Pr(8) \cdot (\$36) + \Pr(\text{Not } 8) \cdot (\$0) = \frac{1}{38} \cdot (\$36) + \frac{37}{38} \cdot (\$0) \approx \$0.95.$$

Since it costs \$1 to play and your expected return is only \$0.95, the overall expected value of betting \$1 in roulette is $-\$0.05$. So roulette, like other casino games, is an unfair bet.

[1]The concept of **variance** attempts to convey this part of the story.

2.1 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 fair coin, one where the probability of getting heads is 50%. 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, your 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 a statistical theorem called the **law of large numbers**. In English, the law of large numbers says that if you flip a (fair) coin a large number of times, odds are that the proportion of heads will be close to 50%. (Details online[2] and online[3].) 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 an individual gambler will play roulette only a few times, but that the casino plays roulette thousands of times each day. The law of large numbers does not apply to the individual gambler, but it does apply to the casino. As a consequence, the 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[4] 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

[2] http://www.stat.berkeley.edu/~stark/Java/lln.htm
[3] http://www.ruf.rice.edu/~lane/stat_sim/binom_demo.html
[4] http://www.theonion.com/onion3920/casino_has_great_night.html

one or two stocks, investors can reduce risk by owning many stocks. (One way to do this is to own shares of mutual funds or index funds, companies whose line of business is investing money in the stock market.)

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). The risk in this situation is clear: either you win the whole $100 or you win nothing.

Now consider flipping two coins, each independently paying $50 for heads and $0 for tails. The four possible outcomes—all equally likely—are HH (paying $100), HT (paying $50), TH (paying $50), and TT (paying $0). Note that the expected value in this situation is $50:

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

This is the same expected value as in the one-coin flip situation described above. Compared to the one-coin flip, however, this two-coin flip is less risky: instead of always winning everything or nothing, there is now a 50% chance of winning something in between. 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 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 directly from the law of large numbers. It is important to note, however, that 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 coin flip, stocks have a tendency to go up or down together. (In statistical terms, coin flips are **independent** while the prices of different stocks 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 affect companies independently. But the law of large numbers does not apply when risks are correlated: diversification cannot reduce the risk of stock market booms or busts 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 Chapter 5

suggests 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; diversifying your portfolio reduces risk, but it may also reduce your expected payoff. The optimal behavior in such situations is the subject of **portfolio selection theory**, a branch of economics whose development helped win James Tobin the Nobel Prize in Economics in 1981. Professor Tobin died in 2002; his obituary in the *New York Times* (available online[5]) included the following 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'."

Problems

1. You roll a six-sided die and win that amount (minimum $1, maximum $6). What is the expected value of this game?

2. With probability 1/3 you win $99, with probability 2/3 you lose $33. What is the expected value of this game?

3. *Fun/Challenge* (The Monty Hall Problem) The Monty Hall 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 *does not* show you what is 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 to from Door #1 to Door #3.] What should you do?

4. 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 the question and possible answers for Problem #23 are in Hungarian. (When you ask your teacher, she claims that the class learned Hungarian on Tuesday. . . .)

[5] http://cowles.econ.yale.edu/archive/people/tobin/nyt_obit.htm

(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 by people like you. 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 the policy you came up with in the previous part 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 1/3rd 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?

5. 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.

6. Howard Raiffa's book *Decision Analysis: Introductory Lectures on Choices Under Uncertainty* (New York: McGraw-Hill, 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 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.

7. *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.)

8. 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 10.)

Chapter 3

Optimization over 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.)

Answer: It depends. Simply comparing dollars today and dollars in the future is like comparing apples and oranges. We need some way of translating values in the future (and the past) into something comparable to values in the present.

The way economists solve this problem is by observing that banks and other financial institutions turn money today into money tomorrow and vice versa. By using banks to save and borrow money, we can use the relevant **interest rates** to put future values (and past values) 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 a dollar last year has a present value of $1.10: if you had put a dollar in the bank last year, you'd have $1.10 today. Similarly, the present value of a dollar a year from now is about $.91: put $.91 in the bank today and in a year you'll have about $1.00.

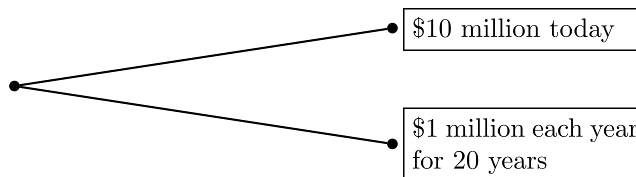

Figure 3.1: A decision tree involving time

So we can use interest rates to compare apples to apples instead of apples to oranges.

3.1 Lump Sums

Question: Say you've got \$100 in a Swiss bank account at a 5% annual interest rate. How much will be in the account after 30 years?

Answer:

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

So the **future value of a lump sum** of $\$x$ invested for n years at interest rate r is $\text{FV} = x(1+r)^n$. Although the future value formula can be useful in and of itself, it also provides insight on the topic of present value:

Question: What 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 $x = \frac{100}{(1.05)^{30}} \approx \$23.14$. If you put \$23.14 in the bank today at 5% interest, after 30 years you'd have about \$100.

So the **present value of a lump sum** payment of $\$x$ received at the end of n years at interest rate r is

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

3.2 Annuities

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%? (Unlike a lump sum payment, we now have a stream of payments. A stream of annual payments is called an **annuity**.)

Idea: Imagine that you put those payments in a savings account when you receive them. Now imagine another savings account into which you put some amount x today, and don't make any more deposits or withdrawals, simply letting the interest accumulate. Imagine that the two accounts have the same

balance after the third and final \$100 deposit in account #1. Then the amount x is the present value of the annuity.

Answer: You've got three \$100 payments. The present value of the first \$100 payment, which comes after one year, is $\frac{100}{(1.05)^1} \approx 95.24$. The present value of the second \$100 payment, which comes after two years, is $\frac{100}{(1.05)^2} \approx 90.70$. And the present value of the third \$100 is $\frac{100}{(1.05)^3} \approx 86.38$. So the present value of the three payments together is about $\$95.24 + \$90.70 + \$86.38 = \$272.32$. One way to calculate the present value of an annuity, then, is to calculate the present value of each year's payment and then add them all together. Of course, this can get pretty tedious, e.g., if you're trying to answer the lottery question at the beginning of this chapter.

Is There an Easier Way?

Answer: Yes, and it involves some pretty mathematics. Here it is: We want to calculate the present value (PV) of a 3-year \$100 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 $.05 \cdot \text{PV}$. The right hand side also simplifies (all the middle terms cancel, hence the moniker **telescoping series** for what mathematicians also call a **geometric series**), yielding

$$.05 \cdot \text{PV} = 100 - \frac{100}{(1.05)^3}.$$

Dividing both sides by .05 and grouping terms, we get

$$\text{PV} = 100 \left[\frac{1 - \frac{1}{(1.05)^3}}{.05}\right].$$

This suggests an easier way of calculating annuities: At interest rate r, the **present value of an annuity** paying \$$x$ at the end of each year for the next

n years is

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

Example: Car Payments

Car dealers usually give buyers a choice between a lump sum payment (the car's sticker price of, say, \$15,000) and a monthly payment option (say, \$400 a month for 36 months). The annuity formula can come in handy here because it allows you to compute the sticker price from the monthly payments, or vice versa. The applicability of the annuity formula stems from the fact that the monthly option essentially involves the dealer loaning you an amount of money equal to the sticker price. At the relevant interest rate, then, *the sticker price is the present value of the stream of monthly payments.*

The only real complication here is that car payments are traditionally made monthly, so we need to transform the dealer's *annual* interest rate—the APR, or Annual Percentage Rate—into a *monthly* interest rate. (Like the other formulas in this chapter, the annuity formula works not just for annual payments but also for payments made daily, weekly, monthly, etc. The only caveat is that you need to make sure that you use an interest rate r with a matching time frame, e.g., a monthly interest rate for payments made monthly.) A good approximation of the monthly interest rate comes from dividing the annual interest rate by 12, so an APR of 6% translates into a monthly interest rate of about 0.5%, i.e., $r \approx .005$.[1]

With a monthly interest rate in hand, we can use the annuity formula to translate monthly payment information (say, \$400 a month for 36 months at 0.5%) into a lump sum sticker price:

$$\text{PV} = 400 \left[\frac{1 - \dfrac{1}{(1.005)^{36}}}{.005} \right] \approx \$13,148.$$

More frequently, you'll want to transform the lump sum sticker price into a monthly payment amount, which you can do by solving the annuity formula for x, the monthly payment:

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

[1] Because of compounding, a monthly interest rate of 0.5% actually corresponds to an annual interest rate of about 6.17%. See problem 11 for details and for information about a completely accurate formula.

Using this formula, we can figure out that a \$15,000 sticker price with a 6% APR (i.e., a monthly interest rate of about 0.5%, i.e., $r \approx .005$) translates into 36 monthly payments of \$456.33, or 60 monthly payments of \$289.99.

3.3 Perpetuities

Question: What is the present value of receiving \$100 at the end of each year *forever* at a 5% interest rate?

Answer: Such a stream of payments is called a **perpetuity**,[2] and it really is forever: after death, you can give the perpetuity away in your will! So what we're looking for is

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

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

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

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

$$1.05 \cdot \text{PV} - \text{PV} = 100.$$

This simplifies to $\text{PV} = \frac{100}{.05} = \$2,000$. In general, then, the **present value of a perpetuity** paying $\$x$ at the end of each year *forever* at an interest rate of r is

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

3.4 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 own a lake with a bunch of fish in it. You want to maximize your profits from selling the fish. How should you manage this resource? What information is important in answering this question?

Partial answer: The growth rate of fish, the cost of fishing (e.g., is fishing easier with a bigger stock of fish?), the market price of fish over time.... Note that one management option, called **Maximum Sustainable Yield** (MSY), is to

[2]Note that the word comes from perpetual, just like annuity comes from annual.

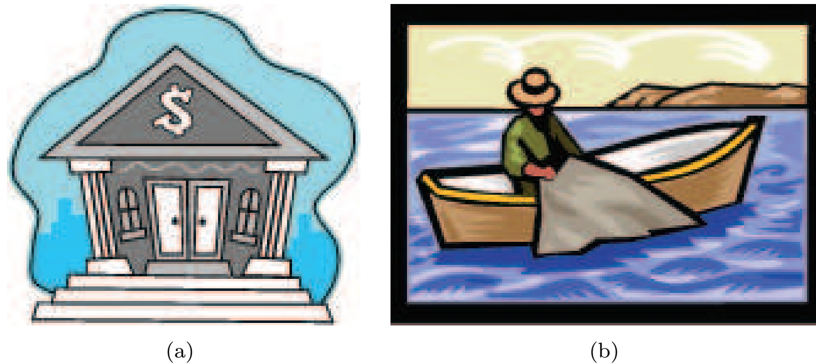

(a) (b)

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

manage the lake so as to get the maximum possible catch that you can sustain year after year forever, as shown in Figure 3.3. We will return to this option shortly.

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 different options, which are to invest in the bank (by catching and selling all the fish and putting the proceeds in the bank), or to invest in the fish (by not catching any of them, thereby allowing them to grow and reproduce to make more fish for next year), or some combination of the two (catching some of the fish and letting the rest grow and reproduce). This suggests that you need to compare the interest rates you get at the Bank of America with the interest rates you get at the Bank of Fish (see Figure 3.2).

This sounds easy, but there are lots of potential complications. To make the problem tractable, let's simplify matters by assuming 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 is 5%. Then one interesting result is that the profit-maximizing catch level is *not* the Maximum Sustainable Yield (MSY) defined above and in Figure 3.3.

To see why, let's calculate the present value of following MSY: each year you start out with 100 pounds of fish, catch 10 pounds of fish, and sell them for $1 per pound, so the perpetuity formula gives us a present value of $\dfrac{10}{.05} = \$200$.

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

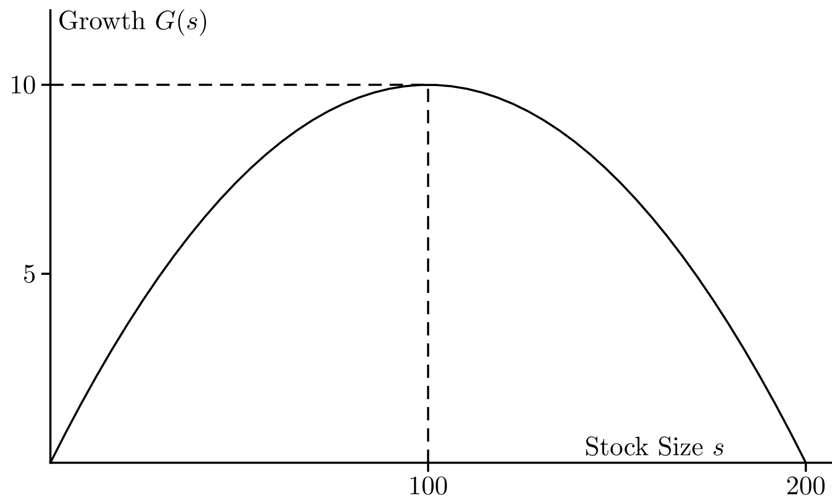

Figure 3.3: With an initial population of 100 pounds of fish, population growth over the course of one year amounts to 10 additional pounds of fish. Harvesting 10 pounds returns the population to 100, at which point the process can begin again. A population of 100 pounds of fish therefore produces a **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 after year.

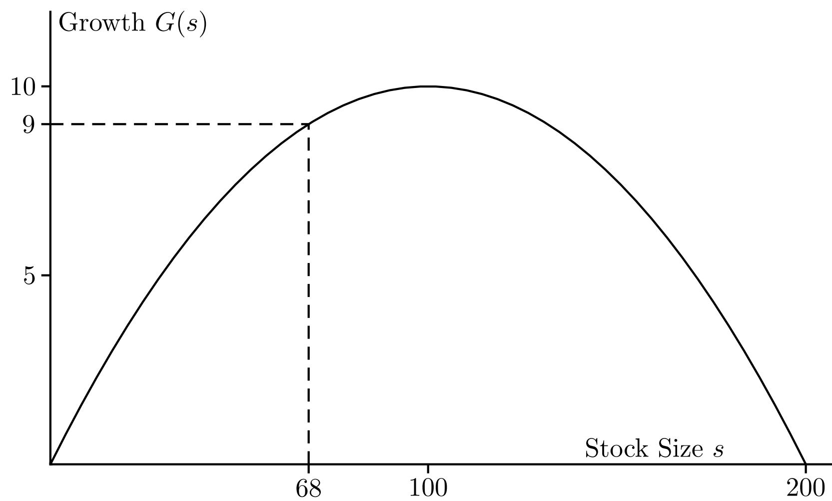

Figure 3.4: With an initial population of about 68 pounds of fish, population growth over the course of one year amounts to 9 additional pounds of fish, meaning that 9 pounds per year is the sustainable yield corresponding to an initial population of 68 pounds of fish.

year (as shown in Figure 3.4), so every year thereafter you can catch 9 pounds of fish. The present value of this harvest policy is $32 + \frac{9}{.05} = \$212$, which is more than the present value from the MSY policy!

So MSY is *not* the optimal 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.[3]

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![4]

So why do people invest in these things? 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" or "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.

Problems

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.

[3] Another part of the explanation 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.

[4] Puzzle: What about stock splits?

(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?

2. Intuitively, how much difference do you think there is between an annuity paying \$100 each year for 1 million 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. Explain the perpetuity formula in terms of "living off the interest".

4. 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!

5. (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, 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?

6. 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?

7. *Fun* Compound interest has been called the Eighth Wonder of the World. Here's why.

   (a) Legend has it that the inventor of chess showed his game to the King of India (or maybe it was the Emperor of China), and the king was so impressed he offered the inventor a reward. The inventor said that all he wanted was 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. How many grains of rice was the king 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 and then calculate this height.

   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 take 70 years to double to $200, and every 70 years thereafter it will double again (to $400, $800,...). So if we call 70 years a *lifetime*, we see that an *annual* interest rate of 1% is equivalent to a *lifetime* interest rate of 100%. So 1% growth and 100% growth are different in degree but not in spirit.

8. *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 $1 trillion ($= 10^{12} = 1,000,000,000,000$) 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%? [Note: If all the zeroes confuse you or your calculator, use $1,000,000 or $1,000 instead.]

(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 8a 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 8b and compute the present value of having that amount of damage occur 100 years in the future. [Note: If you didn't answer part 8b, 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?

9. The U.S. Federal Reserve Board (a.k.a. the Fed) recently took steps to lower interest rates in the U.S.: in January 2001 the Federal Funds rate was 6.0%, and by June it was 4.0%. You'll learn a whole bunch more about the Fed in Econ 201 (and can get details about changes in the Federal Funds rate at http://www.federalreserve.gov/fomc/fundsrate.htm), but we can already understand something about Alan Greenspan and his mysterious ways.

   (a) Pretend (you might not have to) that you're a consumer. You're thinking about borrowing some money to buy a new TV set (i.e., you're thinking about buying on credit). Since you can now get a loan at a lower interest rate, does the Fed's action make it more likely or less likely that you'll decide to buy that TV set?

   (b) There are a lot of consumers out there, all thinking rationally, just like you. Is the Fed's action likely to increase or decrease the total amount of stuff that we buy?

   (c) Pretend that you're running a business. You have made a tidy profit of $100, and you're trying to decide whether to put that $100 in the bank or spend it on expanding your business. Does the Fed's action make spending that money on expanding your business more attractive or less attractive?

   (d) There are a lot of CEOs out there, all thinking rationally, just like you. So will the Fed's action tend to increase or decrease business expansions and economic activity?

   (e) Pretend that you're an investment guru. Somebody's unsuspecting grandmother has given you $100 to invest on her behalf, and you're trying to decide whether to invest that money in the bank or buy stock. Does the Fed's action make investing in the stock market more attractive or less attractive?

(f) There are a lot of investment gurus out there, all thinking rationally, just like you. So is the Fed's action likely to make the stock market go up or down?

10. *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 almost 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: it has been about 200 years since the Louisiana Purchase. 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%?

(b) 8%? (See problem 5 on page 36 for more about the Louisiana Purchase.)

11. 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%. How much money will actually be in the bank at the end of one year (12 months)? Use your answer to determine the actual annual interest rate that is 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%.

(d) *Challenge* Use logarithms to determine the actual monthly interest rate that is equivalent to an annual rate of 0.5%.

12. *Fun* (Thanks to Kea Asato.) The 2001 Publishers Clearing House "$10 million sweepstakes" came with three payment options:

    Yearly Receive $500,000 immediately, $250,000 at the end of Year 1 and every year thereafter through Year 28, and $2.5 million at the end of Year 29 (for a total of $10 million).

    Weekly Receive $5,000 at the end of each week for 2,000 weeks (for a total of $10 million).

    Lump sum Receive $3.5 million immediately.

    Calculate the present value of these payment options if the interest rate is 6% per year. What is the best option in this case? (See problem 6 on page 36 for more about this problem.)

Chapter 4

More Optimization over Time

Present value calculations are complicated by the existence of **inflation**, a general increase in prices over time. The presence of inflation requires us to distinguish between *nominal* interest rates and *real* interest rates.

4.1 Nominal and Real Interest Rates

The **nominal interest rate** is the interest rate the bank pays. If you have $100 and you put it in a bank account paying 5% interest, in one year you'll have 5% more money ($105). 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% higher in one year, your ability to buy stuff (formally called your **purchasing power**) 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, and in doing so it properly accounts for inflation. 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, 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 (i.e., the real interest rate), 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%.

4.2 Inflation

As noted at the beginning of this chapter, inflation is a general increase in prices over time. To measure inflation, economists look at how prices change for a "representative market basket" of goods and services: food, rent, 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%.)

The most commonly used measure of inflation in the United States is the Consumer Price Index (CPI), which is shown in Figure 4.1a. According to the CPI, a market basket of goods and services that cost $100 in 2000 would have cost about $10 in 1925; in other words, the purchasing power of $100 in 2000 was about the same as that of $10 in 1925. Taking inflation into account can therefore produce significant shifts in perspective.

As an example, consider Figures 4.1b and 4.1c, both of which show average U.S. gas prices between 1918 and 2002. Figure 4.1b features *nominal* prices: a gallon of gasoline actually sold for an average of $0.22 in 1925 and $1.56 in 2000. This perspective shows gas prices increasing over time, with the late 1970s marking a striking transition between "low" and "high" gas prices.

The perspective in Figure 4.1c is quite different. It shows gas prices *falling* over time, interrupted—but only temporarily—by price shocks in the 1930s and the 1970s.[2] These conclusions come from analyzing *real* gas prices: Figure 4.1c adjusts for inflation by putting everything in **constant year 2000 dollars**, i.e., in terms of year 2000 purchasing power.

For a more explicit comparison between Figures 4.1b and 4.1c, consider the price of a gallon of gasoline in 1925. Figure 4.1b asks this question: "What was the average price of gasoline in 1925?" The answer: about $0.22 per gallon. In contrast, the question posed in Figure 4.1c is this: "What would the average price of gasoline have been in 1925 *if the general price level in that year had been the same as the general price level in the year 2000*?" The answer: about $2.20 per gallon because, according to the CPI, the general price level in the year 2000 was about 10 times higher than in 1925.

Because we have a good grasp of the purchasing power of today's money (i.e., because we "know what a dollar [today] is worth"), 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 4.1d shows gas prices in constant 1925 dollars, i.e., in terms of 1925 purchasing power. Note that the graph is identical to that in Figure 4.1c except for the labels on the vertical axis (1925 dollars versus year 2000 dollars).

[1] This 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 economics course to learn how to deal with complications caused by changes in the representative market basket over time.

[2] The latter price shock resulted from successful collusion by the OPEC (Organization of Petroleum Exporting Countries) cartel. The cause of the former price shock is not known to this author—anybody want to do a research project?

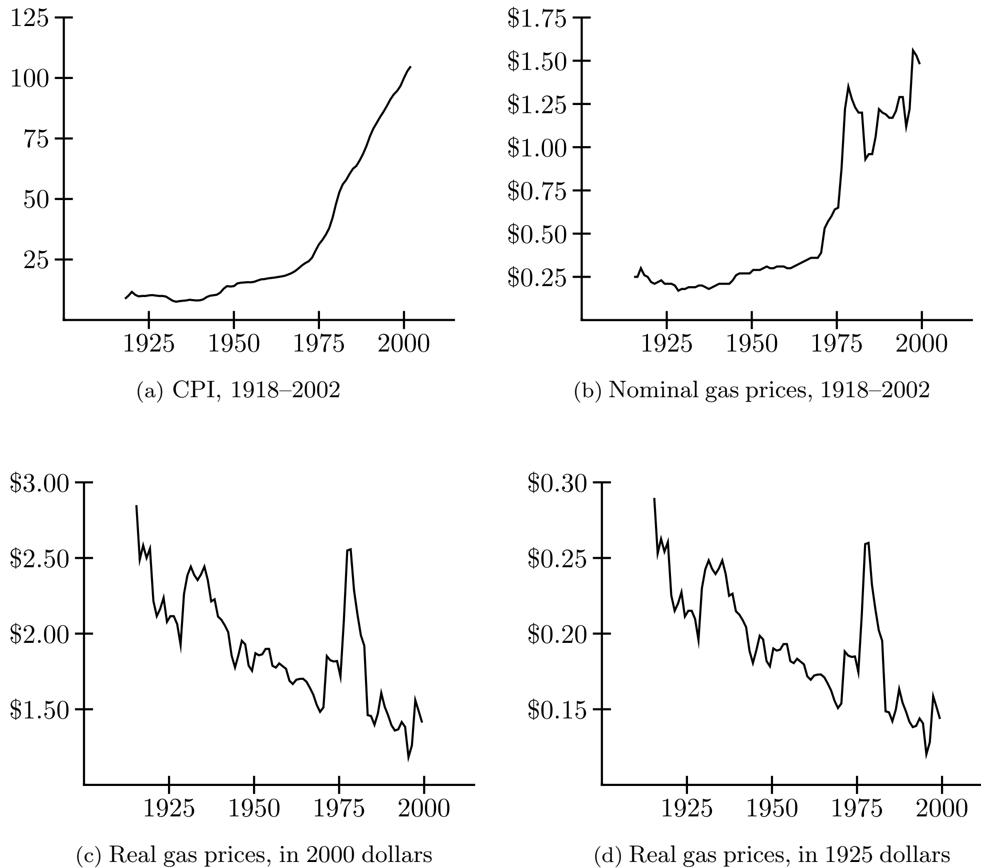

(a) CPI, 1918–2002

(b) Nominal gas prices, 1918–2002

(c) Real gas prices, in 2000 dollars

(d) Real gas prices, in 1925 dollars

Figure 4.1: Inflation and average U.S. gas prices, 1918–2002. Figure (a) shows the Consumer Price Index (CPI) measure of inflation; note that there was a period of **deflation** (i.e., 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 $0.22 in 1925 and $1.56 in 2000. Figure (c) shows average gas prices using *real* year 2000 dollars, meaning that it adjusts for inflation by putting everything in terms of year 2000 purchasing power. According to the CPI—shown in Figure (a)—a market basket that cost $100 in 2000 would have cost about $10 in 1925. It follows that the $0.22 price tag on a gallon of gasoline in 1925 is equivalent to about $2.20 in year 2000 dollars, as shown in Figure (c). Conversely, the $1.56 price tag on a gallon of gasoline in 2000 is equivalent to $0.156 in 1925 dollars; this can be seen in Figure (d), which shows average gas prices using real 1925 dollars. (Source: Figures are based on American Petroleum Institute data.)

4.3 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 4.1's discussion of Colombia. If the nominal interest rate is n and the inflation rate is i, the real interest rate r is given by

$$1 + r = \frac{1+n}{1+i}, \quad \text{i.e.,} \quad r = \frac{1+n}{1+i} - 1.$$

Intuitively, the numerator $(1 + n)$ 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.

There is also a handy rule of thumb that works well when the inflation rate is small (say, less than 10%). To derive it, note that

$$r \;=\; \frac{1+n}{1+i} - 1 \;=\; \frac{1+n}{1+i} - \frac{1+i}{1+i} \;=\; \frac{n-i}{1+i} \;\approx\; n - i$$

when the inflation rate i is small because $1 + i \approx 1$. In short,

$$r \approx n - i,$$

which in words says that the real interest rate is approximately equal to the nominal interest rate minus the inflation rate.

It is important to remember that the rule of thumb is an approximation that works well only when the inflation rate is small. For an example, return to the Colombia calculation above, where the exact formula gives a real interest rate of 3.7%:

$$\frac{1+.13}{1+.09} - 1 = .037.$$

The rule of thumb estimates the real interest rate at $13\% - 9\% = 4\%$, which is pretty close. In contrast, consider a nominal interest rate of 113% and an inflation rate of 109%. The rule of thumb again estimates the real interest rate at 4% ($113\% - 109\% = 4\%$), but this time the real interest rate is actually only 1.9%:

$$\frac{1+1.13}{1+1.09} - 1 = .019.$$

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. 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. But there are a couple of rules of thumb you can follow; two of them are described in Problem 4.

A final note here is that many microeconomic analyses (including many of the examples and descriptions in this text) simplify matters by assuming that the inflation rate is zero. In this case, real and nominal interest rates are equal to each other and it makes sense to refer to "the interest rate" without specifying one or the other.

Problems

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.

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.)

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.

   (a) You put \$100 in the bank today. How much will be in your account after 10 years?

   (b) 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?

   (c) 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.")

   (d) 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.

   (e) 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* (3d). Compare your answer here with the result from question 3c.

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%?

5. *Fun* Recall the question from Chapter 3 concerning the Louisiana Purchase. That problem (#10 on page 28) 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?

6. *Fun* Recall the question from Chapter 3 concerning the Publishers Clearing House sweepstakes. That problem (#12 on page 29) asked you to calculate the present value of different payment options. In calculating those present values, should you use the real interest rate or the nominal interest rate?

Chapter 5

Transition: Arbitrage

> 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 purpose of this transitional chapter is to connect the first part of this book, which analyzes the actions of a single optimizing individual, with later parts, which analyze interactions *between* optimizing individuals. Our focus will be on the idea of **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 for $100 and on the NASDAQ stock exchange for $105, an individual trader can engage in arbitrage by simultaneously "buying low" (purchasing Microsoft stock on the New York Stock Exchange for $100 per share) and "selling high" (selling Microsoft stock on NASDAQ for $105 per share). The net result is a riskless gain of $5 per share.

More generally, the concept of arbitrage applies to assets that are *similar* rather than *identical*. In 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

Arbitrage as a concept 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 preceding chapters in this book had a similar focus. The following chapters focus on the interactions *between* optimizing

individuals. One of the key questions will be: "What would 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, an almost completely 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 is an exercise in futility because all the lanes are moving at the same speed.

These predictions are examples of what David D. Friedman calls "hidden order".[1] 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 the different exchanges, and nobody is "in charge" of the different lanes of traffic during rush hour; to the contrary, each stockbroker is only interested in making money, and each driver is only interested in getting home as fast as possible. Despite this, the aggregate result of these myopic activities is not chaos but order. Finding such "hidden order", one of the pleasures of economics, is the focus of the rest of this book.

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 the victory will be short-lived: 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—end up equalizing the traffic speeds on the different lanes.

A curious note 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 relax rather than stress out over switching lanes all the time. 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 at the beginning of this section—it turns out that explorers *can* get ahead, but only by a little—we now turn our attention to the "expect". This word refers to the expected value calculations of Chapter 2, and as in that chapter the term is

[1] Friedman's excellent book is called *Hidden Order: The Economics of Everyday Life* (Harper Business, 1996).

used to focus 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 can be said is that *on average* you will neither gain nor lose by taking the bet.

Returning to rush hour, imagine that a second (identical) bridge is built right next to the SR-520 bridge that currently connects Seattle and the Eastside. 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 bridges will always be approximately equal: if there is a big accident on one of the bridges right after you get on it, you will get home hours later than the lucky drivers who chose the other bridge. What we can say is that the *expected* travel times on the two bridges will be approximately equal. In hindsight—i.e., after you cross the bridge—it is easy to say whether or not you should have taken the other bridge. Before the fact, though, all that can be said is that *on average* the bridge you choose will get you home in approximately the same amount of time as the other bridge.

5.3 Financial Arbitrage

A nice analogy connects financial decisions and traffic decisions: think of the different lanes of the freeway as different investment options (Microsoft stock, Intel stock, etc.), 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). Given these similarities, it is not surprising that there are similarities in the economic analyses of these two phenomena. A completely correct statement about financial arbitrage turns out to be 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 price of oil over the last 30 years, or the price of internet stocks (amazon.com, webvan, Cisco) over the last 5 years, you'll see that their prices do not steadily increase at the rate of interest. But using hindsight to look back and say that you should have bought or sold XYZ stock five years ago is like using hindsight to say 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, for example, 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 selling stock in companies that are expected to perform poorly (thereby lowering the stock

price, which makes the stock more attractive) and buying stock in companies that are expected to do well (thereby increasing the stock price, which makes the stock less attractive), these explorers equalize the attractiveness of different stocks. (We'll see shortly that equal attractiveness doesn't necessarily mean equal expected rates of return.)

As in the case of rush-hour arbitrage, the presence of financial "explorers" allows other investors (the "sheep") to feel comfortable with just picking some stocks and sticking with them. This is a good thing because many people don't want to (or, because they have day jobs, can't) spend all their time investigating 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 almost completely correct 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.... (See the joke at the beginning of this chapter for an example.)

A complete picture is that there *are* small discrepancies that occasionally appear between the price of Microsoft stock on the New York Stock Exchange and on NASDAQ; and there *are* people who profit from such discrepancies (and, by opposing, end them). 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. As for the idea that these people are making out like bandits, it is worth noting that 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**: as we saw before, people are often unwilling to take fair bets, meaning that they 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 should have higher expected interest rates than U.S. government bonds because they involve uncertainty: if the companies selling those bonds go bankrupt, the bonds will be worthless. (This idea is tricky; see Problem 2 for an example.) If one-year junk bonds and U.S. government bonds have expected returns of 6% and 4% respectively, then the **risk premium** associated with the company's 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 do 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 the following information provided in a recent issue of The Schwab Investor:[2] 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.

Problems

1. Explain (as if to a non-economist) why comparable investments should have comparable expected rates of return.

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?

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.

[2]It would be better to have an unbiased source for this information. . .

Part II

One v. One, One v. Many

Chapter 6

Cake-Cutting

> A recently divorced woman finds a beat-up old bottle while walking 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 to be rich." "Remember," says the Genie, "your ex-husband will have ten times more money than you." "That's fine," the woman replies, "I'll still be able to buy whatever I want." "Alright, whatever you say," 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 optimizing individuals *interact*. The branch of economics that studies strategic interactions between individuals is called **game theory**.

Some key issues to consider when analyzing or describing a strategic interaction are the **players** (Who's involved?); the **strategies** (What options does each player have?); the **payoffs** (What are the players' outcomes from the various strategy combinations?); the **timeline** (Who gets to act when?); and the **information structure** (Who knows what when?).

6.1 Some Applications of Game Theory

One application of game theory is, not surprisingly, to games. Here, for instance, is the introduction to Ken Binmore's economic analysis of poker his book *Fun and Games: A Text on Game Theory*:

> 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.

In other words, poker is all about strategic behavior: 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. . . .

Of course, poker is not the only interesting game of strategy. Others include

Evolutionary games When facing conflict, why are some species aggressive ("fight") while others run ("flight")? Why have some species evolved to be big and strong and others to be small and quick? These are sample 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? (Chapter 10 analyzes various kinds of auctions.)

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. The ideas of game theory helped make John von Neumann, one of its pioneers, into a proponent of "preventative war", i.e., nuking the U.S.S.R. before the Soviets 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 some features of the Cold War, such as the 1972 ABM Treaty (in which the U.S. and the U.S.S.R. agreed to not pursue *defenses* against missile attacks) and the doctrine of Mutually Assured Destruction.

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 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 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 children. (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 comparable to 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 with only a little frosting, parts with lots of chocolate and parts without chocolate, etc. (The different parts of the cake correspond to the different items—cars, houses, kids, etc.—in a divorce or estate settlement.) Second, the people 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 children can divide up the cake fairly and without rancor. In other words, Mom's problem is one of **mechanism design**:

Motivating question: Let's say you have two children, and you want them to divide the cake between them so that each child gets (in his or her opinion) at least 50% of the cake. What should you do?

One answer: There's a time-honored algorithm called Divide and Choose: have one child cut the cake into two pieces, and the other child choose between them. Now each child can guarantee him- or herself at least 50% of the cake.

Another answer: Pass a long knife from left to right over the cake. When one child calls out "Stop", cut the cake at that point and give the slice to the left of the knife to the child that called out. The remainder of the cake goes to the child that was silent. (This is called a Moving Knife algorithm.)

Question: Imagine you're one of the children participating in a Divide and Choose algorithm. Do you want to be the cutter or the chooser?

Answer: Well, 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 also depends on how much you know about the other child's preferences, and how much they know about your preferences.[1]

If you know a lot about what they like, 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; odds are that your valuation is different and you'll be able to choose a piece that you value at more than 50%.

Side note: The importance of **information** will be a recurring theme in our analysis of game theory (and in fact in all of microeconomics).

Question: What can Mom do if she has three or more kids?

One answer: You can generalize the Moving Knife algorithm to deal with any number of children. Simply follow the same procedure (pass the knife over the cake from left to right, stop when one child calls out "Stop", and give the slice to the left to the child that called out), and then repeat with the remaining cake and the remaining children. Each time you complete the procedure there's less cake remaining and one less child. With n children, repeating the Moving Knife algorithm $n-1$ times will successfully divide the cake.

Another answer: You can also generalize the Divide and Choose algorithm to get what's called the Successive Pairs algorithm. It's a little complicated, but here's how it works with three children: Have two of the children play Divide and Choose, so that each of them gets at least 50% of the cake. Then have each of those children divide their part into three pieces, and have the third child choose one piece from each of those. The cake will thus have been cut into six pieces, and each child will end up with two. (Exercise: Can you convince yourself—or someone else—that each child will get at least 33% of the cake in his or her estimation?)

Which of these solutions is better? Well, that depends on your goal. Amusingly, mathematicians concern themselves with how to fairly divide the cake

[1] In fact, it also depends on how much you know they know about your preferences, and how much you know they know you know, and...

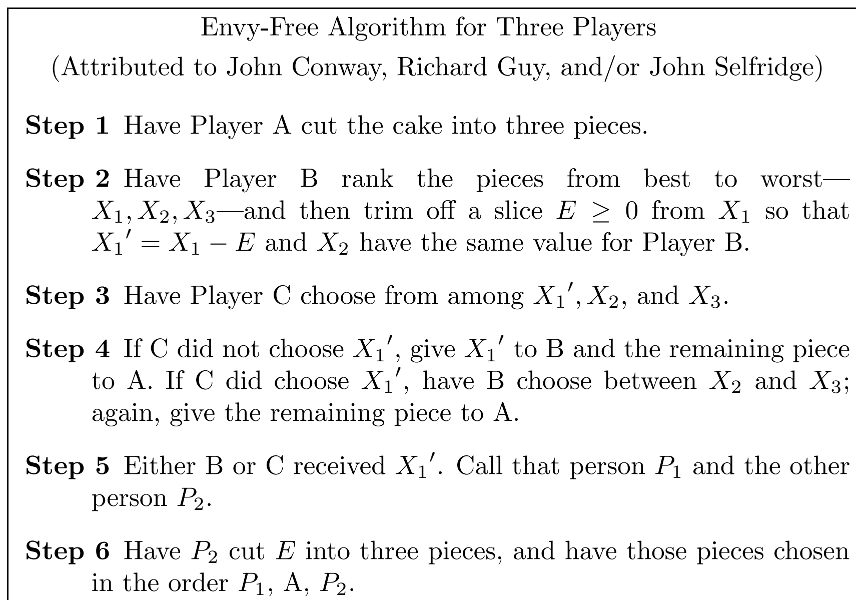

Envy-Free Algorithm for Three Players
(Attributed to John Conway, Richard Guy, and/or John Selfridge)

Step 1 Have Player A cut the cake into three pieces.

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.

Step 3 Have Player C choose from among ${X_1}', X_2$, and X_3.

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.

Step 5 Either B or C received ${X_1}'$. Call that person P_1 and the other person P_2.

Step 6 Have P_2 cut E into three pieces, and have those pieces chosen in the order P_1, A, P_2.

Figure 6.1: An envy-free algorithm for three players.

with the *minimum* number of cuts. This leads (like many questions mathematical) into an amazingly connected realm of results about mathematical topics such as Ramsey partitions and graph theory.

Question: What if the object that you are dividing up is not a *good* (e.g., cake) but a *bad* (e.g., chores)?

Answer: A few algorithms do exist 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 children in an **envy-free** manner, i.e., so that no child covets another child'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, adapted from Section 1.4 of *Cake-Cutting Algorithms*.

[2]Robertson, Jack and William Webb, *Cake-Cutting Algorithms: Be Fair If You Can* (Natick, Mass.: A K Peters, 1998) [A mathematical treatment of the topic; quite readable, but the mathematics get pretty intense]

Fair Division

The issue of envy-freeness 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"?

In their books *The Win-Win Solution: Guaranteeing Fair Shares to Everybody*[3] and *Fair Division: From Cake-Cutting to Dispute Resolution*[4], Steven Brams and Alan Taylor identify the following criteria for fairness in division problems:

Proportionality Each child should get at least $\frac{1}{n}^{th}$ of the cake according to his or her estimation.

Envy-free No child should covet another child'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 child better off and not make any child 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 Camp David peace accord and the Donald and Ivana Trump divorce.

6.3 The Importance of Trade

Observation of sandwich-swapping in any elementary school lunchroom will confirm that the initial allocation of resources is often only the beginning of the story. Just because "Mom" divides the cake in a certain way doesn't mean that things will stay that way.

The possibility of trade is in fact 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 child 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 child would be better off and nobody would be worse off. In such **Pareto inefficient** situations, at least one child would have a strong incentive to push for such a reallocation, e.g., by trading with or bribing the other children to accept that alternative allocation.

[3]Brams, Steven J. and Alan D. Taylor, *The Win-Win Solution: Guaranteeing Fair Shares to Everybody* (New York: W.W. Norton, 1999) [A very easy-to-read popular guide by two game theorists, one a political scientist, the other a mathematician].

[4]Brams, Steven J. and Alan D. Taylor, *Fair Division: From Cake-Cutting to Dispute Resolution* (Cambridge University Press, 1996) [A theoretical treatment by the same folks].

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. (Ronald Coase won the 1991 Nobel Prize in Economics in part for discussing this issue.)

For another perspective on the Coase Theorem, consider what blues musician B.B. King said during a discussion about Napster, the pioneering music-swapping website. Commenting on the copyright-infringement lawsuits filed against Napster by various 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]

The Coase Theorem indicates 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 deceptively simple. For an example that illustrates this point, see problem 5 on page 52; this problem addresses the issue of **comparative advantage**.

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, and the concomitant inattention to equity considerations or other elements of fairness, leads economists to think rather differently than non-economists. (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 child satisfies the criterion of Pareto efficiency—it is not possible to make unlucky children better off without making the lucky child worse off—but that does not necessarily make it a "good" or "fair" way to cut cake.

---

[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, see the book *Information Rules* by Carl Shapiro and Hal Varian.

[6] If 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. In the case of Napster, a variety of nasty issues led them into bankruptcy instead of into a successful bargain with the music industry. Nonetheless, it seems likely that smart people will eventually figure out a method for digital distribution of music over the internet.

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. Brams and Taylor's Adjusted Winner algorithm notwithstanding, 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 child is bigger? What if one child 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 reason and final reason is political feasibility. In the real world, there already exists a certain distribution of resources; 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

1. "Differences in opinion make fair division harder." Do you agree or disagree? Explain why.

2. *Challenge* Can you convince yourself—or someone else—that each child can get at least $\frac{1}{n}$th of the cake (in his or her estimation) with the divide-and-choose algorithm? How about with the moving-knife algorithm?

3. 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.)

4. *Challenge* Show that an envy-free division is also proportional.

5. (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) Is this a Pareto improvement over the no-trade result in question 5a?

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) Is this a Pareto improvement over the no-trade result in question 5d?

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. Fraction 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**.

Chapter 7

Economics and Social Welfare

> 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.[1] On the 8th hole, the three men begin to vent their frustrations. 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?" 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; their blindness explains their slow play." 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?"

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 to pick the best outcome for society as a whole.

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

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 economics is also called **normative economics**. In contrast, **positive economics** seeks to make predictions about the objective 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; normative economics deals with what *should* happen.

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. According to this measure of social welfare, 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 amounts 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 ethical and philosophical issues, 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: its focus on the *quantities* of costs and benefits ignores the *distribution* of those costs and benefits. Distribution is in fact irrelevant to the analysis: a policy that helps the poor and hurts the rich receives the same treatment as a policy that does the reverse. It is easy to argue that cost-benefit analysis is based on the principle of "a dollar is a dollar", and to use that principle to weigh one person's losses against another person's gains in what economists call an **interpersonal**

[2] Economists 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.

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

utility comparison.

It is less easy, but still possible, to argue 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 Nicholas Kaldor's argument that economists should emphasize the *possibility* of everyone being better off, but that *actual* distributional questions are the purview not of economists but of politicians.

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 (and earlier) 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 tolerate losers, thereby avoiding the difficulties facing cost-benefit analysis. 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, there is a concept of "tallest" or "not the tallest", and 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]

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

It is also possible to define these terms without specifically mentioning Pareto improvement. For an example, consider the cake-cutting problem or any other resource allocation problem. Some allocation of resources A is said to be 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 A is said to be Pareto inefficient if there *is* another allocation of resources that would make at least one person better off and make nobody worse off.

Complications

First, the term "Pareto improvement" compares different 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. 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.

Second, the terms "Pareto efficient" and "Pareto inefficient" apply to specific allocations of resources. Every allocation of resources is either Pareto efficient or inefficient. If it is inefficient, there exists (by definition) a Pareto improvement over it. If it is efficient, there does not exist any Pareto improvement over it. Moreover, there is usually not just one Pareto efficient allocation; in most situations there are many Pareto efficient allocations. For example, one Pareto efficient allocation in the cake-cutting game is to give the first child all the cake; then it is not possible to make the second child better off without making the first child worse off. But giving the second child the whole cake is also a Pareto efficient allocation!

7.3 Examples

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. (The advantages and disadvantages of this perspective were discussed in Section 6.3 on pages 50–52.) For example, Pareto efficiency in the cake-cutting problem can be attained simply by letting the children trade, so economists are likely to give less weight to the cake-cutting process itself

as efficient or optimal.

and more weight to the possibility of trade after the cake is cut. The economic perspective is also evident in the following examples.

Taxes and Monopoly Pricing

For the most part, what bothers normal people about taxes is having to pay them, and what bothers normal people about monopolies is that they charge too much. *This is not what bothers economists about taxes or monopolies.* The reason is simple: a policy that results in some people handing money over to the government (or to a monopolist) does not inherently result in Pareto inefficiency.

Consider, for example, a **lump sum transfer** in which the government transfers money from me to you by taxing me $500 (a **lump sum tax**) and giving that money to you (a **lump sum payment**). Such a policy *is* Pareto efficient: it is not possible to make me better off (e.g., by reducing my tax payments) without making you worse off.

What bothers economists, curiously enough, 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., 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 I am 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 $800, so I take two trips each year. Now imagine that the government imposes a $500 travel tax that raises the cost of vacations to $1300; it gives the resulting revenue to you as a lump sum payment. Note, however, that I 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 for the government to only tax me for the first trip. You would not be any worse off (you're not getting any money from the lost trip anyway), and I would be better off because I could take two trips to Hawaii, paying only $800 for the second trip, which I value at $1000.

.

Monopoly Pricing (Price Discrimination)

A firm has a **monopoly** when it is the only seller in a market. As noted previously, the high prices that make normal people dislike monopolies do not make economists dislike monopolies. The reason is simple: high prices—even exorbitant prices—are not evidence of Pareto inefficiency. It is not clear that it is possible to make consumers better off without making someone else (namely, the monopolist) worse off.

Consider, for example, a monopolist record company that can produce additional copies of a Britney Spears CD for $2. Imagine further that that a million

college students are each willing to pay up to \$10 for the CD and that a million college graduates are each willing to pay up to \$20 for the CD. If the monopolist somehow costlessly learns about all this and then charges \$10 for college students and \$20 for college graduates, it is engaging in **group pricing**, also called **third-degree price discrimination**.[5] If the monopolist then sells a million CDs to each group, it will make $[(\$20 - \$2) + (\$10 - \$2)] = \$26$ million. *This outcome is Pareto efficient*: it is not possible to make college students or college graduates better off without making the monopolist worse off.

There are, however, other aspects of price discrimination that can result in Pareto inefficient outcomes. One example involves mail-in rebates or coupons; there is something odd about having to cut out the UPC code from the box and fill out a form and waste a stamp to mail it all to some far-off place in order to get a \$10 rebate check *when they could have just given you a \$10 instant rebate at the time of purchase.* Indeed, this mail-in rebate leads to a Pareto inefficient outcome: having an instant rebate would save both you and the manufacturer time and money. (If they're going to have to send you \$10 anyway, they might as well avoid having to open your envelope, process your rebate, and waste a stamp mailing you a check.)

The economics explanation of mail-in rebates is that they are a form of **versioning**,[6] also called **screening** or **second-degree price discrimination**. Instead of directly charging different prices to different groups, monopolists engaged in versioning offer all potential customers the same menu of options, but they design the menu in such a way that different types of customers will choose different options. Mail-in rebates present customers with a menu consisting of two options: cut out the UPC code &etc in order to get the lower price, or pay \$10 more and walk away with no strings attached. Different customers will choose different options, and the monopolist will have succeeded in increasing its profits by charging different customers different prices.

A second instance of versioning concerns two IBM laser printers: the E-series version was cheaper and printed 5 pages per minute; the F-series version was more expensive and printed 10 pages per minute. In term of engineering, however, a computer magazine discovered that the only difference was that the E-Series printer had a chip in it that made it pause at the end of each line. This is clearly Pareto inefficient.

As a final example, airplane passengers may be interested in Jules Dupuit's analysis of similar circumstances in the railroad industry in France in 1849:

> 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

[5]These terms simply mean that different groups get charged different prices; discounts for students or seniors are common examples. The extreme case of group pricing is when each group consists of only one person, i.e., when each person gets charged a different price; this is called **personalized pricing** or **first-degree price discrimination**.

[6]The terms "versioning", "group pricing", and "personalized pricing" come from Carl Shapiro and Hal Varian's book *Information Rules: A Strategic Guide to the Network Economy* (Harvard Business School Press, 1999).

> 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.

According to this analysis, which is supported by modern economic theory, people buying economy class tickets on airplanes get cramped seats and lousy food in part to "frighten the rich" into riding first class. The resulting outcome is Pareto inefficient: a Pareto improvement would be to force the rich to ride first class and then provide economy travelers with a much more comfortable trip that costs only slightly more than before.

Monopoly (Uniform Pricing)

Unfortunately for monopolists (and perhaps for some potential consumers as well), engaging in price discrimination is not always an option. The monopolist may not know enough about different consumers to charge them different prices, or it may not be able to prevent **resale**. (If college students can buy CDs for \$10 and then resell them on eBay, record companies will have difficulty selling CDs to college graduates for \$20.) In such cases monopolists have little choice but to adopt **uniform pricing** by charging everybody the same price. In the CD discussed earlier, the profit-maximizing choice under uniform pricing is for the monopolist to charge \$20. (With a \$20 price, college students aren't going to buy the CD, so the monopolist makes and sells only one million CDs; profits are $(\$20 - \$2) = \$18$ million. With a \$10 price, students and graduates both buy CDs, so the monopolist makes and sells two million CDs; profits are $[(\$10 - \$2) + (\$10 - \$2)] = \$16$ million.)

The Pareto inefficiency here is clear: college students would pay up to \$10 for the CD, and the monopolist would incur costs of only \$2 to produce additional CDs. A Pareto improvement over the uniform pricing outcome is for the monopolist to continue to charge \$20 to college graduates but to charge students a price between \$2 and \$10; then nobody is worse off and either students or the monopolist (or both) are better off.

In the case of uniform pricing, then, what bothers economists is not the high prices *per se*, but the fact that high prices cause some 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 in uniform pricing.

Problems

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

2. "A Pareto efficient outcome may not be good, but a Pareto inefficient outcome is in some meaningful sense bad."

   (a) Give an example or otherwise explain, as if to a non-economist, the first part of this sentence, "A Pareto efficient outcome may not be good."

   (b) Give an example or otherwise explain, as if to a non-economist, the second part of this sentence, "A Pareto inefficient outcome is in some meaningful sense bad."

3. "Any Pareto efficient allocation is better than any Pareto inefficient allocation."

   (a) How would an economist define "better"? Can you add anything else to that definition?

   (b) Assume that we're using the economist's definition of "better". Do you agree with this claim? If so, explain. If not, provide a counter-example or otherwise explain.

4. Many magazines and newspapers have on-line archives containing articles from past issues. As with many information commodities, is costs essentially nothing for the publisher to provide readers with archived material. But in many cases they charge a fee (usually about $1) to download an article from the archive.

   (a) Sometimes the maximum I'm willing to pay is $.25, so instead of downloading it I just do without. Is this outcome efficient? If not, identify a Pareto improvement.)

   (b) Imagine that the publisher could engage in perfect price discrimination, i.e., could figure out the maximum amount that each reader is willing to pay for each article, and then charge that amount. Would the result be efficient?

   (c) Explain briefly why the outcome described in question 4b is unlikely to happen.

Chapter 8

Sequential Move Games

A **sequential move game** is a game in which the players take turns moving; examples include chess and poker. These games can be analyzed using a **game tree**, essentially a multi-person extension of the decision trees we saw in Chapter 1. The label at each node of the game tree indicates which player has to move. The various branches coming out of that node identify the different options for that player. When the player chooses, we move along the chosen branch, arriving at either another node (where it's another player's turn to move) or at the end of the game, where we find the different outcomes for the various players.

Figure 8.1 shows an example, called the grenade game. This game has two players: Player 2 has threatened to pull the pin on a grenade unless Player 1 hands over $1 million. In the first round of the game Player 1 gives Player 2 either $0 or $1 million. In the second round, Player 2 either pulls the pin on a grenade (killing both players) or does not. The outcome for each player depends on how much money they gain (or lose) and their health status.

8.1 Backward Induction

The fundamental solution concept in sequential move games is called **backward induction**. The main idea is to *look ahead and reason back*. In other words, you start at the end of the game and move back towards the beginning, solving along the way by analyzing the choices of the different players. In this way the players can anticipate their opponents' moves (and their opponents' anticipations!). Here's how backward induction works in the grenade game:

In order to figure out what to do, Player 1 needs to *anticipate* what Player 2 will do. For example, if Player 1 pays the $1 million, we get the **subgame** shown in Figure 8.2. It seems safe to assume that Player 2 will not pull the pin. So if Player 1 pays the $1 million, he can anticipate that his payoff will be (Alive, −$1 million).

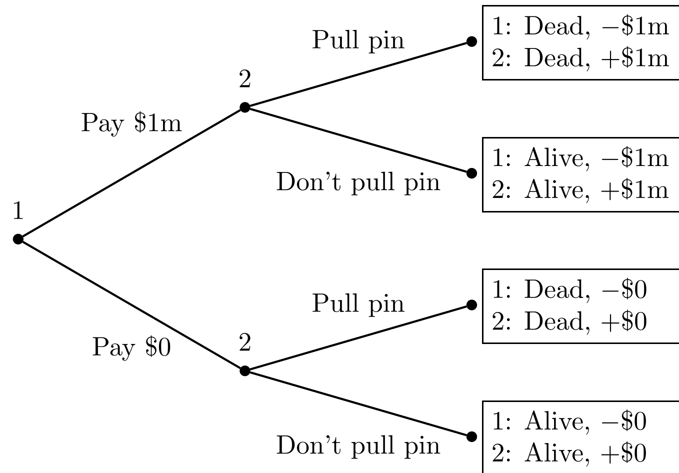

Figure 8.1: The grenade game

Similarly, Player 1 can anticipate what will happen if he doesn't pay the $1 million. In this case, we get the subgame shown in Figure 8.3. As long as Player 2 is not suicidal, it again seems safe to assume that she will not pull the pin. So if Player 1 refuses to pay the $1 million, he can anticipate that his payoff will be (Alive, −$0 million).

Combining these results, Player 1 can anticipate that the game tree boils down to the decision tree shown in Figure 8.4: if he pays the $1 million, he'll be alive and out $1 million; if he doesn't pay the $1 million, he'll be alive and out nothing. Clearly, Player 1 would rather keep the $1 million himself, so we can

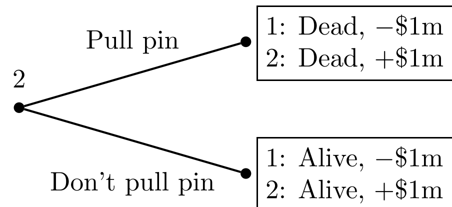

Figure 8.2: The upper subgame of the grenade game

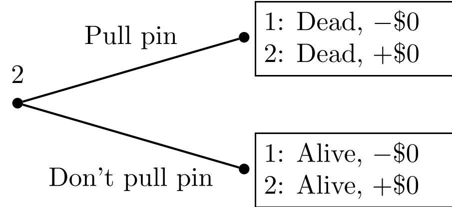

Figure 8.3: The lower subgame of the grenade game

anticipate that the outcome of this game is for Player 1 to refuse to pay the $1 million and for Player 2 to not pull the pin.

Why shouldn't Player 1 be worried about ending up dead? Because Player 2's threat to pull the pin is a **non-credible threat**. Player 2 can threaten all she wants, but when the time comes it's not going to be in her self-interest to pull the pin.

Question: How might Player 2 make this threat credible?

Answer: Change the structure of the game tree to eliminate her ability to opt-out; in other words, create a **credible threat**. For example, Player 2 could, prior to hitting Player 1 up for money, borrow $1m from the mafia, explicitly requesting that the mafia torture her if she doesn't repay. Then Player 2 could go out and blow that money on wild parties. The day before the mafia comes calling, Player 2 hits up Player 1 for the money. The threat to pull the pin is now credible, because Player 2 might prefer death by grenade to torture at the hands of the mafia.

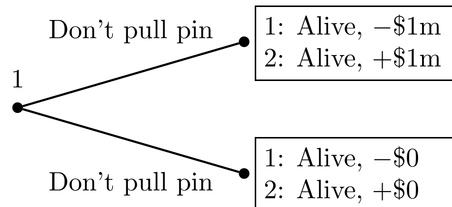

Figure 8.4: The grenade game reduced to a decision tree

Problems

1. *Challenge* Explain (as if to a non-economist) why backward induction makes sense.

2. 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.

3. (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.

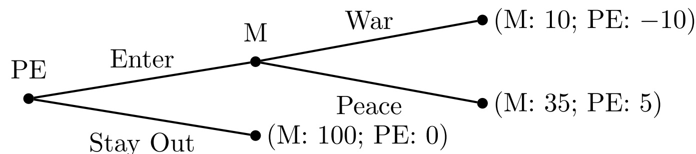

Figure 8.5: Story #1

Story #1 (See figure 8.5): 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.)

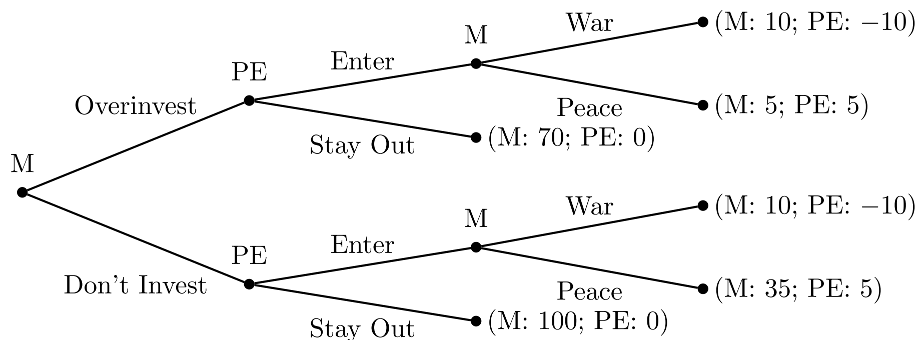

Figure 8.6: Story #2

Story #2 (See figure 8.6): 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.

(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.)

4. (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? [Hint: Use backwards induction!]

   (b) *Challenge* If there are n sticks on the table, which player would you rather be? Can you describe a general strategy?

   (c) *Super Challenge* (Note: I haven't thought much about the second question here, and in any case haven't solved it.) Analyze this problem when there are m piles of sticks (each turn one of the players picks one of the piles and removes one or two sticks from it) or when there are more than two players (the losing player must pay $1 to each of the others).

5. (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 nth period (at which point Mom eats the remaining pie, meaning that the players get nothing).

   (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.

6. Make up some game trees (players, options, payoffs, etc.) and solve them using backward induction. It may help to exchange games and solutions with a neighbor.

7. (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: First 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 8.7.*

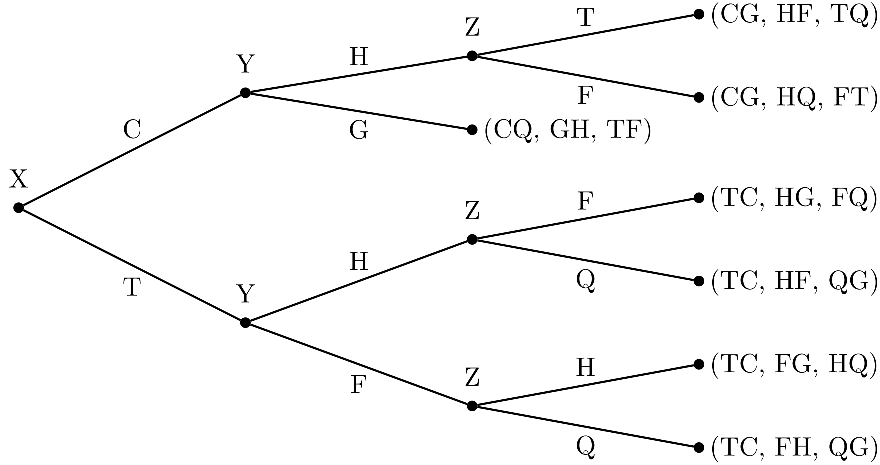

Figure 8.7: The draft game

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 #2), and team Z gets the Tailback and the Fullback (#1 and #4). 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.

(a) 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?

(b) 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.

(c) What outcome do you expect from this game using backward induction?

(d) Is the expected outcome you identified Pareto efficient? If so, explain. If not, identify a Pareto improvement.

(e) 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.

(f) *Challenge* Prove that the game tree really does boil down to what's shown on the previous page.

8. *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. You 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.

9. *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.

10. (The Centipede Game) There are 6 $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?

11. *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.

Chapter 9

Simultaneous Move Games

Another kind of game is one of **simultaneous moves**: two or more players move at the same time (an example here is the game "Rock, Paper, Scissors"). Although it is possible to analyze these games using a game tree, it is often easier to use a **payoff matrix**. An example is the **duopoly**[1] game shown in Figure 9.1. There are two firms, Coke and Pepsi. The two firms simultaneously choose "high" or "low" prices. If they both choose high prices, they both make large profits: $4 million each. If they both choose low prices, they both make small profits: $3 million each. If one chooses a high price and one chooses a 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. (Note: By convention, the row-playing firm (in this case, Coke) is called Firm 1, and the column-playing firm (Pepsi) is called Firm 2. The payoffs in each box are listed in the order of the firms: (5, 1) indicates that Firm 1 (Coke) gets a payoff of 5 and Firm 2 (Pepsi) gets a payoff of 1.)

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

Figure 9.1: The duopoly game

[1] A market with only two sellers is called a duopoly, just like a market with only one seller is called a monopoly.

9.1 Dominance

One perspective on this game comes from the observation that Coke always makes more money by choosing a low price *regardless of what Pepsi does*: 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. In other words, "low price" is a **strictly dominant strategy** for Coke: no matter what Pepsi does, choosing the low price produces a higher payoff for Coke.

Similarly, we can see that "low price" is a strictly dominant strategy for Pepsi. We can predict, therefore, that both firms will choose the low price, yielding profits of $3 million for each firm.

9.2 The Prisoner's Dilemma

Figure 9.1 is one version of an important game called the Prisoner's Dilemma. Another version (which explains the game's name) is shown in Figure 9.2. The story is this: You are visiting an unnamed foreign country, waiting in line to buy a Coke 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 (e.g., 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 on his record, 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 (big surprise) 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? Well, then you both get locked up for 5 years (for armed robbery). And if you both keep quiet, you both get locked up for 1 year (for littering). (This information is summarized in Figure 9.2.)

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

Figure 9.2: The Prisoner's Dilemma

Question: What is a good prediction for the Prisoner's Dilemma game if both players are optimizing?

Answer: The only reasonable prediction is for both of you to confess because confessing is a strictly dominant strategy: no matter what the other guy does, the best thing for you to do is confess, and no matter what you do, the best thing for the other guy to do is confess. Curiously, there is another outcome which both of you strictly prefer to this outcome: you'd both be better off (i.e., it would be a Pareto improvement) if you both kept quiet. But the incentive structure is such that you both confess!

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

- You and everybody else might like to drive your own cars to work, but perhaps we would all be better off with carpools or buses instead of traffic jams.
- Your firm and your 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. (Note: This is exactly what happened in the Coke and Pepsi game we saw earlier.)
- 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, forcing your classmate to do all the work, and getting credit for their effort. Of course, if both of you do this you'll both get horrible grades...
- Finally, it is worth noting that the Prisoner's Dilemma game is a reasonable representation of some of the dangers of plea-bargaining. There really is a dilemma for those accused of crimes: suspects might agree to testify against others out of their own self-interest, with harmful ramifications for the integrity of our criminal justice system.

How can we reach the cooperative outcome in the Prisoner's Dilemma game? To examine this possibility, we'll use the version of the Prisoner's Dilemma game shown in Figure 9.3; this version involves cash prizes rather than prison terms.

		Player 2	
		D	C
Player 1	D	0,0	10,-5
	C	-5,10	1,1

Figure 9.3: Another version of the Prisoner's Dilemma

The idea of dominant strategies makes for a clear prediction for this game: Player 1 plays D, Player 2 plays D, and each player gets nothing. Clearly this outcome is not Pareto efficient: the strategy combination (C, C) would allow each player to get one dollar. So we can think of (C, C) as the **cooperative strategy** and ask under what conditions we might be able to see this kind of cooperative behavior.

9.3 Finitely Repeated Games

One possibility is to play the same game over and over again: perhaps the players can establish reputations and learn to trust each other, and then maybe we can get cooperative behavior. So we turn our attention to **repeated games**, in which a smaller game (called the **stage game**) is played over and over again (n times). We will take the Prisoner's Dilemma game in Figure 9.3 as the stage game, and see what the potential is for cooperation.

We know that if $n = 1$ (i.e., the stage game is only played once) the potential for cooperation is quite slim: playing D is a strictly dominant strategy for each player: no matter what Player 2 does, Player 1 is better off playing D, and vice versa. So we should not expect to see cooperative behavior if the game is played only once.

An extension of this logic suggests that we should not expect to see cooperative behavior if the game is played twice, or 10 times, or even 100 times. To see this, apply backward induction: in the final game (whether it's the second playing of the stage game or the 100th), playing D is a strictly dominant strategy for both players, so both players should expect the outcome (D, D) in the final game. We now look at the penultimate (second-to-last) game: knowing that (D, D) will be the outcome in the final game, playing D becomes a strictly dominant strategy in the penultimate game as well! We continue working backwards all the way to the beginning of the game, and conclude that the only rational prediction is for the players to play (D, D) in each of the stage games. So we cannot expect any cooperative behavior.

Digression: This theory does not match up well with reality. In classroom experiments (or experiments on "regular people" conducted by **experimental economists**), what tends to happen is this: if the stage game is repeated 10 times, players will cooperate for about 5 or 6 stages, and then one player (or

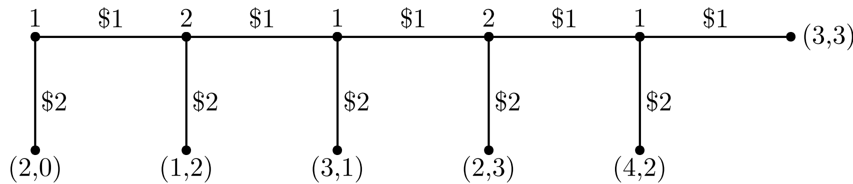

Figure 9.4: The centipede game

possibly both players) will stop cooperating and the noncooperative outcome (D, D) occurs for the remainder of the game. Maybe people's deductive powers simply don't work that well? Or maybe people's deductive powers are stronger than we give them credit for! For example, most people playing the game in Figure 9.4 do better than game theory predicts. This game is called the centipede game (because the game tree looks like a stick-figure centipede), and the story behind it is this: There are six $1 bills on a table. Players 1 and 2 take turns moving. On each turn the player moving takes either $2 from the table (in which case the game ends) or $1 (in which case it becomes the other player's turn). The payoffs (which become much more impressive if you think of a game involving six $100 bills instead of six $1 bills) are given at each terminal node of the game tree. (Exercise: Predict the outcome of this game using the tools of game theory. Then ask yourself what strategy you would follow if you were Player 1, or if you were Player 2 and Player 1 started off by taking $1. Then, if you're really looking for a challenge, try to make sense of the whole mess. Is the theory wrong, or are the players wrong, or what?)

Problems

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!

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 in the simultaneous-move game, i.e., 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!

3. Consider the following game.

		Player 2	
		D	C
Player 1	D	$-2, -2$	$10, -5$
	C	$-5, 10$	$1, 1$

   (a) What outcome do you expect if this game is played once? Explain briefly.

   (b) (5 points) What outcome do you expect if this game is played twice? Explain briefly.

Chapter 10

Application: Auctions

> José robs a bank in Texas and hightails it across the border to Mexico with the Texas Rangers in hot pursuit.[1] 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."

The morals of this story are that people don't always do what we want them to do, and that people don't always tell the truth. These morals are at the heart of our next topic: auctions.

Motivating question: Why use auctions?

Answer[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 of sale is also important for commodities with a short shelf life. The Aalsmeer Flower Auction in the Netherlands—the largest flower auction in the world, with $1 billion in sales each year—completes about 50,000 auctions each morning, with each auction taking only a few seconds. (More on this below.)

Second, auctions prevent dishonest dealing between the seller's agent and the buyer. This can be quite important in government and business contracts.

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

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

For example, let's say I'm a procurement manager for Company X, and my company 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 at Company X (and the company's shareholders) might have some questions about whether my brother's copier was really the best choice, whether I got the best deal I could for Company X, &etc. 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.

10.1 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, used in the wholesale flower market in the Netherlands (and elsewhere), is often conducted with the help of a **Dutch auction clock**. (You can get a virtual tour of the Aalsmeer Flower Auction, which completes 50,000 auctions each morning, sells $1 billion in flowers every year, and holds the Guinness record for the world's largest commercial building, online.[3]) 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 (and agrees to pay whatever price it displays) by pushing their button.

(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 but 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 but stands up when they're willing to claim the object, the winner will be the first bidder standing.)

[3] http://www.vba.nl

A third type 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 bid. So if you bid $10 and are the highest bidder then you get the object 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 gets the object, but pays only $7.

Question: If you want to auction off an item, which kind of auction should you use?

Answer: Tough! We'll come back to this question at the end. To get an idea of the complexities involved, however, consider this question: What kind of moron would run a second-price sealed-bid auction instead of a first-price sealed bid auction? To answer this, let's ask yet another question:

Question: Do the bidders have identical strategies in the first- and second-price sealed bid auctions?

Answer: No! Let us say that a buyer's **true value** for some object is $100 if that buyer is indifferent between having $100 and having the object. In a first-price sealed bid auction, then, each buyer has 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 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.

To see why, 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 less than $80 would produce a worse outcome (you'd end up losing an auction that you would have rather won.)

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). 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 rather lost).

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: bidders have an incentive to shade their bids below their true value in a first-price sealed bid auction, but not in a second-price sealed bid auction. So it's no longer clear that first-price auctions will yield higher profits than second-price auctions!

10.2 Auction Equivalences

Imagine that you're a bidder in a descending price (Dutch) 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 your button.

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: 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 (English) 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 continue bidding. If the auctioneer could see these sit-down prices for all of the bidders, he 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

> 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:
>
> 1. Decide the **maximum** you're willing to pay and enter this amount.
> 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.
> 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 10.1: eBay's Proxy Bidding feature

$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.[4]) For further supporting evidence, Figure 10.1 duplicates the online[5] description of the Proxy Bidding feature on the auction house 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 being the ascending price auction and the second-price sealed bid auction, the other being the descending price auction and the first-price sealed bid auction—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 dominant strategy, namely to continue bidding as long as the asking price is less than your 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

[4] The 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.

[5] http://pages.ebay.com/help/buyerguide/bidding-prxy.html

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:

Question: If you want to auction off an item, which kind of auction should you use?

Answer: We've already seen that ascending-price (English) and descending-price (Dutch) auctions are equivalent to second- and first-price sealed bid auctions, respectively. We've also seen that a first-price sealed-bid auction is not (as it seems as first) clearly superior to a second-price sealed bid auction because bidders will shade their bids in a first-price auction. A more difficult result that requires some calculus (and some further assumptions) is that 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 remarkable result is called the **Revenue Equivalence Theorem**.

10.3 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 lobbying groups don't get their money back!

Queueing (waiting in line) For example, for tickets to crowded shows. In this case the potential buyers pay with their time (time is money, remember), and the folks 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 (often large) sums of money in the hopes of being the first firm to invent something. 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 (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**.

The thing that makes a common values auction interesting is that none of the bidders know what that common value is. For example, the oil companies bidding for the rights to the oil under a given plot of land might not know how much oil is really there. Instead, the different bidders must rely on estimations, some of which are probably higher than the true amount, some probably lower. We can imagine that *on average* they are correct, i.e., that their estimates fluctuate around the true amount. But if each bidder bids his or her own estimate, the winner will be the bidder with the highest estimate—an estimate that almost certainly exceeds the true amount! Hence the winner's curse: because her estimate is higher than the average estimate, the winner gets the rights to the oil; but because her estimate is higher than the average estimate, the amount of oil in the ground is almost certainly less than she thinks, meaning that she might very well end up losing money in the end.[6]

Multi-Unit Auctions

The FCC (Federal Communications Commission) and similar bodies in other countries have and continue to conduct spectrum auctions to allocate various wavelengths of airspace (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 10.2 and the accompanying text, both excerpted from a working paper[7] on the 1996-1997 FCC auction:

> [Figure 10.2] 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.

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

[7]Peter 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 10.2: Code bidding in the FCC spectrum auction

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

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)?

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.

3. You're a bidder in a first-price sealed bid auction. 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.

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.

5. We can use the expected value calculations from Chapter 2 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 Value(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?

Chapter 11

Application: Marine Affairs

> 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 2[1]

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 caught by a fisherman trying to recover lost equipment off the coast of Georgia.

[1]This 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!

The fishery rapidly expanded, with catch totalling 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.

11.1 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.

11.2 A Brief History of 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 23.) Although economists would quibble with the establishment of MSY as the desired objective (see Section 3.4 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

[2]For 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 rammings 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.

11.3 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 12

Transition: Game Theory v. Price Theory

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 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 op-

[1] For 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.

timizing 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.

12.1 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 invest-

ments 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**.

12.2 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.

Problems

1. Would you say that that market for new cars is a competitive market? Why or why not? How about the market for used cars?

Part III

Many v. Many

Chapter 13

Supply and Demand: The Basics

If we go to the grocery store or to the gas station, what can we observe? Well, we can observe the **quantity transacted** (the number of units of milk or gas bought and sold) and we can observe the **market price** (the amount the buyers pay the sellers for each unit). Figure 13.1a graphs these observables: the market price P is on the vertical axis and the quantity transacted Q is on the horizontal axis. (This somewhat counterintuitive set-up is explained further in Chapter 17.)

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

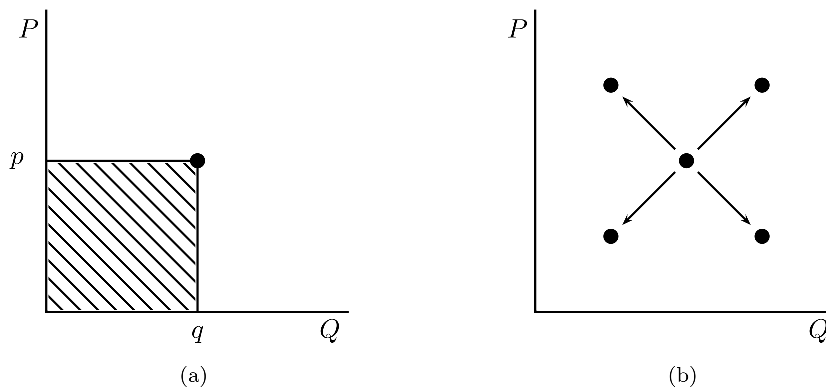

Figure 13.1: (a) the observables, p and q, and pq = total revenue (and/or total expenditure); (b) the observables changing over time. Note that we have the quantity transacted Q on the x-axis and the market price P on the y-axis.

that changes hands—i.e., 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 or the gas station, we can see these observables change over time: milk goes on sale and the quantity transacted increases, the price of gasoline increases and the quantity transacted decreases, etc. Graphically, we get Figure 13.1b.

13.1 The Story of Supply and Demand

The story of supply and demand attempts to explain what's going on with p and q. The story is that the dot we see is really the intersection of two curves: the market supply curve and the market demand curve. (See Figure 13.2.) When the dot moves around, it is because these curves are moving.

The **market supply curve** (see Figure 13.2a) is a graphical presentation of *hypothetical* questions like these: *If* the market price were $6, how many units of this good would sellers want to sell? *If* the market price were $7, how many units of this good would sellers want to sell? Put all these hypothetical questions together, and you get the market supply curve, which answers the following question: *If* the market price were p, how many units of this good would sellers want to sell?

Similarly, the **market demand curve** (see Figure 13.2b) is a graphical presentation of *hypothetical* questions like these: *If* the market price were $6, how many units of this good would buyers want to buy? *If* the market price were $7, how many units of this good would buyers want to buy? Put all these hypothetical questions together, and you get the market demand curve, which answers the following question: *If* the market price were p, how many units of this good would buyers want to buy?

Note from Figure 13.2 that supply and demand curves move in opposite directions. The supply curve is *upward sloping*, indicating that as the price rises sellers want to sell more (or, equivalently, that as the price falls sellers want to sell less). The demand curve is *downward sloping*, indicating that as the price rises buyers want to buy less (or, equivalently, that as the price falls buyers want to buy more).[1] We will examine these slopes in more detail in Chapter 16.

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

Answer: Look at the price and quantity amounts at the intersection point in Figure 13.2c. The quantity that buyers are willing to buy at that price is equal to the quantity that sellers are willing 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

[1] This latter result is known as the **Law of Demand**.

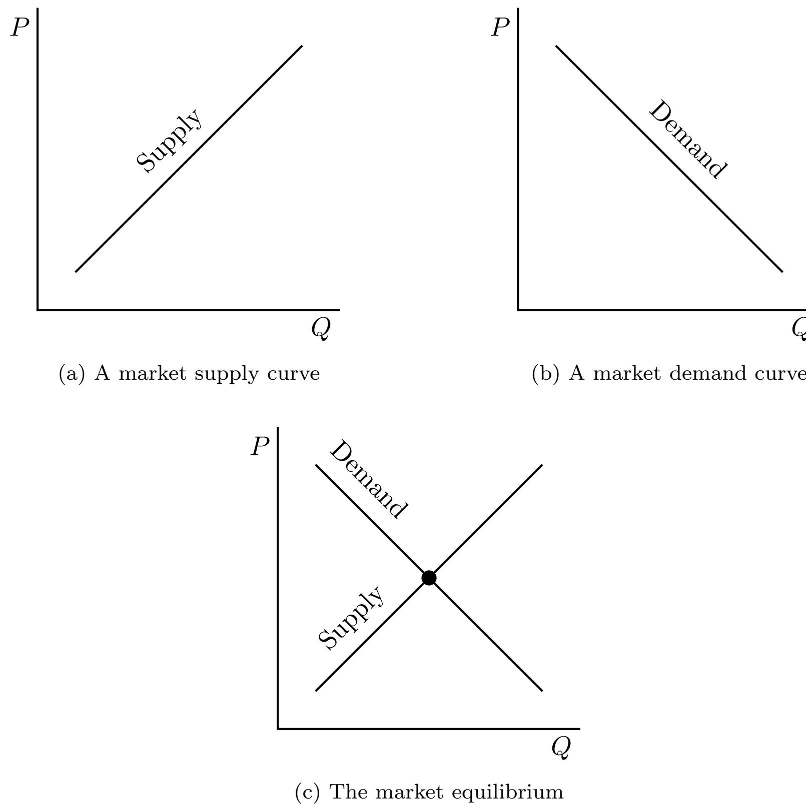

(a) A market supply curve

(b) A market demand curve

(c) The market equilibrium

Figure 13.2: Supply and demand

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 a price lower than the market-clearing price, the quantity that buyers want to buy exceeds 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 supplied equal the quantity demanded.

13.2 *Math*: 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 simultaneously by using the demand curve to substitute for q in the supply curve:

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

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 $p = 1$ and a quantity of $q = 17$.

13.3 Shifts in Supply and Demand

A good way to think about the market equilibrium is to imagine that the demand curve is blue and the supply curve is yellow. The market equilibrium (which you can imagine in green) comes at the point where the two curves intersect. 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 many 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 13.3:

Demand increases An increase in demand means that buyers want to buy more than before at any given price. The demand curve shifts to the *right*, and the equilibrium price and quantity both increase. (See Figure 13.3a.)

Demand decreases A decrease in demand means that buyers want to buy less than before at any given price. The demand curve shifts to the *left*, and the equilibrium price and quantity both decrease. (See Figure 13.3b.)

Supply increases An increase in supply means that sellers want to sell more than before at any given price. The supply curve shifts to the *right*; the equilibrium price decreases and the equilibrium quantity increases. (See Figure 13.3c.)

Supply decreases A decrease in supply means that sellers want to sell less than before at any given price. The supply curve shifts to the *left*; the equilibrium price increases and the equilibrium quantity decreases. (See Figure 13.3d.)

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

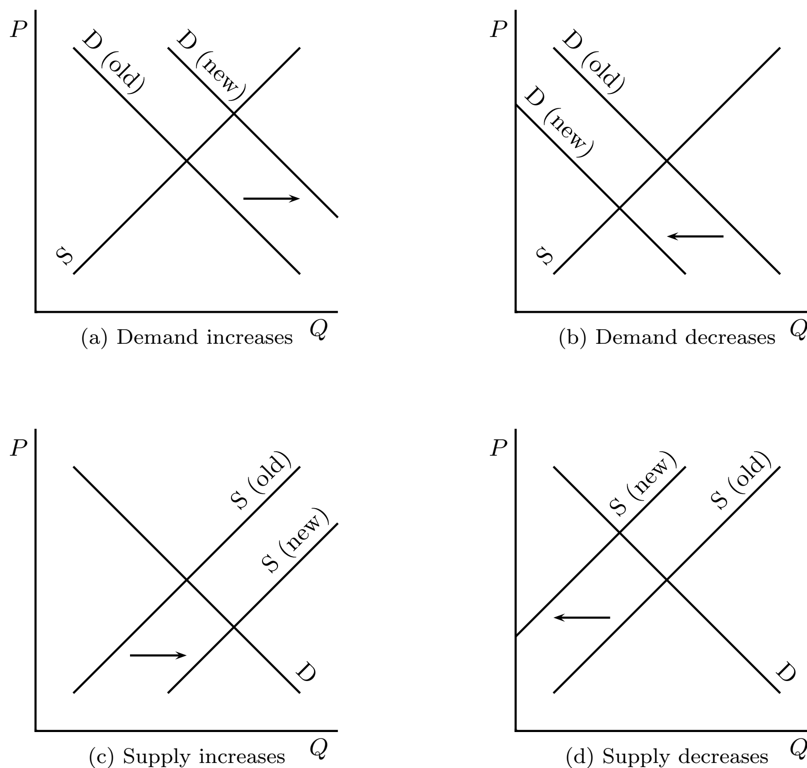

(a) Demand increases

(b) Demand decreases

(c) Supply increases

(d) Supply decreases

Figure 13.3: 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.

people want to buy or sell less). The curves move left or right, *not up or down*! This counterintuitive behavior stems from the fact that we have the quantity transacted Q on the x-axis and the market price P on the y-axis. Chapter 17 returns to this issue.

Problems

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.

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.)

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.

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.

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.

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.)

7. A few quarters ago a libertarian named Joel Grus came to class, and one of the things he discussed was drug legalization. So: for each item below, indicate the likely impact on the supply and demand for cocaine. Then indicate the effect on the equilibrium price and quantity. It may help to use a graph.

   (a) Military aid to Columbia to help that country exterminate the coca plant from which cocaine is made.

   (b) Sentences of life in prison for cocaine traffickers and members of drug-selling gangs.

   (c) Sentences of life in prison for buyers of cocaine.

   (d) Drug treatment programs to try to help addicts stop using drugs.

   (e) News reports that eating nutmeg has the same effect as snorting cocaine.[2]

   (f) Finally: Imagine (hypothetically) that legalizing cocaine would have no effect on the demand for that drug. Describe the impact of legalization on the market for cocaine.

   (g) This question is a follow-up to the previous one: Would the legalization of cocaine necessarily lead to more money being spent on cocaine, i.e., would total expenditures necessarily go up?

8. For each item, indicate the likely impact on the supply and demand for beef. Then indicate the effect on the equilibrium price and quantity. It may help to use a graph.

   (a) Mad cow disease scares away meat-eating shoppers.

   (b) All the cows in England are slaughtered and thrown into the ocean.

   (c) New drugs make cows grow faster at lower cost to farmers.

   (d) The price of chicken falls.

   (e) News reports that eating red meat is bad for you.

[2]It doesn't, really. But nutmeg *is* a hallucinogen, if taken by the tablespoonful. Based on my experience as a camp counselor, I'd recommend against it: you're almost sure to end up in the hospital, and maybe in the morgue. . . .

9. Read the following excerpt from the *New York Times* of October 27, 2000.

> For 10 years, British officials consistently misled the public by deliberately playing down the possibility that mad-cow disease could be transmitted to humans, an official report said today. . . The 4,000-page report, published after a three-year investigation. . . severely criticized the "culture of secrecy" that characterized the government's response to a crisis that has wreaked havoc with Britain's once-proud beef industry, forced the slaughter of almost four million cows and led to the deaths so far of 77 Britons. . . "My own personal belief would be that we are more likely looking in the region of a few hundred to several thousand more" victims, Prof. Peter Smith, acting head of the government's advisory committee on the disease, said on television this morning. "But it must be said that we can't rule out tens of thousands."
>
> [It was] "a consuming fear of provoking an irrational public scare". . . that caused a government veterinary pathologist to label "confidential" his first memo on mad-cow disease in 1986; that led John Gummer, then the agriculture minister, to make a show of publicly feeding a hamburger to his 4-year-old daughter, Cordelia, in 1990; and that led Britain's chief medical officer in 1996 to declare, "I myself will continue to eat beef as part of a varied and balanced diet." . . . At the same time, government policy was marred by bureaucratic bungling, a lack of coordination between departments and the fact that the Ministry of Agriculture, Fisheries and Food had two somewhat contradictory missions: to protect consumers and to support the beef industry.

(a) What was the effect of mad-cow disease on the demand curve for British beef? Draw a supply and demand graph and indicate the effect on the equilibrium.

(b) What was the effect of mad-cow disease on the supply curve for British beef? Draw a supply and demand graph and indicate the effect on the equilibrium.

(c) Combining your answers to the previous two questions, can you predict with certainty what happened to the equilibrium price for British beef? What about the equilibrium quantity?

(d) The British government put together quite a PR campaign in their effect to avoid "an irrational public scare". Would such a public scare have pleased or displeased the following groups?

   i. Die-hard beef eaters
   ii. Die-hard chicken eaters
   iii. Beef suppliers

iv. Chicken suppliers

10. 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. 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 Artic 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 fuzzy math and 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?

Chapter 14

Taxes

Changes in the tax structure also affect the market equilibrium. 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 14.1.

14.1 A Per-Unit Tax Levied on the Sellers

Imagine that the government imposes a **per-unit tax** of $.40 per gallon on the sellers of gasoline. (Later in this chapter we'll discuss **ad valorem taxes**, which are taxes based on the sale amount, e.g., a 10% sales tax.) One thing is obvious: the demand curve for gasoline is not going to change because of a tax

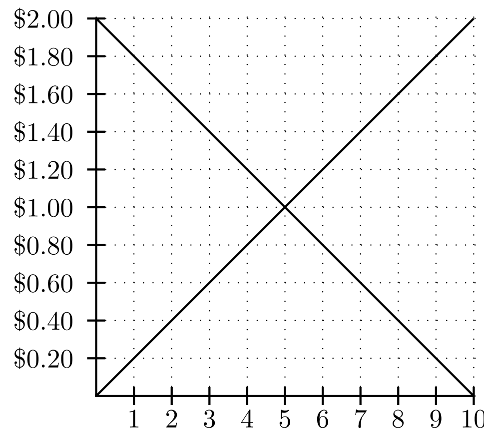

Figure 14.1: A hypothetical market for gasoline, with price in dollars per gallon and quantity in millions of gallons per day.

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.

Consider how much sellers want to sell at a price of \$1.40. Figure 14.1 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 \$.40 per gallon.

But this is easy! If the sellers receive \$1.40 per gallon but have to pay a tax of \$.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 \$.40 tax as they wanted to sell at \$1.00 per gallon without the tax. That amount, as we can see from Figure 14.1, is 5 million gallons.

The same logic applies along the entire supply curve. With a tax of \$.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 14.2.

The punch line: 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. First, *even though the amount of the tax is \$.40, the market price doesn't change by \$.40.* The market price only changes by \$.20. Second, *even though the \$.40 per gallon tax is levied on the sellers, the sellers do not end up worse off by \$.40 per gallon.* The sellers receive \$1.20 per gallon from the buyers, so after paying the tax of \$.40 per gallon the sellers end up with \$.80 per gallon. This is only \$.20 per gallon less than the \$1.00 per gallon that they received before the tax. Third, *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 \$1.20 per gallon; this is \$.20 more per gallon than they paid before the imposition of the tax. We will return to these issues later in this chapter.

Fourth and finally, Figure 14.2 shows that *a per-unit tax of \$.40 on the sellers shifts the market supply curve* up *by \$.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 17.

14.2 A Per-Unit Tax Levied on the Buyers

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 \$.10 in tax, so they should want to sell the same quantity they wanted to sell at a price of \$.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 \$.40, sellers should want to sell

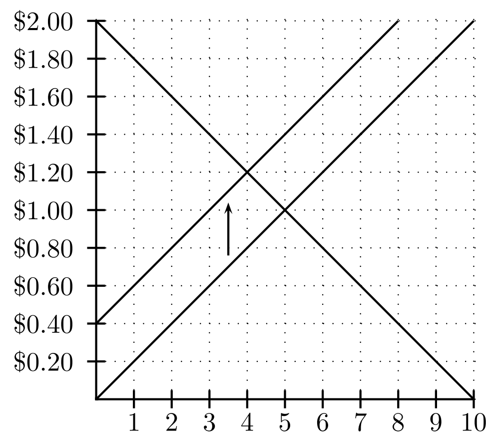

Figure 14.2: Effect of a $.40 per gallon tax on the sellers

at a market price of $1.00 the same amount that they wanted to sell at a price of $1.40 without the subsidy.

The same logic also applies for taxes placed on the buyers instead of the sellers. Let's do the $.40 per gallon gasoline tax example 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.* When the tax is on the seller, the buyer pays the seller the market price p and then the *seller* pays the government the tax. When the tax is on the buyer, the buyer pays the seller the market price p and then the *buyer* pays the government the tax.

With that in mind, imagine the government imposes a $.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 14.1 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 $.40 per gallon.

Again, this is easy. If the buyers pay the sellers $1.40 per gallon but have to pay an additional tax of $.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 $.40 tax as they wanted to buy at $1.80 per gallon without the tax. That amount, as we can see from Figure 14.1, is 1 million gallons.

The same logic applies along the entire demand curve. With a tax of $.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 14.3.

The punch line: 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

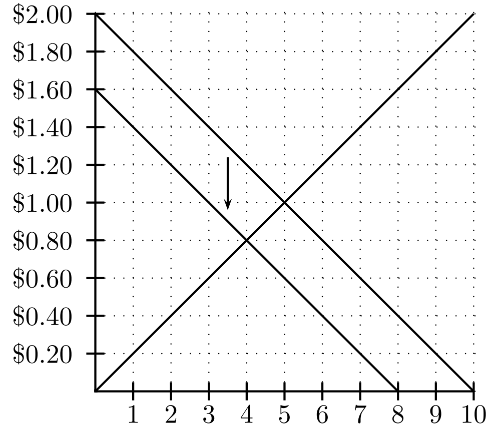

Figure 14.3: Effect of a $.40 per gallon tax on the buyers

per gallon; after the tax, the market-clearing price is $.80 per gallon.

As with the result in the case of a tax on the sellers, this result has a number of curious features. First, *even though the amount of the tax is $.40, the market price doesn't change by $.40.* The market price only changes by $.20. Second, *even though the $.40 per gallon tax is levied on the buyers, the buyers do not end up worse off by $.40 per gallon.* The buyers pay $.80 per gallon to the sellers and $.40 per gallon in taxes, so in the end the buyers end up paying $1.20 per gallon. This is only $.20 per gallon more than the $1.00 per gallon that they paid before the tax. Third, *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 $.80 per gallon; this is $.20 less per gallon than they received before the imposition of the tax. We will return to these issues later in this chapter.

Fourth and finally, Figure 14.3 shows that *a per-unit tax of $.40 on the buyers shifts the market demand curve* down *by $.40, the amount of the tax.* Again, this is not a coincidence, and we will return to this issue in Chapter 17.

14.3 Tax Equivalence

Motivating questions: Which is better for buyers, a tax on buyers or a tax on sellers? What determines who bears the ultimate economic burden of the tax?

To answer this question, 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. The tax results are as follows.

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 $.40 in tax, so what the sellers really get is $.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 $.80 per gallon. After the imposition of the tax, the sellers receive $.80 for each gallon of gasoline. The buyers pay $.80 per gallon to the seller, plus $.40 to the government, so what the buyers really pay is $1.20 per gallon.

Conclusion: 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.

Question: Prior to the imposition of the $.40 tax, the buyers paid $1.00 per gallon and the sellers got $1.00 per gallon. After the imposition of the tax (on either the buyers or the sellers), the buyers end up paying $1.20 per gallon, and the sellers end up receiving $.80 per gallon. So the buyers and sellers are both worse off by $.20 per gallon, meaning that they share the tax burden equally. Do buyers and the sellers always share the tax burden equally?

Answer: No.[1] The guiding principle turns out to be this: *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, the party that is more sensitive to price changes will bear less of the tax burden, and the party that 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, "No! If you put the tax burden on me, I'll buy a lot less!" (or "I'll sell a lot less!"). The party with the weaker sensitivity to price changes can only say, "Well, if you put the tax burden on me, I'll buy a little bit less" (or "I'll sell a little bit less"). The result is that the party with the strongest sensitivity to price changes effectively pushes the lion's share of the tax burden onto the other side.

[1] Note, however, that the *combination* of the two tax burdens always equals the amount of the tax. In the examples above, the buyers and sellers combined are worse off by $.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 $.40 per gallon, and that money has to come from somewhere.

14.4 *Math*: 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 buyers of \$1 per unit, what buyers really have to pay is $p+1$. 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

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

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

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

The same logic applies to taxes on the sellers. 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

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

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., $.8p$; if the before-tax supply curve is $q = 6000 + 800p$, the after-tax supply curve will be

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

Problems

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.)

2. Figure 14.4 shows a hypothetical market for oranges. Use it (and the replicas on the following pages) to answer the questions in the remaining problems in this chapter.

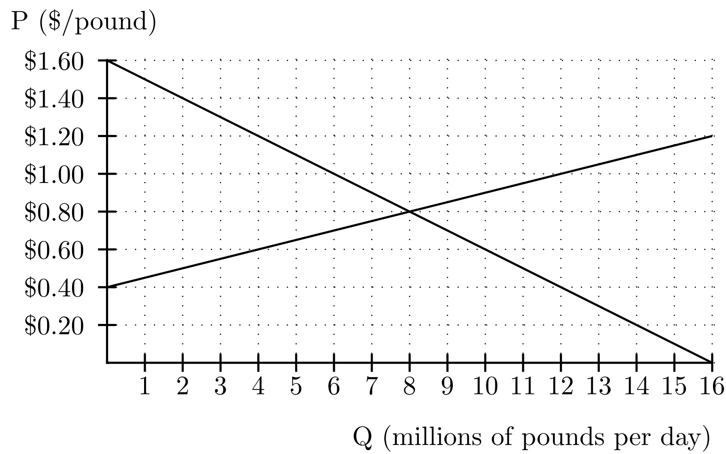

Figure 14.4: A hypothetical market for oranges

(a) What is the equilibrium price and quantity? (Use correct units!)

(b) 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 $\left(\frac{S_D}{S_S}\right)$.

3. Suppose that the government imposes an excise tax of $.60 per pound on the sellers of oranges.

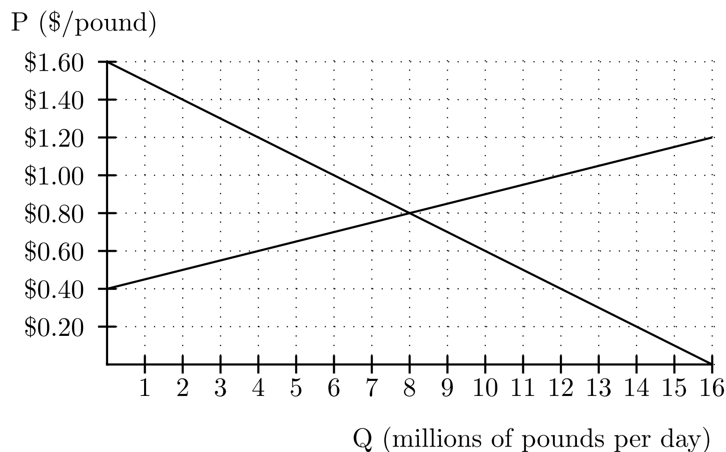

(a) Show the impact of this tax on the supply and demand curves.

(b) At the new equilibrium, how many oranges will people eat? (Use correct units!)

(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) Use your answers to the previous two questions to determine the economic incidence of the tax. In other words, calculate the amount of the tax burden borne by the buyers (T_B) and by the sellers (T_S), and the ratio $\frac{T_B}{T_S}$.

(g) Compare the tax burden ratio with the ratio of the slopes from problem 2b. Can you explain the intuition behind this result?

4. Answer the same questions as in problem 3, but now suppose that the government imposes a per-unit tax of $.60 per pound on the buyers of oranges. (Recall that the buyer now has to pay the government *in addition to* paying the seller.)

5. How do your answers in problem 4 compare with those in problem 3? What does this suggest about the difference between a per-unit tax on buyers and a per-unit tax on sellers?

6. Answer the same questions as in problem 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 $.50 and have to pay the government $.50; if they sell a pound of oranges for $2, they get to keep $1 and have to pay the government $1.)

7. Answer the same questions as in problem 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.)

8. How do your answers in problem 7 compare with those in problem 6? What does this suggest about the difference between a sales tax on buyers and a sales tax on sellers?

Chapter 15

Elasticities

Let's go back and look at the demand curve in a little more detail. One interesting question is this: how sensitive are buyers to changes in price? In other words, how much less (or more) will buyers hypothetically want to buy if the price hypothetically goes up (or down) by a little bit?

This question is relevant, for example, in addressing global warming. One way to get people to use less gasoline or other fossil fuels linked to global warming is to increase the price. This idea has led some environmentalists (and many economists) to advocate the use of a **carbon tax** on gasoline and other 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.

One way to measure the sensitivity of demand is to just eyeball the slope of the demand curve. Figure 15.1a shows a "flat" demand curve, one which appears to be very sensitive to price changes: a small decrease (or increase) in price makes the quantity demanded go way up (or way down). Figure 15.1b shows a "steep" demand curve, one which appears to be very insensitive to price

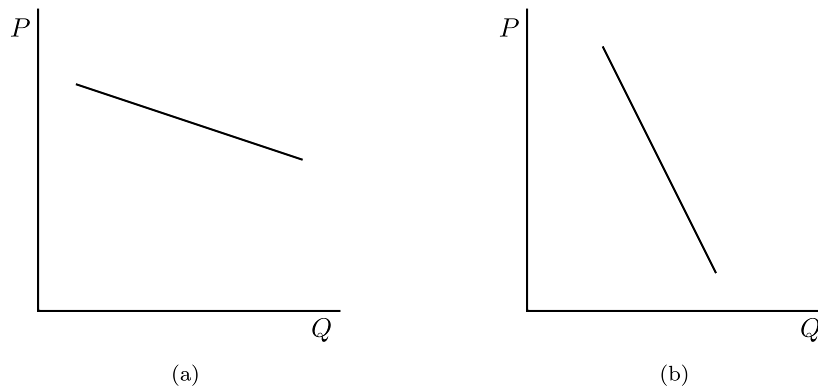

Figure 15.1: (a) a "flat" demand curve; (b) a "steep" demand curve

changes: a small decrease (or increase) in price has only a small impact on the quantity demanded.

But there are some problems with the eyeballing approach. The most important is that it is highly susceptible to changes in the graph's scale. Graph the demand curve for gasoline in dollars and it will look like Figure 15.1a: buyers appear to be very sensitive to price changes. But graph the demand curve for gasoline in pennies and it will look like Figure 15.1b: buyers will appear to be insensitive to price changes. Similar difficulties affect the x-axis: the curve will look different depending on whether we measure gasoline in gallons or barrels.

15.1 The Price Elasticity of Demand

It turns out that a more solid way of analyzing sensitivity is to put everything in percentages. When we do this, the units become irrelevant: we get the same answer whether we're using dollars or pennies, gallons or liters.

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 (from the price at point A) results in a 5% reduction in quantity demanded (relative to that at point A), we say that the price elasticity of demand at point A is -5. Furthermore,

- We say that demand at point A is **unit elastic** or has **unitary elasticity** 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.

- We say that demand at point A is **elastic** 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 very sensitive to price changes. In math, this means that the elasticity of demand is more negative than -1, e.g., -2: 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 $-\infty$.

- We say that demand at point A is **inelastic** 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 not very sensitive to price changes. In math, this means that the elasticity of demand is more positive than -1, e.g., $-\frac{1}{2}$: a 1% increase in price results in only a $\frac{1}{2}$% decrease in quantity demanded. Of occasional theoretic

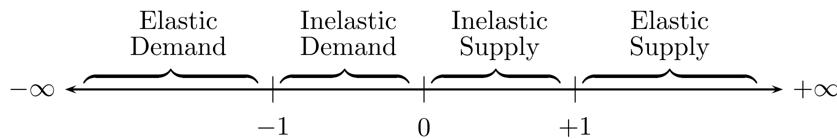

Figure 15.2: Elasticities (with thanks to Anh Nyugen)

interest is the case of completely insensitivity to price changes: demand is **perfectly inelastic** at point A when the elasticity is 0.

Question: What would it mean for the price elasticity of demand to be a positive number such as +1?

Answer: It would mean that a 1% *increase* in price would result in a 1% *increase* in quantity demanded. This means that the demand curve is upward sloping! We never see upward sloping demand curves, so price elasticities of demand are always negative numbers. Because demand elasticities are always negative, economists often talk about elasticities in terms of magnitude or absolute value. When we say that demand is more elastic at point A than at point B, we mean 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 15.2. (Elasticities of supply are discussed in the next section.)

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 them[1]:

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

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!

Application: Elasticities and 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.

To get a better sense of things, imagine that the small increase in the market price amounts to, say, one percent. *If the quantity remained unchanged*, then, total revenue would increase by one percent. But quantity does not remain unchanged. In fact, we have a name for the percentage change in quantity demanded resulting from a one percent increase in price: the price elasticity of demand.

So 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 one percent increase in p will lower q by more than one percent; the combined effect on pq will be to lower total revenue. If demand is inelastic at our original equilibrium, a one percent increase in p will lower q by less than one percent; the combined effect on pq will be to raise total revenue.[2]

Application: Elasticities and Monopoly Behavior

We can also use elasticities to get some results concerning monopoly behavior. Monopolists maximize profits, which are total revenue minus total costs. Although we need calculus to figure out exactly what price a monopolist will charge, we can see here that *monopolists will never choose a price at which demand is inelastic*. To show this, imagine that the monopolist chooses a price at which demand is inelastic. Consider increasing the price by one percent. 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, and we conclude that a monopolist will always increase its price until it reaches the elastic portion of the demand curve.

---

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

[2] It follows from this that total revenue is maximized at the point of unit elasticity, where the elasticity of demand is -1.

15.2 Elasticities of Supply (and Beyond)

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 measures the percentage change in Y resulting from a 1% increase in X. (Again, elasticities are always measured at a point, so the matter at hand is always the X elasticity of Y at some 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, so the matter at hand is always the price elasticity of supply *at some point* A. Unlike demand elasticities, supply curves are positive: a 1% increase in price results in an increase in the quantity supplied. Accordingly, the terms "elastic", "inelastic", and "unit elasticity" refer to the magnitude of the price elasticity of supply relative to +1. Supply at point A is **elastic** if the price elasticity of supply at point A is greater than +1 and **inelastic** if the price elasticity of supply at point A is less than +1; supply at point A has **unit elasticity** if the price elasticity of supply at point A is equal to +1. (See Figure 15.2 for a pictorial summary of supply and demand elasticities.)

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.* ***Long run*** *supply curves can be perfectly elastic at every point.*

Figure 15.3a shows a supply curve that is perfectly inelastic at every point. This 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. Consider, for example, the short run supply curve for 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 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 5,000 apartments in Seattle today, there will be 5,000 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, Figure 15.3b shows a supply curve that is perfectly

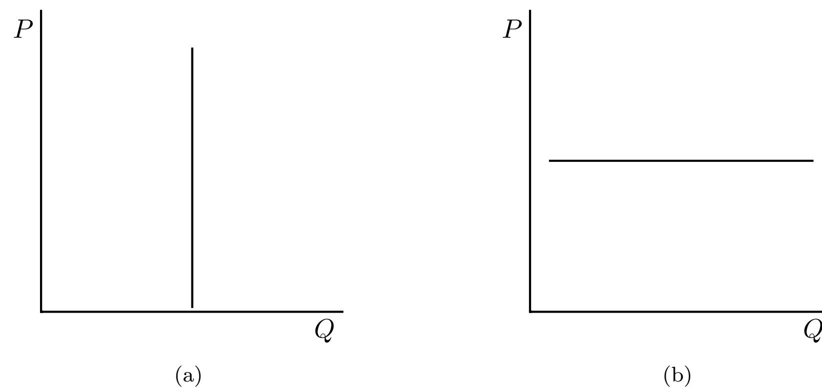

Figure 15.3: (a) A supply curve that is **perfectly inelastic** at every point; (b) a supply curve that is **perfectly elastic** at every point

elastic at every point. This 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 certain long run supply curves. Consider, for example, the long run supply curve for 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 one hundred gallons, one 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 would generate higher rates of return than comparable investments, so everybody would want to sell apple cider; at any price higher than p, then, suppliers *in the long run* would want to sell an infinite amount of apple cider. At any price lower than p, the apple cider business would generate lower rates of return than comparable investments, so nobody would want to sell apple cider; at any price lower than p, then, suppliers *in the long run* would want to sell zero.

The end result is a 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.

Application: Elasticities and Tax Incidence

A final application of elasticities is tax incidence: the ratio of the elasticities of supply and demand matches the ratio of the tax burdens. (See problem 3e.) 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).

Problems

1. Go over previous problems (and/or make up some of your own) and calculate the elasticities of demand and supply at various points.

2. Let's compare short-run elasticities and long-run elasticities.

   (a) Draw a demand curve. Label the axes (p and q) and label this curve "SR". Pick a point on your demand curve (any point!). Assume that the curve you have drawn is a short-run demand curve, and that the point you have drawn is the short-run equilibrium. Now, assume that the market equilibrium in the long run is at the same point. Will the long-run demand curve be flatter or steeper than the short-run curve? Draw in what the long-run demand curve might look like, labeling it "LR". Explain. (Hint: Think about whether buyers' responsiveness to price changes is likely to increase or decrease over time.)

   (b) Draw a supply curve. Label the axes (p and q) and label this curve "SR". Pick a point on your supply curve (any point!). Assume that the curve you have drawn is a short-run supply curve, and that the point you have drawn is the short-run equilibrium. Now, assume that the market equilibrium in the long run is at the same point. Will the long-run supply curve be flatter or steeper than the short-run curve? Draw in what the long-run supply curve might look like, labeling it "LR". Explain. (Hint: Think about whether sellers' responsiveness to price changes is likely to increase or decrease over time.)

3. Figure 15.4 shows a hypothetical demand curve for oranges.

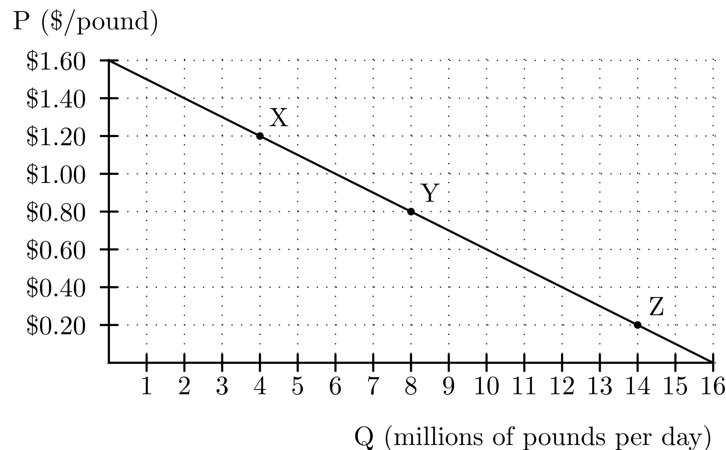

Figure 15.4: A hypothetical market for oranges

(a) Calculate the price elasticity of demand at point Y.

(b) 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?

(c) What is the total revenue at points X, Y, and Z? (Use correct units!)

(d) 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?

(e) This demand curve is the same as in problems 2–8 in Chapter 14. 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 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.

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., $\left(\frac{\epsilon_S}{\epsilon_D}\right) = \left(\frac{S_D}{S_S}\right)$.

5. Very long run supply curves are often assumed to be perfectly elastic.

(a) Explain the intuition behind perfectly elastic long run supply curves. (Hint: Recall that comparable assets should have comparable rates

of return. Think about whether or not a random firm will decide to start producing widgets.)

(b) Draw a graph of a perfectly elastic supply curve.

(c) 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?

Chapter 16

Supply and Demand: Some 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. Our focus will be on two questions. First, 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? Second, 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?

16.1 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. This section deconstructs market supply and demand curves into individual supply and demand curves, and then shows how the cornerstone for all of this analysis is 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? The link to individual optimization is this: we can imagine the firm answering this question by using a decision tree to write down all the different options (sell 4, sell 5, sell 6,. . .) and pick the best one. (Recall decision trees from Chapter 1.) The answer that the firm comes up with at any particular price (e.g., "At a

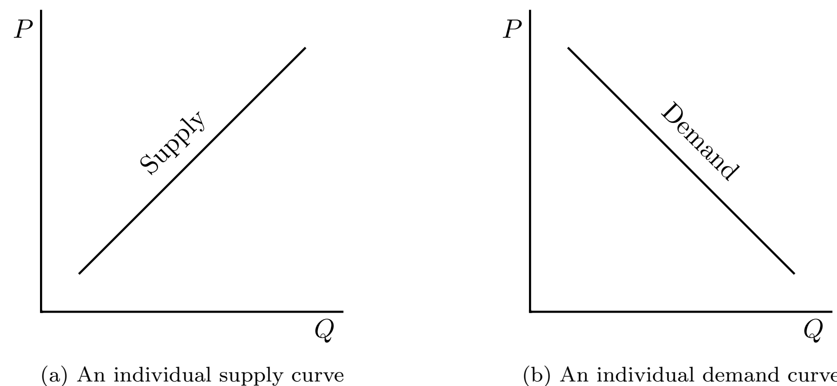

(a) An individual supply curve

(b) An individual demand curve

Figure 16.1: Individual supply and demand curves

price of $6, I would want to sell 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 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 pick 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 16.1) look just like market supply curves (see Figure 13.2c). The only difference is that individual curves are scaled down. All the sellers together may want to sell one million units at a price of $5 per unit, but one individual seller may only want to sell one thousand units at that price. Similarly, all the buyers together may want to buy one million units at a price of $5 per unit, but one individual buyer may only want to buy ten units at that price.

16.2 Reconstructing Supply and Demand

The reconstruction begins 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 (Table 16.2 and Figure 16.2), and over all the individual buyers to get the market demand curve (Table 16.3 and Figure 16.3). Finally, we can put it all together by combining the market supply curve and the market demand curve (Table 16.4 and Figure 16.4).

16.3 *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.

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 is

$$q_M = q_1 + q_2 + \ldots + q_{500} = 500(15 - p) \Longrightarrow q_M = 7500 - 500p.$$

Similarly, 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.

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

$$q_S = q_1 + q_2 + \ldots + q_{400} = 400(15 + 2p) = 6000 + 800p.$$

16.4 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 people want to buy more and that as prices rise people want to buy less. A correct (if unsatisfying) explanation for this assumption is this: 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 another question: why do we assume that *individual* demand curves are downward sloping? There are two possible explanations:

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

Table 16.2: 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

Table 16.3: Aggregating demand

If the price were	0	1	2	3	4	5
Sellers would want to sell	0	3	6	9	12	15
Buyers would want to buy	15	12	9	6	3	0

Table 16.4: Supply and demand together

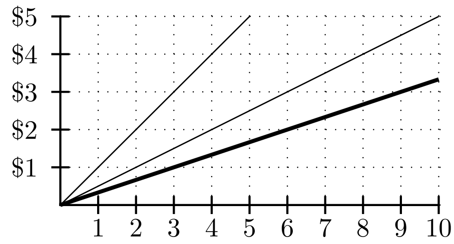

Figure 16.2: Aggregating supply

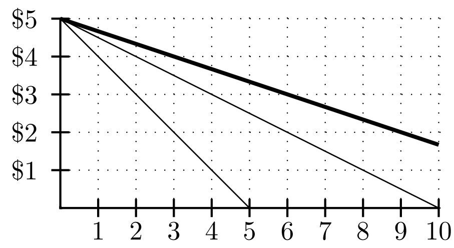

Figure 16.3: Aggregating demand

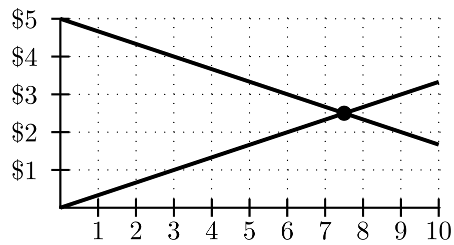

Figure 16.4: Supply and demand together

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. In the lingo: 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.

Obviously, the income effect is not very important when it comes to soft drinks. Indeed, the substitution effect is usually more important than the income effect. 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 *relative to its substitutes* goes up, 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.[1] For practical purposes, then, it is perfectly reasonable to assume—as we will throughout this book—that demand curves are downward sloping.

16.5 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 15.3.) As on the demand

[1] The focus of many such arguments is the Irish potato famine of 1845–49.

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 profits by producing 1,000 units at a price of $1 per unit, it must produce *at least* that many units in order to maximize profits at a price of $2 per unit. (For a mathematical proof of this, see problem 1.)

16.6 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.[2]

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

[2]One 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 15.3b. Demand curves, in contrast, do not exhibit perfect elasticity in the long run.

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 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

1. *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 profits are $pq - C(q)$. To show that downward sloping supply curves are impossible, assume that q^H maximizes profits at market price p^L and then show that a firm facing market price p^H will make higher profits by producing q^H instead of q^L.)

Chapter 17

Margins

Recall that when we draw our familiar supply and demand graph, we put the quantity Q on the x-axis and the price P on the y-axis. This is sort of funny, since 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?" So:

Question: Why don't we graph P on the x-axis and Q on the y-axis?

Answer #1: It was a mistake, a typographical error in the first economics textbook ever written (Alfred Marshall's *Principles of Economics*, originally published in 1890). So it could be that we're locked into this weird way of looking at supply and demand in the same way that physicists are locked into drawing circuit diagrams as if the current is composed of positively-charged particles rather than negatively-charged electrons.

Answer #2 (not necessarily contradictory): Just as we can look at Q as a function of P (i.e., ask, "At a price of 5, how many units will buyers want to buy? How many units will sellers want to sell?"), we can also look at P as a function of Q by asking about **marginal benefit** and **marginal cost**.

17.1 Reinterpreting the Supply Curve

Consider, for example, a firm whose supply curve indicates that (1) at a price of $5 the firm wants to sell 10 units; and (2) at any price below $5 the firm wants to sell less than 10 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 equal to $5. Since the firm wants to sell 10 units at a price of $5, the marginal cost of the 10th unit cannot be greater than $5. (If the marginal cost of the 10th unit was, say, $6, the firm wouldn't be

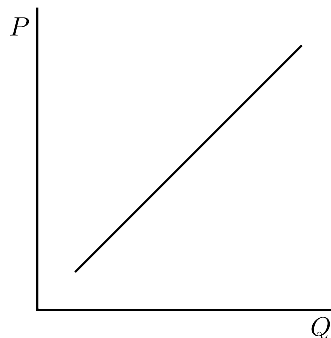

Figure 17.1: Every point on an individual (or market) supply curve is also a point on an individual (or social) marginal cost curve

maximizing profits by selling 10 units at a price of $5; it would make $1 more in profits by selling only 9 units.) And since the firm doesn't want to sell 10 units at a price below $5, the marginal cost of the 10th unit must be at least $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 market prices under $5; for example, with a market price of $4.50, the firm would make $.50 more in profits by selling 10 units instead of 9.) So the firm's marginal cost of producing the 10th unit 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 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 1 million units at a price of $5 but less than 1 million units at any price under $5, 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 17.1. If you start with a given p, an individual (or market) supply curve will tell you how many units that individual supplier (or all the suppliers together) would want to sell at that price; and if you start with a given q, the individual (or social) marginal cost curve will tell you the marginal cost for that individual supplier (or all the

[1] This statement is sufficient for present purposes. As an advanced aside, however, note that while the statement is true, its reverse is not. 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.

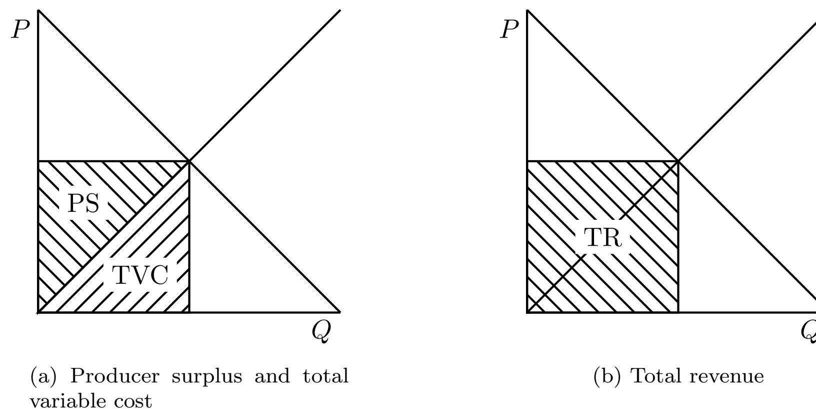

(a) Producer surplus and total variable cost

(b) Total revenue

Figure 17.2: Producer surplus = total revenue − total variable cost

suppliers together) of producing that qth unit.

Example: Taxes Revisited

Let's briefly go back and look at taxes again. Imagine that the government imposes a tax on the sellers of $1 per unit. We know what will happen: the supply curve will shift up by $1. Why? Well, one explanation comes from looking at the market supply curve: at a price of, say, $4, the firms in the market are only going to get to keep $3, 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 another explanation comes from thinking about the supply curve as the marginal cost curve. 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!

17.2 Total Variable Cost and Producer Surplus

If the marginal cost of producing the first unit is $1, that of the second unit is $2, and that of the third unit is $3, then the **total variable cost** of producing three units is $1 + $2 + $3 = $6. Graphically (see Figure 17.2a), the total variable cost can be represented by the area under the supply curve (a.k.a., the marginal cost curve).[2] This area represents all costs except for fixed costs, e.g., the cost of building the factory in the first place. (When you add the fixed costs to the total variable costs, you get total costs.)

If the market price is $3.50 and the firm produces three units, its total revenue will be $\$3.50 \cdot 3 = \$10.50$, and its total variable cost of production will

[2]If you do calculus, this should make good sense: the marginal cost curve is the derivative of the total cost curve, so integrating under the marginal cost curve brings us back to total costs (plus a constant representing fixed costs).

be $6. The difference between total revenue and total variable costs (in this case, $4.50) is **producer surplus**: this is the benefit that the seller gets from selling. Note that producer surplus is not the same thing as profit because we haven't accounted for fixed costs. Once we account for fixed costs, the profit that is left over should be comparable to profits from comparable investments. (See Chapter 5.)

17.3 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 16 concerning downward sloping demand curves, it turns out that there are some theoretical complications.[3] 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 a price of $5 she wants to buy 10 units; and (2) at any price above $5 she wants to buy less than 10 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 equal to $5. Since she wants to buy 10 units at a price of $5, 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 price of $5; she would be better off by buying only 9 units.) And since she doesn't want to buy 10 units at a price above $5, the marginal benefit of the 10th unit must be at most $5. (If the marginal benefit of the 10th unit was, say, $6, she wouldn't be optimizing if she bought only 9 units at market prices above $5; for example, with a market price of $5.50, she'd be better off buying 10 units instead of 9.) So her marginal benefit from obtaining the 10th unit 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 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 a price of $5 but less than 1 million units

[3]Again, the complications stem from the income effect.

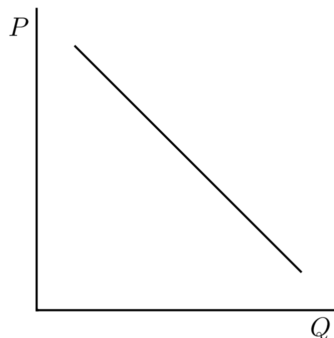

Figure 17.3: Every point on an individual (or market) demand curve is also a point on an individual (or social) marginal benefit curve

at any price over $5, 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 17.3. If you start with a given p, an individual (or market) demand curve will tell you how many units that individual buyer (or all the buyers together) would want to buy at that price; and if you start with a given q, the individual (or social) marginal benefit curve will tell you the marginal benefit for that individual buyer (or all the buyers together) of obtaining that qth unit.

Example: Taxes Revisited

Let's briefly go back and look at taxes again. Imagine that the government imposes a tax on the buyers of $1 per unit. We know what will happen: the demand curve will shift down by $1. Why? Well, one explanation comes from looking at the market demand curve: at a price of, say, $4, 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 another explanation comes from thinking about the demand curve as the marginal benefit curve. The marginal benefit of each additional unit has decreased by $1 because of the tax. So the marginal benefit curve shifts down by $1. Same result, different story!

17.4 Total Benefit and Consumer Surplus

If the marginal benefit of the first unit is $6, that of the second unit is $5, and that of the third unit is $4, then the **total benefit** of having three units is $6 + $5 + $4 = $15. Graphically (see Figure 17.4a), the total benefit can be represented by the area under the demand curve (a.k.a., the marginal benefit

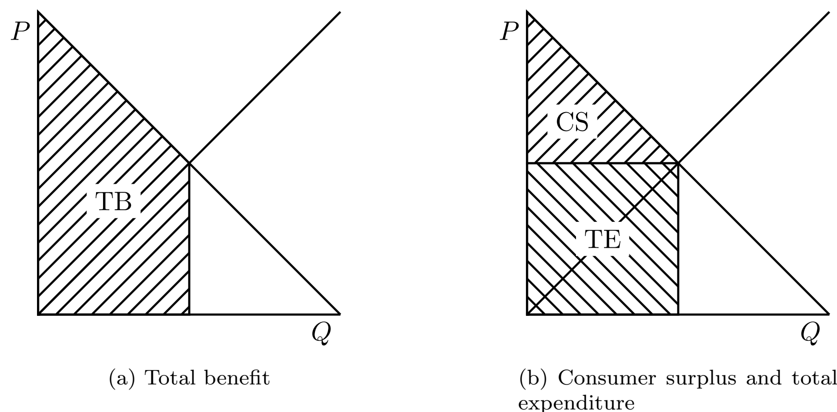

(a) Total benefit

(b) Consumer surplus and total expenditure

Figure 17.4: Consumer surplus = total benefit − total expenditure

curve).[4]

If the market price is \$3.50 and our individual buys three units, her total expenditure will be $\$3.50 \cdot 3 = \$10.50$ and her total benefit will be \$15. The difference between total benefit and total expenditure (in this case, \$4.50) is **consumer surplus**: this is the benefit that the buyer gets from buying.[5]

The Whole Enchilada

Putting supply and demand together, we get Figure 17.5. Total revenue (TR, a.k.a. total expenditure, TE)—is producer surplus (PS) plus total variable cost (TVC). Total benefit (TB) is the area under the demand curve (a.k.a. the marginal benefit curve), so consumer surplus (CS) is $\text{CS} = \text{TB} - \text{TE}$. Total variable cost is the area under the supply curve (a.k.a. the marginal cost curve), so producer surplus is $\text{PS} = \text{TR} - \text{TVC}$. Together producer surplus and consumer surplus constitute the **gains from trade**.

17.5 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). [6] You can think of the marginal benefit

[4] Again, if you do calculus, this should make good sense: the marginal benefit curve is the derivative of the total benefit curve, so integrating under the marginal benefit curve brings us back to total benefits (plus a constant).

[5] Just as producer surplus doesn't account for the seller's fixed costs, consumer surplus doesn't account for the buyer's fixed costs. For example, the consumer surplus you get from travelling around Hawaii doesn't account for the cost of your plane ticket to get there.

[6] To get a donkey to move forward you can either tie a carrot to a pole and dangling 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.

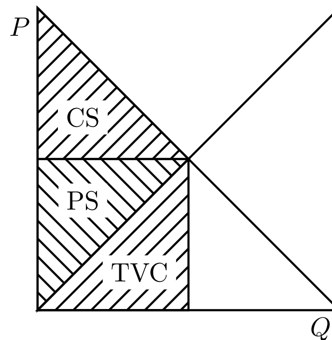

Figure 17.5: The whole enchilada

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.

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 18

Transition: Welfare Economics

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.

18.1 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 reallocation resources via lump sum transfers can make life easy. Unfortunately, the obverse is also true: the inability to reallocation 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.

Appendices &etc.

Appendix A

Government in Practice

Perhaps the most obvious government activities from an economic perspective are the ones more directly related to money: the **fiscal** activities of taxing and spending. Table A.1 provides details on the size of governments—federal, state, and local—relative to the size of the U.S. economy as a whole (as measured by Gross Domestic Product). As with the other tables in the chapter, data are provided in three forms: dollars per person, percentages, and total dollar amounts. For example, the first of these shows that private spending in the U.S. amounted to over $25,000 per person in 2000; the second shows that governments accounted for about 28% of U.S. GDP, with the federal government making up more than half of that amount; and the third shows that the size of the U.S. economy in 2000 was almost $10 trillion.

More information can be found in Tables A.2, A.4, and A.3, which provide overviews and select details of recent annual budgets for the U.S. (federal) government, the state of Washington, and the city of Seattle. Some comments:

- Almost half of federal spending went toward Social Security, Medicare (health care for the elderly), Medicaid (health care for the poor), and other so-called *mandatory* programs.

	Per person	**Percent**	**Total (billions)**
Private spending	$25,300	72%	$7,110
State and local government	4,200	12%	1,180
Federal government	5,630	16%	1,580
Total U.S. GDP	**$35,140**	**100%**	**$9,870**

Table A.1: Breakdown of U.S. Gross Domestic Product (GDP), fiscal year 2000. (Population: 281 million. Items may not add to totals because of rounding.) Source: Office of Management and Budget.

- The category of "non-defense discretionary" encompasses the Environmental Protection Agency, the FBI, the Departments of Agriculture, Commerce, Education, and Transportation, and just about everything else the government does except for the military and mandatory spending.

- Almost one-quarter of Washington state revenue comes from **block grants** and other payments from the federal government. Such government-to-government payments limit the accuracy of Tables A.2–A.4 because of double-counting: money that the federal government gives to the state shows up twice (as a federal expenditure and a state expenditure) even though it is really only spent once.

- The Seattle figures are for the general subfund, which does not include capital funds or special funds. Also excluded are budget figures for the City-operated utilities that provide electricity and water and dispose of sewage and garbage; the budget for these utilities in 2002 totaled $1.6 billion, or $2,860 per person.

A.1 Rules and Regulations

Taxing and spending are only the most obvious economic activities of governments. Another kind of activity is exemplified by minimum wage legislation; such laws have a significant impact on society and on the economy even though they do not entail much direct government spending. Another example is the Clean Air Act, an environmental law first passed in 1970. Among other mandates that affected individuals and private companies, the law required some power plants to put pollution-reducing scrubbers on their smokestacks. A recent analysis by the Environmental Protection Agency found that the costs of the Clean Air Act during its first twenty years (1970–1990) amounted to $523 billion. It also estimated the benefits of that law over the same time period as somewhere between $6 trillion and $50 trillion.

From a broader perspective, governments establish the rules by which society operates. The court system settles disputes and acts to protect rights, including private property rights. The U.S. Treasury prints money, and (as detailed in courses on macroeconomics) the Federal Reserve Board oversees its circulation. And anti-trust laws attempt to prevent mergers or collusion from weakening the forces of competition in economic activity.

U.S. (federal) government, fiscal year 2000

	Per person	**Percent**	**Total (billions)**
MONEY IN			
Personal income taxes	$3,570	49.6%	$1,000
Payroll taxes	2,320	32.2%	650
Corporate income taxes	740	10.2%	210
Excise taxes	250	3.4%	70
Estate and gift taxes	100	1.4%	30
Other revenue	223	3.1%	60
TOTAL REVENUE	**$7,210**	**100.0%**	**$2,030**
MONEY OUT			
“Mandatory” spending, e.g.:	**$3,380**	**46.9%**	**$950**
Social Security	1,440	20.0%	410
Medicare and Medicaid	1,110	15.4%	310
“Discretionary” spending, e.g.:	**2,190**	**30.4%**	**620**
National defense	1,050	14.6%	300
Non-defense discretionary	1,140	15.8%	320
Interest payments	**790**	**11.0%**	**220**
TOTAL EXPENDITURES	**$6,360**	**88.3%**	**$1,790**
BUDGET SURPLUS	**$840**	**11.7%**	**$240**

Table A.2: U.S. federal government revenue and expenditures, fiscal year 2000. (Population: 281 million. Not all detail-level items are shown. Items may not add to totals because of rounding.) Source: Office of Management and Budget.

Washington State, fiscal year 1999

	Per person	Percent	Total (millions)
MONEY IN			
Sales and use taxes, e.g.:	**$1,200**	**36.1%**	**$7,070**
Retail sales and use tax	980	29.5%	5,790
Motor fuels taxes	130	3.9%	760
Alcohol and tobacco taxes	70	2.2%	430
Federal grants	**790**	**23.3%**	**4,570**
Business taxes, e.g.:	**400**	**12.1%**	**2,370**
Business and Occupations tax	310	9.4%	1,850
Property taxes, e.g.:	**300**	**9.0%**	**1,770**
State property tax	220	6.8%	1,330
Vehicle excise tax	60	1.9%	380
Other revenue, e.g.:	**690**	**19.5%**	**3,820**
Workers' comp & unemployment	290	8.6%	1,680
School tuition and fees	94	2.8%	545
Lotto	20	0.6%	110
TOTAL REVENUE	**$3,380**	**100.0%**	**$19,600**
MONEY OUT			
Education, e.g.:	**$1,450**	**43.0%**	**$8,430**
K–12 education	920	27.1%	5,310
Higher education	530	15.7%	3,080
Human services, e.g.:	**1,300**	**38.4%**	**7,530**
Dept. of Social & Health Services	1,040	30.7%	6,020
Dept. of Corrections	80	2.4%	470
Other expenditures, e.g.:	**630**	**18.6%**	**3,650**
Government operations	230	6.8%	1,330
Transportation	130	3.8%	740
Bond retirement and interest	130	3.8%	740
TOTAL EXPENDITURES	**$3,380**	**100.0%**	**$19,600**

Table A.3: Approximate Washington State government receipts and expenditures, fiscal year 1999. (Population: 5.8 million. Not all detail-level items are shown. Items may not add to totals because of rounding.) Source: Washington State Department of Revenue and other.

City of Seattle, fiscal year 2002

	Per person	**Percent**	**Total**
MONEY IN			**(millions)**
Taxes, e.g.:	**$950**	**83.1%**	**$535**
Property taxes	305	26.8%	172
Retail sales taxes	231	20.2%	130
Utilities business taxes	199	17.4%	112
Business and Occupations tax	195	17.1%	110
Other revenue, e.g.:	**171**	**15.0%**	**96**
Service charges	73	6.4%	41
Fines and forfeitures	30	2.6%	17
Government and private grants	21	1.9%	12
Licenses and permits	18	1.6%	10
Parking meter revenue	18	1.6%	10
Fund balance from 2001	**23**	**2.0%**	**13**
TOTAL REVENUE	**$1,142**	**100.0%**	**$643**
MONEY OUT			
Public safety	$557	48.8%	$314
Administration	165	14.5%	93
Arts, culture, and recreation	140	12.3%	79
Utilities and transportation	73	6.4%	41
Health and human services	71	6.2%	40
Neighborhoods/development	48	4.2%	27
Other expenditures	82	7.2%	46
TOTAL EXPENDITURES	**$1,136**	**99.5%**	**$640**
BUDGET SURPLUS	**$5**	**0.5%**	**$3**

Table A.4: City of Seattle general subfund receipts and expenditures, fiscal year 2002. (Population: 560,000. Not all detail-level items are shown. Items may not add to totals because of rounding.) Source: City of Seattle Budget Office.

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 10 for descriptions of different kinds of auctions.

auction A method of selling an item that involves pitting potential buyers against each other. See Chapter 10 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 **prisoner's 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 10 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 10 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 10 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 10 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**)

prisoner's dilemma A simultaneous move game concerning two prisoners who must choose whether or not to accept a plea-bargain that involves testifying against the other prisoner. This game is a classic example of a **collective action problem**.

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 10 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 **Prisoner's Dilemma**. These games can be analyzed with **payoff matrices**. (Compare with **simultaneous 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

Aalsmeer Flower Auction, 79
absolute advantage, 53
annuity, 18, 19, 21, 25
arbitrage, 37–41
  equity premium puzzle, 41
  explorers and sheep, 38–40
  financial, 39–41
  no arbitrage result, 38
  rush-hour, 38–39
auction, 15, 46, 79–88
  all-pay, 84
  code bidding, 85
  common values, 85
  eBay, 83
  equivalences, 82–84
  FCC spectrum, 85
  independent values, 85
  multi-unit, 85
  Revenue Equivalence Theorem, 84
  shading one's bid, 81
  types of, 80–81
  winner's curse, 85

backward induction, 63
Bank
  of America, 22, 24, 41
  of Fish, 22, 24
  of Gold, 24
  of Microsoft, 24
  of Oil, 24
barriers to entry, 97
Bazerman, M.H., 85
benefit
  total, 141
Binmore, Ken, 45
blue
  and yellow makes green, 104
Brams, Steven, 50, 68
business
  game theory applied to, 46, 66

cake cutting algorithms
  applied to divorce and Middle East, 50
cake-cutting algorithms, 45–54
  Adjusted Winner, 50
  books about, 49, 50
  Divide and Choose, 47, 48
    divide or choose?, 48
  Moving Knife, 47, 48
  objectives, 50
  Successive Pairs, 48
  three player envy-free, 49
  with many kids, 48
calculus, 84, 122, 139, 142
Camp David peace accord, 50
capital
  fish are, 22
  theory, 21–24
carbon tax, 119
Cassidy, Butch, 70
climate change, 26, 119
Coase Theorem, 51
collective action problem, 75
comparative advantage, 54
competitive markets, 95
complementary good, 106
compound interest, 26
consumer surplus, 141
Conway, John, 49
cost
  fixed, 139, 140, 142
  opportunity, 5

sunk, 5, 8
total variable, 139
Crampton, Peter, 85

decision tree, 4, 5, 7, 9, 17, 63–65, 129, 130
demand, 101, 129
demand curve
in terms of margins, 140
income effect, 134
increases and decreases, 104
individual, 130
market, 102, 130
substitution effect, 134
differentiation, *see* derivative, calculus
diversification, 11
divorce settlement, 50
Dixit, Avinash K., 68
duopoly, 73

eBay, 79, 83
economic models
Robinson Crusoe, 52
economics
basic assumption of, 3
experimental, 76
elasticity, 119–127
X elasticity of Y, 123
and monopoly behavior, 122
and taxes, 125
and total revenue, 122
elastic, 120
formula, 121
inelastic demand, 120
price elasticity of demand, 120
price elasticity of supply, 123
short-run versus long-run, 125
summed up
in English, 121
in math, 121
in pictures, 121
unit, 120
versus slope, 119
equilibrium
market, 102
equity premium puzzle, 41
evolution
game theory applied to, 46
expected value, 9–15, 27, 38, 39, 87
experimental economics, 76

fair bet, 10, 14
federal government, *see* government
Federal Reserve Board, 27, 152
First Welfare Theorem, 145
fish
are capital, 22
Maximum Sustainable Yield, 21, 22, 24
open-access fishery, 24
slow-growing, 24
free-rider, 75
Friedman, David D., 38
future value, 18, 25

game
subgame, 63
theory, 45
game theory
applications of, 46–47
game tree, 63
games
centipede, 71
centipede game, 77
draft, 68
duopoly, 73
finitely repeated, 76
grenade, 63–65
ice cream pie, 68
investment, 78
pedestrian in Rome, 70
poker, 45
Prisoner's Dilemma, 74–77
Rock, Paper, Scissors, 73
sequential move, 63–71
simultaneous move, 73–78
sticks, 67
surprise exam paradox, 71
global warming, 26, 119
government
block grants, 152

City of Seattle budget, 155
federal budget, 153
fiscal activities, 151
regulatory activities, 152
size of, 151
Washington State budget, 154
green
yellow and blue makes, 104
Guy, Richard, 49

Hall, Monty, 13
hand
invisible, 4, 145, 146
Harriman, E.H., 70

income effect, 134
induction
backward, 63
inferior good, 134
inflation, 31–35
integral, 139, 142
interest
compound, 26
nominal rate of, 31–35
real rate of, 31–35
rule of 72, 25
invisible hand, 4, 145, 146

jokes
$10 bill on the ground, 37
José the bank robber, 79
three wishes, 45
Zen monk, 3

latex (or LaTeX), viii
Linux, vii
liquidity, 40
local government, *see* government
lump sum, 17, 18, 25, 29, 36, 146, 147

macroeconomics, vii, 152
marginal
analysis, 5
benefit, 5, 137, 140–143
cost, 5, 137–143
margins, 137–143
and demand curve, 140
and supply curve, 137
in terms of sticks and carrots, 142
market
competitive, 95
demand curve, 102, 130
equilibrium, 102
market-clearing price, 102
perfectly competitive, 96
price, 101
quantity transacted, 101
supply curve, 102, 129
total expenditure, 101
total revenue, 101
Marshall, Alfred, 137
Maximum Sustainable Yield, 21, 22, 24
mechanism design, 47
microeconomics, vii, 48
monopoly, vii, 6, 59, 73, 96, 97, 122, 160
and elasticity, 122
barriers to entry, 97
government-regulated, 6
in the long run, 96
uniform pricing, 61
Monty Hall Problem, 13

Nalebuff, Barry J., 68
Napster, 51
Newman, Paul, 70
Nobel Prize, 41
Coase, Ronald, 51
King, B.B., 51
normal good, 134

Olson, Mancur, 3
OPEC, 75
open source, vii
opportunity cost, 5

Pareto
efficient, 55–62
improvement, 55–62
inefficient, 55–62

Pareto, Vilfrado, 55
patent race, 84
payoff matrix, 73
perpetuity, 21, 22, 25
poker, 45
present value, 17–19, 22, 24–29, 31, 35, 36
  of annuity, 19
  of Louisiana Purchase, 28, 36
  of perpetuity, 21
price theory, 95
price-taker, 96
Prisoner's Dilemma, 74–77
  finitely repeated, 76
producer surplus, 139
profit, 4
  accounting, 4
  economic, 4

quantum, vii, 162

Raiffa
  Howard, 14
Ramen noodles, 134
Raymond, Eric S., viii
Redford, Robert, 70
repeated game, 76
risk, 4, 9, 14, 40, 41
  premium, 40, 41
    equity premium puzzle, 41
  reducing with diversification, 11
Robertson, Jack, 49
Robinson Crusoe model, 52
Rock, Paper, Scissors, 73
rule of 72, 25

Samuelson, W.F., 85
Schwartz, Jesse A., 85
Second Welfare Theorem, 146
Selfridge, John, 49
Shapiro, Carl, 60
slope, 117
  and elasticity, 126
  and tax burdens, 118, 126
  of supply/demand curve, 117, 119, 126
  versus elasticity, 119
Smith
  Adam, 145
sports
  game theory applied to, 46, 68
stage game, 76
state government, *see* government
subgame, 63
substitute, 106
substitution effect, 134
supply, 101, 129
supply curve
  in terms of margins, 137
  increases and decreases, 104
  individual, 129
  market, 102, 129
surplus
  consumer, 141
  producer, 139

tax
  carbon, 119
  equivalence, 114–115
  per-unit, 111–114
taxes, 111–118
  and elasticity, 125
  in terms of margins, 139, 141
Taylor, Alan, 50, 68
total benefit, 141
tree
  game, 63
Trump, Donald, 50

uncertainty, 40
unfair bet, 10

Varian, Hal, 60
variance, 10, 12
von Neumann, John, 46

war
  game theory applied to, 46
Webb, William, 49
welfare economics, 145
  First Welfare Theorem, 145
  policy implications, 147

Second Welfare Theorem, 146
winner's curse, 85
Wolfstetter, Elmar, 79

yellow
and blue makes green, 104