

Quantum Microeconomics *with Calculus*

Version 3.2

January 2005

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	5
1.2 Example: Monopoly	6
1.3 <i>Math</i> : Optimization and Calculus	7
2 Optimization and Risk	19
2.1 Reducing Risk with Diversification	21
2.2 <i>Math</i> : Indifference Curves	23
3 Optimization over Time	29
3.1 Lump Sums	30
3.2 Annuities	31
3.3 Perpetuities	33
3.4 Capital Theory	34
3.5 <i>Math</i> : Present Value and Budget Constraints	37
4 <i>Math</i>: Trees and Fish	45
4.1 Trees	45
4.2 Fish	48
5 More Optimization over Time	55
5.1 Nominal and Real Interest Rates	55
5.2 Inflation	56
5.3 Mathematics	58
6 Transition: Arbitrage	61
6.1 No Arbitrage	61
6.2 Rush-Hour Arbitrage	62
6.3 Financial Arbitrage	63

II	One v. One, One v. Many	67
7	Cake-Cutting	69
7.1	Some Applications of Game Theory	69
7.2	Cake-Cutting: The Problem of Fair Division	71
7.3	The Importance of Trade	74
8	Economics and Social Welfare	79
8.1	Cost-Benefit Analysis	80
8.2	Pareto	81
8.3	Examples	82
9	Sequential Move Games	89
9.1	Backward Induction	89
10	Simultaneous Move Games	99
10.1	Dominance	100
10.2	The Prisoners' Dilemma	100
10.3	Finitely Repeated Games	102
11	Iterated Dominance and Nash Equilibrium	105
11.1	Iterated Dominance	105
11.2	Nash Equilibrium	108
11.3	Infinitely Repeated Games	110
11.4	Mixed Strategies	111
11.5	<i>Math</i> : Mixed Strategies	112
12	Application: Auctions	121
12.1	Kinds of Auctions	122
12.2	Auction Equivalences	124
12.3	Auction Miscellany	126
13	Application: Marine Affairs	131
13.1	An Economic Perspective	132
13.2	A Brief Look at Government Intervention	132
13.3	ITQs to the Rescue?	135
14	Transition: Game Theory v. Price Theory	137
14.1	Monopolies in the Long Run	138
14.2	Barriers to Entry	139
III	Many v. Many	141
15	Supply and Demand Basics	143
15.1	The Story of Supply and Demand	144
15.2	Shifts in Supply and Demand	146

15.3 <i>Math</i> : The Algebra of Markets	147
16 Taxes	153
16.1 A Per-Unit Tax Levied on the Sellers	153
16.2 A Per-Unit Tax Levied on the Buyers	154
16.3 Tax Equivalence	156
16.4 <i>Math</i> : The Algebra of Taxes	158
17 Elasticities	163
17.1 The Price Elasticity of Demand	164
17.2 Elasticities of Supply (and Beyond)	167
17.3 <i>Math</i> : Elasticities and Calculus	169
18 Supply and Demand Details	173
18.1 Deconstructing Supply and Demand	173
18.2 Reconstructing Supply and Demand	174
18.3 <i>Math</i> : The Algebra of Markets	175
18.4 On the Shape of the Demand Curve	175
18.5 On the Shape of the Supply Curve	178
18.6 Comparing Supply and Demand	179
19 Margins	183
19.1 Reinterpreting the Supply Curve	183
19.2 Total Variable Cost and Producer Surplus	185
19.3 Reinterpreting the Demand Curve	186
19.4 Total Benefit and Consumer Surplus	187
19.5 Conclusion: Carrots and Sticks	188
20 <i>Math</i>: Deriving Supply and Demand Curves	191
20.1 Cost Minimization	191
20.2 Supply Curves	199
20.3 Demand Curves	201
21 Transition: Welfare Economics	213
21.1 From Theory to Reality	215
Appendices &etc.	217
A Government in Practice	219
A.1 Rules and Regulations	220
B <i>Math</i>: Monopoly and Oligopoly	225
B.1 Monopoly: The Vanilla Version	225
B.2 Monopoly: The Chocolate-with-Sprinkles Version	226
B.3 Duopoly	228
B.4 The Transition to Perfect Competition	233

C Answer Key for Part I	237
D Answer Key for Part II	251
E Answer Key for Part III	265
F Answer Key for Appendixes	281
Glossary	287
Index	294

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. 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 L^AT_EX 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 L^AT_EX, 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-L^AT_EX format.

¹<http://www.opensource.org>

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>

⁶Another 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 and/or chronological) do not imply endorsement.

Yoram Bauman
 yoram@smallparty.org
 Walla Walla, Washington
 January 2005

Melissa Andersen	Aaron Finkle	Noelwah Netusil
Kea Asato	Jason Gasper	Anh Nguyen
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
Kathryn Bergh		

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 10. For an optimistic view, see Chapter 21 or consider the words of Adam Smith, who wrote in 1776 that “[man is] led by an invisible hand to promote an end which was no part of his intention. . . . By pursuing his own interest he frequently promotes that of society more effectually than when he really intends to promote it.”)

Optimization

Economics assumes that individuals try to do the best they can. Although economics is unwavering in the assumption that individuals are optimizing—i.e., that each has some objective—there *is* flexibility in determining exactly what those objectives are. In particular, economics does not need to assume that individuals are selfish or greedy; their objectives may well 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.

An aside about individual firms Economists often treat companies as optimizing individuals whose goal is *profit-maximization*. (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. While some branches of economics do analyze the goings-on inside firms, in many cases it is valuable to simplify matters by assuming—as we will throughout this book—that firms act like optimizing individuals and that their objective is to maximize profit.

An aside about individual people Economists say that optimizing individuals pursue *utility maximization*. Although it is not incorrect to think of **utility** as simply a fancy word for happiness, some complex philosophical and mathematical ideas lie under the surface. At the heart of modern microeconomic theory is the idea of *preferences*: given two options, an individual can say that she prefers one over the other (or that she likes them both equally). If we make certain assumptions (e.g., that if the individual prefers A over B and prefers B over C then she prefers A over C), it turns out that we can represent

¹The 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

individual preferences using a **utility function** that assigns a number (a “happiness level”) to each option. Economists around the turn of the 20th century thought that utility might actually be real, something that could be measured using “hedonmeters” or “psychogalvanometers”. In contrast, modern economic theory treats utility as simply a handy mathematical technique for representing an individual’s preferences: saying that option A produces a utility of 20 and option B produces a utility of 10 means no more (and no less) than saying that the individual prefers A over B.²

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 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

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

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.³ 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

³For this reason, monopolists are also known as **price-setting** firms, whereas firms in a competitive market are known as **price-taking** firms.

lazy doesn't find much support in this model.⁴

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.

1.3 Math: Optimization and Calculus

Derivatives in theory

The **derivative** of a function $f(x)$, written $\frac{d}{dx}[f(x)]$ or $\frac{df(x)}{dx}$ or $f'(x)$, measures the *instantaneous rate of change* of $f(x)$:

$$\frac{d}{dx}[f(x)] = \lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h}.$$

Intuitively, derivatives measures **slope**: $f'(x) = -3$ intuitively means that if x increases by 1 then $f(x)$ will decrease by 3. This intuition matches up with setting $h = 1$, which yields

$$f'(x) \approx \frac{f(x+1) - f(x)}{1} = f(x+1) - f(x).$$

All of the functions we use in this class have derivatives (i.e., are **differentiable**), which intuitively means that they are smooth and don't have kinks or discontinuities. The maximum and minimum values of such functions must either be **corner solutions**—such as $x = \infty$, $x = -\infty$, or (if we are trying to maximize $f(x)$ subject to $x \geq x_{min}$) $x = x_{min}$ —or **interior solutions**. The vast majority of the problems in this class will have interior solutions.

At an interior maximum or minimum, the slope $f'(x)$ must be zero. Why? Well, if $f'(x) \neq 0$ then either $f'(x) > 0$ or $f'(x) < 0$. Intuitively, this means that you're on the side of a hill—if $f'(x) > 0$ you're going uphill, if $f'(x) < 0$ you're heading downhill—and if you're on the side of a hill then you're not at the top (a maximum) or at the bottom (a minimum). At the top and the bottom the slope $f'(x)$ will be zero.

⁴More 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 shows the monopolist's total revenue if it chooses the given price; note that the highest price (\$6 per unit) yields *less* revenue than lower prices! The fourth line shows the monopolist's costs for producing the amount that buyers want to buy at the given price; 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.

To say the same thing in math: If $f'(x) \neq 0$ then either $f'(x) > 0$ or $f'(x) < 0$. If $f'(x) > 0$ then $f(x+h) > f(x)$ (this comes from the definition of derivative), so $f(x)$ isn't a maximum; and $f(x-h) < f(x)$ (this follows from continuity, i.e., the fact that $f(x)$ is a smooth function), so $f(x)$ isn't a minimum. Similarly, if $f'(x) < 0$ then $f(x+h) < f(x)$ (this comes from the definition of derivative), so $f(x)$ isn't a minimum; and $f(x-h) > f(x)$ (this follows from continuity), so $f(x)$ isn't a maximum. So the only possible (interior) maxima or minima must satisfy $f'(x) = 0$, which is called a **necessary first-order condition**.

In sum: to find candidate values for (interior) maxima or minima, simply take a derivative and set it equal to zero, i.e., find values of x that satisfy $f'(x) = 0$.

Such values do not *have* to be maxima or minima: the condition $f'(x) = 0$ is *necessary* but not *sufficient*. This is a more advanced topic that we will not get into in this course, but for an example consider $f(x) = x^3$. Setting the derivative ($3x^2$) equal to zero has only one solution: $x = 0$. But $x = 0$ is neither a minimum nor a maximum value of $f(x) = x^3$. The **sufficient second-order condition** has to do with the second derivative (i.e., the derivative of the derivative, written $f''(x)$). For a maximum, the sufficient second-order condition is $f''(x) < 0$; this guarantees that we're on a hill, so together with $f'(x) = 0$ it guarantees that we're on the top of the hill. For a minimum, the sufficient second-order condition is $f''(x) > 0$; this guarantees that we're in a valley, so together with $f'(x) = 0$ it guarantees that we're at the bottom of the valley.)

Partial derivatives

For functions of two or more variables such as $f(x, y)$, it is often useful to see what happens when we change one variable (say, x) without changing the other variables. (For example, what happens if we walk in the north-south direction without changing our east-west position?) What we end up with is the *partial derivative with respect to x* of the function $f(x, y)$, written $\frac{\partial}{\partial x} [f(x, y)]$ or $f_x(x, y)$:

$$\frac{\partial}{\partial x} [f(x, y)] = \lim_{h \rightarrow 0} \frac{f(x+h, y) - f(x, y)}{h}.$$

Partial derivatives measure rates of change or slopes *in a given direction*: $f_x(x, y) = 3y$ intuitively means that if x increases by 1 and y doesn't change then $f(x, y)$ will increase by $3y$. Note that "regular" derivatives and partial derivatives mean the same thing for a function of only one variable: $\frac{d}{dx} [f(x)] = \frac{\partial}{\partial x} [f(x)]$.

At an (interior) maximum or minimum of a smooth function, the slope must be zero *in all directions*. In other words, the necessary first-order conditions are that *all* partials must be zero: $f_x(x, y) = 0$, $f_y(x, y) = 0$, etc. Why? For the same reasons we gave before: if one of the partials—say, $f_y(x, y)$ —is not zero, then moving in the y direction takes us up or down the side of a hill, and so we cannot be at a maximum or minimum value of the function $f(x, y)$.

In sum: to find candidate values for (interior) maxima or minima, simply take partial derivatives with respect to all the variables and set them equal to zero, e.g., find values of x and y that simultaneously satisfy $f_x(x, y) = 0$ and $f_y(x, y) = 0$.

As before, these conditions are necessary but not sufficient. This is an even more advanced topic than before, and we will not get into it in this course; all I will tell you here is that (1) the sufficient conditions for a maximum include $f_{xx} < 0$ and $f_{yy} < 0$, *but these aren't enough*, (2) you can find the sufficient conditions in most advanced textbooks, e.g., Silberberg and Suen's *The Structure of Economics*, and (3) an interesting example to consider is $f(x, y) = \frac{\cos(x)}{\cos(y)}$ around the point $(0, 0)$.

One final point: Single variable derivatives can be thought of as a degenerate case of partial derivatives: there is no reason we can't write $f_x(x)$ instead of $f'(x)$ or $\frac{\partial}{\partial x}f(x)$ instead of $\frac{d}{dx}f(x)$. All of these terms measure the same thing: the rate of change of the function $f(x)$ in the x direction.

Derivatives in practice

To see how to calculate derivatives, let's start out with a very simple function: the constant function $f(x) = c$, e.g., $f(x) = 2$. We can calculate the derivative of this function from the definition:

$$\frac{d}{dx}(c) = \lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h} = \lim_{h \rightarrow 0} \frac{c - c}{h} = 0.$$

So the derivative of $f(x) = c$ is $\frac{d}{dx}(c) = 0$. Note that *all* values of x are candidate values for maxima and/or minima. Can you see why?⁵

Another simple function is $f(x) = x$. Again, we can calculate the derivative from the definition:

$$\frac{d}{dx}(x) = \lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h} = \lim_{h \rightarrow 0} \frac{(x+h) - x}{h} = \lim_{h \rightarrow 0} \frac{h}{h} = 1.$$

So the derivative of $f(x) = x$ is $\frac{d}{dx}(x) = 1$. Note that *no* values of the function $f(x) = x$ are candidate values for maxima or minima. Can you see why?⁶

A final simple derivative involves the function $g(x) = c \cdot f(x)$ where c is a constant and $f(x)$ is any function:

$$\frac{d}{dx}[c \cdot f(x)] = \lim_{h \rightarrow 0} \frac{c \cdot f(x+h) - c \cdot f(x)}{h} = c \cdot \lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h}.$$

The last term on the right hand side is simply the derivative of $f(x)$, so the derivative of $g(x) = c \cdot f(x)$ is $\frac{d}{dx}[c \cdot f(x)] = c \cdot \frac{d}{dx}[f(x)]$.

⁵Answer: All values of x are maxima; all values are minima, too! Any x you pick gives you $f(x) = c$, which is both the best and the worst you can get.

⁶Answer: There are no interior maxima or minima of the function $f(x) = x$.

More complicated derivatives

To differentiate (i.e., calculate the derivative of) a more complicated function, use various differentiation rules to methodically break down your problem until you get an expression involving the derivatives of the simple functions shown above.

The most common rules are those involving the three main binary operations: addition, multiplication, and exponentiation.

- **Addition** $\frac{d}{dx} [f(x) + g(x)] = \frac{d}{dx} [f(x)] + \frac{d}{dx} [g(x)].$

Example: $\frac{d}{dx} [x + 2] = \frac{d}{dx} [x] + \frac{d}{dx} [2] = 1 + 0 = 1.$

Example: $\frac{d}{dx} [3x^2(x + 2) + 2x] = \frac{d}{dx} [3x^2(x + 2)] + \frac{d}{dx} [2x].$

- **Multiplication** $\frac{d}{dx} [f(x) \cdot g(x)] = f(x) \cdot \frac{d}{dx} [g(x)] + g(x) \cdot \frac{d}{dx} [f(x)].$

Example: $\frac{d}{dx} [3x] = 3 \cdot \frac{d}{dx} [x] + x \cdot \frac{d}{dx} [3] = 3(1) + x(0) = 3.$

(Note: this also follows from the result that $\frac{d}{dx} [c \cdot f(x)] = c \cdot \frac{d}{dx} [f(x)].$)

Example: $\frac{d}{dx} [x(x + 2)] = x \cdot \frac{d}{dx} [(x + 2)] + (x + 2) \cdot \frac{d}{dx} [x] = 2x + 2.$

Example: $\frac{d}{dx} [3x^2(x + 2)] = 3x^2 \cdot \frac{d}{dx} [(x + 2)] + (x + 2) \cdot \frac{d}{dx} [3x^2].$

- **Exponentiation** $\frac{d}{dx} [f(x)^a] = a \cdot f(x)^{a-1} \cdot \frac{d}{dx} [f(x)].$

Example: $\frac{d}{dx} [(x + 2)^2] = 2(x + 2)^1 \cdot \frac{d}{dx} [x + 2] = 2(x + 2)(1) = 2(x + 2).$

Example: $\frac{d}{dx} [(2x + 2)^{\frac{1}{2}}] = \frac{1}{2}(2x + 2)^{-\frac{1}{2}} \cdot \frac{d}{dx} [2x + 2] = (2x + 2)^{-\frac{1}{2}}.$

Putting all these together, we can calculate lots of messy derivatives:

$$\begin{aligned} \frac{d}{dx} [3x^2(x + 2) + 2x] &= \frac{d}{dx} [3x^2(x + 2)] + \frac{d}{dx} [2x] \\ &= 3x^2 \cdot \frac{d}{dx} [x + 2] + (x + 2) \cdot \frac{d}{dx} [3x^2] + \frac{d}{dx} [2x] \\ &= 3x^2(1) + (x + 2)(6x) + 2 \\ &= 9x^2 + 12x + 2 \end{aligned}$$

Subtraction and division

The rule for addition also works for subtraction, and can be seen by treating $f(x) - g(x)$ as $f(x) + (-1) \cdot g(x)$ and using the rules for addition and multiplication. Less obviously, the rule for multiplication takes care of division: $\frac{d}{dx} \left[\frac{f(x)}{g(x)} \right] = \frac{d}{dx} [f(x) \cdot g(x)^{-1}]$. Applying the product and exponentiation rules to this yields the quotient rule,⁷

$$\bullet \text{ Division } \quad \frac{d}{dx} \left[\frac{f(x)}{g(x)} \right] = \frac{g(x) \cdot \frac{d}{dx} f(x) - f(x) \cdot \frac{d}{dx} g(x)}{[g(x)]^2}.$$

$$\text{Example: } \frac{d}{dx} \left[\frac{3x^2 + 2}{-e^x} \right] = \frac{-e^2 \cdot \frac{d}{dx} [3x^2 + 2] - (3x^2 + 2) \cdot \frac{d}{dx} [-e^x]}{[-e^x]^2}.$$

Exponents

If you're confused about what's going on with the quotient rule, you may find value in the following rules about exponents, which we will use frequently:

$$x^a \cdot x^b = x^{a+b} \quad (x^a)^b = x^{ab} \quad x^{-a} = \frac{1}{x^a} \quad (xy)^a = x^a y^a$$

Examples: $2^2 \cdot 2^3 = 2^5$, $(2^2)^3 = 2^6$, $2^{-2} = \frac{1}{4}$, $(2 \cdot 3)^2 = 2^2 \cdot 3^2$.

Other differentiation rules: e^x and $\ln(x)$

You won't need the chain rule, but you may need the rules for derivatives involving the exponential function e^x and the natural logarithm function $\ln(x)$. (Recall that e and \ln are inverses of each other, so that $e^{(\ln x)} = \ln(e^x) = x$.)

$$\bullet \text{ The exponential function } \quad \frac{d}{dx} [e^{f(x)}] = e^{f(x)} \cdot \frac{d}{dx} [f(x)].$$

$$\text{Example: } \frac{d}{dx} [e^x] = e^x \cdot \frac{d}{dx} [x] = e^x.$$

$$\text{Example: } \frac{d}{dx} [e^{3x^2+2}] = e^{3x^2+2} \cdot \frac{d}{dx} [3x^2 + 2] = e^{3x^2+2} \cdot (6x).$$

$$\bullet \text{ The natural logarithm function } \quad \frac{d}{dx} [\ln f(x)] = \frac{1}{f(x)} \cdot \frac{d}{dx} [f(x)].$$

$$\text{Example: } \frac{d}{dx} [\ln x] = \frac{1}{x} \cdot \frac{d}{dx} [x] = \frac{1}{x}.$$

$$\text{Example: } \frac{d}{dx} [\ln(3x^2 + 2)] = \frac{1}{3x^2 + 2} \cdot \frac{d}{dx} [3x^2 + 2] = \frac{1}{3x^2 + 2} (6x).$$

⁷Popularly remembered as $d \left[\frac{\text{Hi}}{\text{Ho}} \right] = \frac{\text{Ho} \cdot d\text{Hi} - \text{Hi} \cdot d\text{Ho}}{\text{Ho} \cdot \text{Ho}}.$

Partial derivatives

Calculating partial derivatives (say, with respect to x) is easy: just treat all the other variables as constants while applying all of the rules from above. So, for example,

$$\begin{aligned}\frac{\partial}{\partial x} [3x^2y + 2e^{xy} - 2y] &= \frac{\partial}{\partial x} [3x^2y] + \frac{\partial}{\partial x} [2e^{xy}] - \frac{\partial}{\partial x} [2y] \\ &= 3y \frac{\partial}{\partial x} [x^2] + 2e^{xy} \frac{\partial}{\partial x} [xy] - 0 \\ &= 6xy + 2ye^{xy}.\end{aligned}$$

Note that the partial derivative $f_x(x, y)$ is a function of both x and y . This simply says that the rate of change with respect to x of the function $f(x, y)$ depends on where you are in both the x direction and the y direction.

We can also take a partial derivative with respect to y of the same function:

$$\begin{aligned}\frac{\partial}{\partial y} [3x^2y + 2e^{xy} - 2y] &= \frac{\partial}{\partial y} [3x^2y] + \frac{\partial}{\partial y} [2e^{xy}] - \frac{\partial}{\partial y} [2y] \\ &= 3x^2 \frac{\partial}{\partial y} [y] + 2e^{xy} \frac{\partial}{\partial y} [xy] - 2 \\ &= 3x^2 + 2xe^{xy} - 2.\end{aligned}$$

Again, this partial derivative is a function of both x and y .

Integration

The integral of a function, written $\int_a^b f(x) dx$, measures the *area* under the function $f(x)$ between the points a and b . We won't use integrals much, but they are related to derivatives by the Fundamental Theorem(s) of Calculus:

$$\int_a^b \frac{d}{dx} [f(x)] dx = f(b) - f(a) \qquad \frac{d}{ds} \left[\int_a^s f(x) dx \right] = f(s)$$

Example: $\int_0^1 x dx = \int_0^1 \frac{d}{dx} \left[\frac{1}{2}x^2 \right] dx = \frac{1}{2}(1^2) - \frac{1}{2}(0^2) = \frac{1}{2}$

Example: $\frac{d}{ds} \left[\int_0^s x dx \right] = s.$

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 \$1 billion it spent on research and development (R&D).
 - (a) Will a profit-maximizing firm pay attention to R&D costs when determining its pricing?
 - (b) If you said "Yes" above: Do you think the company would have charged less for the drug if it had discovered it after spending only \$5 million instead of \$1 billion?
If you said "No" above: Do R&D costs affect the company's behavior *before* they decide whether or not to invest in the R&D, *after* they invest in the R&D, both before and after, or neither?

Calculus Problems

- C-1. Explain the importance of taking derivatives and setting them equal to zero.
- C-2. Use the definition of a derivative to prove that constants pass through derivatives, i.e., that $\frac{d}{dx}[(c \cdot f(x))] = c \cdot \frac{d}{dx}[f'(x)]$.
- C-3. Use the product rule to prove that the derivative of x^2 is $2x$. (*Challenge:* Do the same for higher-order integer powers, e.g., x^{30} . *Do not* do this the hard way.)
- C-4. Use the product and exponent rules to derive the quotient rule.

C-5. For each of the following functions, calculate the first derivative, the second derivative, and determine maximum and/or minimum values (if they exist):

- (a) $x^2 + 2$
- (b) $(x^2 + 2)^2$
- (c) $(x^2 + 2)^{\frac{1}{2}}$
- (d) $-x(x^2 + 2)^{\frac{1}{2}}$
- (e) $\ln \left[(x^2 + 2)^{\frac{1}{2}} \right]$

C-6. Calculate partial derivatives with respect to x and y of the following functions:

- (a) $x^2y - 3x + 2y$
- (b) e^{xy}
- (c) $e^x y^2 - 2y$

C-7. Consider a market with demand curve $q = 20 - 2p$.

- (a) Calculate the slope of the demand curve, i.e., $\frac{dq}{dp}$. Is it positive or negative? Does the sign of the demand curve match your intuition, i.e., does it make sense?
- (b) Solve the demand curve to get p as a function of q . This is called an **inverse demand curve**.
- (c) Calculate the slope of the inverse demand curve, i.e., $\frac{dp}{dq}$. Is it related to the slope of the demand curve?

C-8. Imagine that a monopolist is considering entering a market with demand curve $q = 20 - p$. Building a factory will cost F , and producing each unit will cost 2 so its profit function (if it decides to enter) is $\pi = pq - 2q - F$.

- (a) Substitute for p using the inverse demand curve and find the (interior) profit-maximizing level of output for the monopolist. Find the profit-maximizing price and the profit-maximizing profit level.
- (b) For what values of F will the monopolist choose not to enter the market?

C-9. (Profit maximization for a firm in a competitive market) Profit is $\pi = p \cdot q - C(q)$. If the firm is maximizing profits and takes p as given, find the necessary first order condition for an interior solution to this problem, both in general and in the case where $C(q) = \frac{1}{2}q^2 + 2q$.

- C-10. (Profit maximization for a non-price-discriminating monopolist) A monopolist can choose both price and quantity, but choosing one essentially determines the other because of the constraint of the market demand curve: if you choose price, the market demand curve tells you how many units you can sell at that price; if you choose quantity, the market demand curve tells you the maximum price you can charge while still selling everything you produce. So: if the monopolist is profit-maximizing, find the necessary first order condition for an interior solution to the monopolist's problem, both in general and in the case where the demand curve is $q = 20 - p$ and the monopolist's costs are $C(q) = \frac{1}{2}q^2 + 2q$.

- C-11. (Derivation of Marshallian demand curves) Considered an individual whose preferences can be represented by the utility function

$$U = \left(\frac{100 - p_Z \cdot Z}{p_B} \cdot Z \right)^{\frac{1}{2}}.$$

If the individual chooses Z to maximize utility, find the necessary first order condition for an interior solution to this problem. Simplify to get the Marshallian demand curve $Z = \frac{50}{p_Z}$.

- C-12. *Challenge.* (Stackleberg leader-follower duopoly) Consider a model with two profit-maximizing firms, which respectively produce q_1 and q_2 units of some good. The inverse demand curve is given by $p = 20 - Q = 20 - (q_1 + q_2)$. Both firms have costs of $C(q) = 2q$. In this market, firm 1 gets to move first, so the game works like this: firm 1 chooses q_1 , then firm 2 chooses q_2 , then both firms sell their output at the market price $p = 20 - (q_1 + q_2)$.

- Write down the maximization problem for Firm 1. What is/are its choice variable(s)?
- Write down the maximization problem for Firm 2. What is/are its choice variable(s)?
- What's different between the two maximization problems? (Hint: Think about the timing of the game!) *Think hard about this question before going on!*
- Take the partial derivative (with respect to q_2) of Firm 2's objective function (i.e., the thing Firm 2 is trying to maximize) to find the necessary first order condition for an interior solution to Firm 2's problem.
- Your answer to the previous problem should show how Firm 2's choice of q_2 depends on Firm 1's choice of q_1 , i.e., it should be a **best response function**. Explain why this function is important to Firm 1.

- (f) Plug the best response function into Firm 1's objective function and find the necessary first order condition for an interior solution to Firm 1's problem.
- (g) Solve this problem for q_1 , q_2 , p , and the profit levels of the two firms. Is there a **first mover advantage** or a **second mover advantage**? Can you intuitively understand why?

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




Figure 2.1: A decision tree involving risk

concept of expected value provides a valuable perspective by condensing this risky situation into a single number. Note, however, that it conveys only one perspective. 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

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

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

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

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.

¹The 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 [online²](#) and [online³](#).) 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 [online⁴](#) 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>

³http://www.ruf.rice.edu/~lane/stat_sim/binom_demo.html

⁴http://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 6

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](http://cowles.econ.yale.edu/archive/people/tobin/nyt_obit.htm)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’.”

2.2 Math: Indifference Curves

A more mathematical treatment of this material allows us introduce **indifference curves**: we can ask when an individual is indifferent (i.e., has no preference) between one risky situation and another, or between a risky situation and a riskless one. For example, consider a coin-flipping bet in which someone wins \$100 for heads and \$0 for tails. Presumably this person would be indifferent between this bet and a symmetric bet in which he wins \$100 for tails and \$0 for heads. (If we refer to the first bet using the notation $(100, 0)$, then the notation $(0, 100)$ describes the second bet.) In fact, there are probably an infinite number of pairs (x, y) such that he’d be indifferent between the bet $(100, 0)$ and the bet (x, y) , i.e., a bet in which he wins \$ x for heads and \$ y for tails. If we graph all these pairs we get an indifference curve like the one in Figure 2.2. Using the utility-maximization idea described in Chapter 1, we can say that all the points on this indifference curve have the same utility.

Of particular interest on this indifference curve is the point where $x = y$, because at this point the individual is not exposed to any uncertainty: regardless of the outcome of the coin flip, he gets the same amount. This amount is called the **certainty equivalent wealth** corresponding to that indifference curve. For the curve in Figure 2.2, the certainty equivalent wealth is \$37, meaning that our individual is indifferent between getting \$37 *with certainty* and facing the uncertain bets—such as $(100, 0)$ and $(0, 100)$ —that are on the same indifference curve.

Note that the expected value of the bet $(100, 0)$ is \$50, which is *more* than the individual’s certainty equivalent wealth. This implies that he is risk-averse; a risk-neutral person would have a certainty equivalent wealth of \$50 for the

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




Figure 2.2: An indifference curve for a coin-flipping bet

same original bet of $(100, 0)$. The difference between the certainty equivalent wealth of a risk-averse person (in our example, \$37) and that of a risk-neutral person (\$50) makes it possible for insurance companies to provide a benefit to society and still make a profit. If an insurance company offers to “buy” the bet $(100, 0)$ from our risk-neutral individual for, say, \$40, he will accept the offer because he is indifferent between the bet $(100, 0)$ and having \$37 with certainty, and having \$40 with certainty is better than having \$37 with certainty. (In the language of economics, accepting the offer of \$40 will put him on a higher indifference curve, i.e., one corresponding to a higher utility level.) What’s in it for the insurance company? Well, they spend only \$40 to buy a bet that has an expected value of \$50; by buying many such bets, the insurance company has an extremely good chance of making a profit.

Problems

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....)
 - (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 12.)

9. Although many services are provided on a fee-up-front basis (for example, you pay before you enter the movie theater), some services are provided on the "honor system". Parking meters are one example: instead of paying someone when you park your car, you are "on your honor" to put money in the meter. (Some cities also do this with their public transportation systems.) Of course, it isn't just a matter of honor: there are also enforcement officers ("meter maids") who show up from time to time and penalize rule-breakers. So:

- (a) If you “risk it” by putting nothing in the meter, there’s an 80% chance that you’ll get away with it (so your cost will be \$0) and a 20% chance that you’ll get a \$20 ticket. What is the expected cost of taking this risky option?
- (b) Imagine that the city managers want to save money by cutting in half the number of enforcement officers (so that the chance of getting a ticket is only 10%). Can you suggest a way to do this without drastically increasing the attractiveness of cheating?

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




Figure 3.1: A decision tree involving time

is about \$.91: put \$.91 in the bank today and in a year you'll have about \$1.00. So we can use interest rates to compare apples to apples instead of apples to oranges.

It is important to note that **inflation**—a general increase in prices, which will be discussed in detail in the next chapter—is one reason why money today and money tomorrow are not commensurate. However, inflation is not the only reason, and to isolate its effects *you should assume for the remainder of this chapter that there is no inflation*, i.e., that the general price level remains constant over time. This assumption *does not* make money today commensurate with money tomorrow, for the simple reason that people generally prefer not to wait for things: offered a choice between a TV today and a TV next year, most people prefer the TV today. This time preference produces a positive interest rate even in the absence of inflation.

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:

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

So the **future value of a lump sum** of \$ x invested for n years at interest rate s is $FV = x(1+s)^n$. 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 s is

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

Note that the present value formula can be used to evaluate lump sums received in the *past* as well as the future: simply use, say, $n = -1$ to calculate

the present value of a lump sum received one year ago.¹ The present value formula can also be used to calculate lump sums received in the *present*: simply use $n = 0$ to find that the present value of receiving $\$x$ today (i.e., zero years from today) is $\$x$.²

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:

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

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

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

¹The result here is identical to using the future value formula with $n = 1$: in both cases the idea is that you can put that money in the bank and allow it to gain interest for one year.

²Again, the intuition is straightforward: you need to put $\$x$ in the bank today in order to have $\$x$ in the bank today. Thanks to the fall 2004 students at Whitman College for the observations addressed in this footnote and the previous footnote.

Now we subtract the first equation from the second equation:

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

The left hand side of this equation simplifies to $.05 \cdot PV$. 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

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

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

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

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

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

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., $s \approx .005$.³

³Because 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.

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:

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

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:

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

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., $s \approx .005$) 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**,⁴ and it really is forever: after death, you can give the perpetuity away in your will! So what we're looking for is

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

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

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

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

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

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

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

⁴Note that the word comes from perpetual, just like annuity comes from annual.



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

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 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 $\frac{10}{.05} = \$200$.

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 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.⁵

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!⁶

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

⁵Another 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.

⁶Puzzle: What about stock splits?




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.




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

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.

3.5 *Math*: Present Value and Budget Constraints

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

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

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

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

Continuous compounding

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

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

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

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

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

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

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

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

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

Answer: We want an amount x such that $x \cdot e^{(.1)(3)} = 100$. The solution is $x = 100 \cdot e^{-.3} \approx \74.08 .

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 .
 - (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 s 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 $s\%$ (e.g., 10%), then your money will double in $\frac{72}{s}$ 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. Explain (as if to a non-economist) the phrases “fish are capital,” “trees are capital,” and/or “oil is capital,” or otherwise explain the importance of the interest rate at the Bank of America in management decisions regarding natural resources such as fish, trees, and oil.
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 60 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 6%.
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 60 for more about this problem.)




Figure 3.6: A population growth function for fish.

13. Imagine that you own a lake and that you're trying to maximize the present value of catching fish from the lake, which currently has 1000 fish in it. The population growth function of the fish is described in Figure 3.6.
 - (a) The maximum sustainable yield policy is to catch 100 fish at the end of this year, 100 fish at the end of the next year, and so on, forever. If the price of fish is \$1 per fish, what is the resulting present value at a 5% interest rate?
 - (b) An alternative policy is to catch 400 fish *today* (so that 600 remain in the lake), and then catch 84 fish at the end of this year, 84 fish at the end of the next year, and so on, forever. What is the resulting present value? (Assume as above a price of \$1 per fish and an interest rate of 5%.) Is it higher or lower than the present value of the maximum sustainable yield policy?
14. Assume that you've just bought a new carpet. The good news is that the carpet will last forever. The bad news is that you need to steam-clean it at the end of every year (i.e., one year from today, two years from today, etc.). What you need to decide is whether to buy a steam-cleaner or just rent one every year. *You can use the bank to save or borrow money at a 5% interest rate.*
 - (a) Will the amount you paid for the carpet affect your decision regarding renting versus buying?
 - (b) One year from today (i.e., when you first need to clean the carpet), you'll be able to buy a steam-cleaner for \$500; like the carpet, the steam-cleaner will last forever. Calculate the present value of this cost.

- (c) The alternative to buying is renting a steam-cleaner, which will cost you \$20 at the end of every year forever. Calculate the present value of this cost. Is it better to rent or buy?
- (d) Imagine that your friend Jack calls to tell you that steam-cleaners are on sale (today only!) for \$450: “You’d have to be a moron to pay \$20 every year forever when you can just pay \$450 today and be done with it!” Write a brief response explaining (as if to a non-economist) why you do or do not agree.

Chapter 4

Math: Trees and Fish

4.1 Trees

Let $f(t)$ be the size of a tree planted at time $t = 0$ (or, more correctly, the amount of lumber that can be harvested from such a tree). For the simplest case, consider a tree farmer who *cannot* replant his tree, has no alternative use of the land, and has no harvesting or planting costs. We also assume that there is no inflation, that the price of timber is constant at p , and that the real interest rate is $100 \cdot r\%$, compounded continuously.

The farmer's problem is to choose the harvest time t to maximize the present value of the lumber, i.e., choose t to maximize

$$PV = pf(t)e^{-rt}.$$

Taking a derivative with respect to the choice variable t and setting it equal to zero we get

$$pf'(t)e^{-rt} + pf(t)e^{-rt}(-r) = 0 \implies pf'(t) - r \cdot pf(t) = 0.$$

This necessary first-order condition (NFOC) has two intuitive interpretations.

The first comes from rewriting it as $pf'(t) = r \cdot pf(t)$. The right hand side here is the interest payment we will get if we cut the tree down and invest the proceeds $pf(t)$ in the bank at interest rate r . The left hand side is the "interest" we will get from letting the tree grow, namely the value $pf'(t)$ of the additional timber we will get. Our NFOC says that we should cut the tree down when these interest payments are equivalent. Otherwise, we are not maximizing the present value of the lumber: if $pf'(t) > r \cdot pf(t)$ then we should hold off on cutting down the tree because its interest payment is greater than the bank's; if $pf'(t) < r \cdot pf(t)$ then we should cut the tree down earlier because its interest payment is less than the bank's.

The second intuitive interpretation comes from rewriting the NFOC as $\frac{f'(t)}{f(t)} = r$. Here the right hand side is the (instantaneous) interest rate you

get from investing in the bank, and the left hand side is the (instantaneous) interest rate you get from investing in the tree: if the tree adds 10 board-feet worth of lumber each year and currently contains 500 board-feet, then the tree's current growth rate is 2%: $\frac{f'(t)}{f(t)} = \frac{10}{500} = .02$. In this formulation the NFOC says to cut the tree down when the interest rate from the tree equals the interest rate from the bank. If $\frac{f'(t)}{f(t)} > r$ then we should hold off on cutting down the tree because it's a better investment than the bank; if $\frac{f'(t)}{f(t)} < r$ then the tree is a worse investment than the bank and we should have cut it down earlier.

Harvesting Costs

What if there's a cost c_h to cut down the tree? Intuitively, you should cut down the tree *later* than in section 4.1 because by waiting you push the harvesting cost further into the future, which benefits you because the present value of the cost goes down.

Mathematically, our problem is to choose the harvest time t to maximize the present value of the lumber, i.e., choose t to maximize

$$PV = [pf(t) - c_h]e^{-rt}.$$

Taking a derivative with respect to the choice variable t and setting it equal to zero we get

$$pf'(t)e^{-rt} + [pf(t) - c_h]e^{-rt}(-r) = 0 \implies pf'(t) - r \cdot [pf(t) - c_h] = 0.$$

We can rewrite this as $pf'(t) + r \cdot c_h = r \cdot pf(t)$, which has a nice intuitive interpretation. The left hand side is the return $pf'(t)$ we get from leaving the tree in the ground for another year, *plus* the return $r \cdot c_h$ we get from another year of investing the amount c_h we need to cut the tree down. The right hand side is the return $r \cdot pf(t)$ we get from cutting the tree down immediately and putting the proceeds in the bank.

Both the intuition and the math support our original guess, which is that harvesting costs should induce the farmer to wait longer to cut down the tree. (Mathematically, this is because the term $r \cdot c_h$ "boosts" the tree's interest rate, making it look like a more valuable investment relative to the bank.)

Planting Costs

What if there's a cost c_p to planting the trees in the first place? Intuitively, the amount c_p becomes sunk once you plant the trees; therefore it should not affect the optimal rotation time.

Mathematically, our problem is to choose the harvest time t to maximize the present value of the lumber, i.e., choose t to maximize

$$PV = pf(t)e^{-rt} - c_p.$$

Since c_p is a constant, the choice of t that maximizes this objective function will be the same one that maximizes the objective function in section 4.1. So we get the same NFOC and the same optimal rotation time.

Having said that, it is important to note that c_p is *not* irrelevant. For high enough values of c_p , your present value from planting trees will become negative. In this case, your profit-maximizing choice is to not plant anything, which gives you a present value of zero.

Replanting

Now assume there are no harvesting or replanting costs, but that you can replant the tree. What is your optimal management policy? Well, first we need to figure out our objective function, because it's no longer $PV = pf(t)e^{-rt}$. There are a couple of ways of figuring this out. One is to use brute force: assume that you cut the tree down and replant at times $t, 2t, 3t, \dots$. Then the present value of your timber is

$$\begin{aligned} PV &= pf(t)e^{-rt} + pf(t)e^{-2rt} + pf(t)e^{-3rt} + \dots \\ &= pf(t)e^{-rt} [1 + e^{-rt} + e^{-2rt} + \dots] \\ &= pf(t)e^{-rt} \cdot \frac{1}{1 - e^{-rt}} \\ &= \frac{pf(t)}{e^{rt} - 1} \text{ (multiplying through by } \frac{e^{rt}}{e^{rt}} \text{ and cancelling).} \end{aligned}$$

A second way to look at this is to realize that cutting the tree down and replanting essentially puts you in the same place you started, so that

$$PV = pf(t)e^{-rt} + PVe^{-rt}.$$

Rearranging this we get

$$PVe^{rt} = pf(t) + PV \implies PV(e^{rt} - 1) = pf(t) \implies PV = \frac{pf(t)}{e^{rt} - 1}.$$

Having determined the thing we want to maximize (PV), we now take a derivative with respect to the choice variable t and set it equal to zero:

$$\frac{(e^{rt} - 1) \cdot pf'(t) - pf(t) \cdot e^{rt}(r)}{(e^{rt} - 1)^2} = 0.$$

Simplifying (multiplying through by the denominator and rearranging) we get

$$pf'(t) = r \cdot pf(t) \cdot \frac{e^{rt}}{e^{rt} - 1} \iff \frac{f'(t)}{f(t)} = r \cdot \frac{e^{rt}}{e^{rt} - 1}.$$

The exact intuition behind this answer is not very clear, but we *can* learn something by comparing this answer with the answer we got previously (at the

start of section 4.1, when replanting was not possible). Since $\frac{e^{rt}}{e^{rt} - 1} > 1$, replanting “boosts” the interest rate paid by the bank, i.e., makes the bank look like a more attractive investment and the tree look like a less attractive investment.

To see why, consider the optimal harvest time t^* from the no-replanting scenario, i.e., the time t^* such that $\frac{f'(t^*)}{f(t^*)} = r$. At time t^* with no replanting, the farmer is indifferent between investing in the bank and investing in the tree. But at time t^* *with* replanting, our NFOC says that something is wrong: $\frac{f'(t^*)}{f(t^*)} = r < r \cdot \frac{e^{rt^*}}{e^{rt^*} - 1}$, meaning that the farmer should have cut the tree down earlier and put the money in the bank. Why? Because cutting the tree down earlier allows the farmer to replant earlier, and replanting earlier increases the present value of the harvests from replanting. So there *is* some intuition behind the result that replanting should reduce the optimal harvest time.

All Together Now

What if there is replanting and also harvesting and/or planting costs? Then your problem is to choose t to maximize

$$PV = [pf(t) - c_h]e^{-rt} - c_p + PVe^{-rt}.$$

Solving this for PV (ugly!) and setting a derivative equal to zero would show that harvesting costs and planting costs both increase the optimal rotation time compared to section 4.1: increasing the rotation time pushes your harvesting and/or planting costs further into the future, which benefits you by decreasing their present value.

Alternative Use of the Land

What if there’s an alternative use of the land (e.g., building a factory) that you can choose instead of planting trees? In this case you have two choices: either build the factory or plant trees. If you build the factory, your present value is, say, x . If you plant trees, your present value is given by the equations above. You should choose the option with the higher present value. If your optimal choice is planting trees, the amount x will have no impact on your optimal rotation time.

4.2 Fish

Let’s say you inherit a small lake with a bunch of fish in it. Every fishing season you catch a bunch of fish, and the ones that are left grow and reproduce, meaning that you’ll have more fish at the start of the next fishing season. How many fish should you catch each year in order to maximize the present value of your fish?




Figure 4.1: A growth function for fish, $G(s) = .2s - .001s^2$

There are lots of important factors in this problem, starting with the population dynamics of the fish. We can crudely model those population dynamics with a **growth function**, $G(s)$: if at the end of the fishing season you have a **stock size** of s pounds of fish, by the following fishing season you'll have $s + G(s)$ pounds of fish, i.e., $G(s)$ more pounds than you started with. For our example, we'll use $G(s) = .2s - .001s^2$, which is graphed in Figure 4.1. Is this a reasonable growth function? Well, you have to ask the biologists to find out for sure. But one good sign is that $G(0) = 0$; this says that if you start out with zero fish then you won't get any growth. Another good sign is that the lake has a **maximum carrying capacity**, i.e., a maximum number of fish that its ecology can support. To find the carrying capacity, we set $G(s) = 0$ and find that one solution is $s = 200$; in other words, if you have 200 pounds of fish in the lake at the end of the fishing season, you'll still have 200 pounds of fish in the lake at the beginning of the next season because that's the maximum amount the lake can support.

A Two-Period Model

To make the model as simple as possible, we also make the following assumptions. First, there are only two time periods, this year ($t = 1$) and next year ($t = 2$). Second, the price of fish is $\$p$ per pound and is constant over time. Third, there are no costs associated with fishing; you just whistle and the fish jump into your boat. Fourth, the interest rate at the bank is $100 \cdot r\%$, compounded annually. Finally, we assume that you inherit the lake at the beginning of this year's fishing season, and that the current stock size is s_0 .

Given the stock size at the end of your first season (call it s_1), your stock size at the beginning of the second season is given by $s_2 = s_1 + G(s_1)$. Since

there are only two periods and no fishing costs, whatever you don't catch this year you should catch next year. Your maximization problem, then, is to choose s_1 to maximize the present value of your harvest,

$$PV = p(s_0 - s_1) + p \left[\frac{s_1 + G(s_1)}{1 + r} \right].$$

The first term here, $p(s_0 - s_1)$, is your profit from catching and selling $s_0 - s_1$ pounds of fish now. The second term is the present value of your profit from catching and selling the remaining fish next year.

Taking a derivative with respect to s_1 and setting it equal to zero we get a necessary first-order condition of

$$p(-1) + p \cdot \frac{1}{1 + r} \cdot [1 + G'(s_1)] = 0.$$

Multiplying through by $\frac{1 + r}{p}$ we get

$$-1 - r + 1 + G'(s_1) = 0 \implies G'(s_1) = r.$$

This NFOC has a nice interpretation. The right hand side, r , is the (instantaneous) interest rate paid by the Bank of America. The left hand side, $G'(s_1)$, is the (instantaneous) interest rate paid by the Bank of Fish: if you leave an extra pound of fish in the lake, after one year you'll get a "fish interest payment" of $G'(s_1)$. So the NFOC says to equate the interest rate at the two banks: if $G'(s_1) > r$, you should increase s_1 by leaving more fish in the lake, where they'll grow at a faster rate than money in the bank; if $G'(s_1) < r$, you should decrease s_1 by catching more fish and putting the proceeds in the bank, where the money will grow at a faster rate than the fish population.

What happens if the value of s_1 that satisfies this equation is greater than s_0 ? Since we cannot catch negative amounts of fish, the best we can do is choose the corner solution $s_1 = s_0$, i.e., catch zero fish this period. So the complete answer to our problem is to find s^* that satisfies $G'(s^*) = r$ and then set $s_1 = \min\{s_0, s^*\}$. Technically, this is called a **bang-bang solution**. (Non-technical explanation: if $s^* < s_0$, you shoot the fish until $s_1 = s^*$; if $s^* > s_0$, you shoot the fishermen!)

An Infinite-Period Model

Now let's do an infinite-period fish model, using $G(s) = .2s - .001s^2$ as our growth function. To make the model as simple as possible, we also make the following assumptions. First, the price of fish is \$1 per pound and is constant over time. Second, there are no costs associated with fishing; you just whistle and the fish jump into your boat. Third, the interest rate at the bank is $r = .05$, i.e., 5%, compounded annually. Finally, we assume that you inherit the lake at the beginning of this year's fishing season, and that the current stock size is $s = 110$.

Given the current fish population, there is an intuitively appealing solution to your fisheries management problem: follow a policy of **Maximum Sustainable Yield (MSY)**. This says that you should harvest 10 pounds of fish this year, leaving you with a population of $s = 100$, which will grow to $100 + G(100) = 110$ pounds by next fishing season, at which point you can harvest 10 pounds of fish, leaving you with a population of $s = 100$. . . We can see from the graph of $G(s)$ (and from setting $G'(s) = 0$) that harvesting 10 pounds of fish every year is the maximum number of fish you can catch year after year, i.e., is the maximum sustainable yield.

If you follow this policy, the present value of your lake is (using the perpetuity formula)

$$PV = 10 + \frac{10}{1.05} + \frac{10}{(1.05)^2} + \frac{10}{(1.05)^3} + \dots = 10 + \frac{10}{.05} = 210.$$

This sounds pretty good, but now let's do a little marginal analysis to see if we can do better. Let's change the MSY policy a little bit, as follows: let's "borrow" some extra fish (say, 5 pounds) this year, and "return" them next year. Effectively what we're doing here is borrowing from the Bank of Fish (by catching the fish and selling them) and using that borrowed money to invest in the Bank of America (by depositing it and letting it gain interest); in a year we're going to unwind our investments by taking the proceeds from the Bank of America and using them to repay our loan from the Bank of Fish. Whether or not this is a good idea depends on how the "interest rate" at the Bank of Fish compares to the interest rate at the Bank of America.

Some numbers may make this more clear. What we're going to do this year is increase the harvest by 5 pounds of fish (to 15 pounds), so that at the end of the season there will be $110 - 15 = 95$ pounds of fish in the lake. In a year there will be $95 + G(95) = 104.975$ pounds of fish in the lake. Next year we limit our harvest to 4.975 pounds, leaving 100 pounds of fish in the lake at the end of the second season. This returns us to the MSY policy, and each year thereafter we harvest 10 pounds of fish.

What's the present value of this? Well, it's

$$\begin{aligned} PV &= 15 + \frac{4.975}{1.05} + \frac{10}{(1.05)^2} + \frac{10}{(1.05)^3} + \frac{10}{(1.05)^4} + \dots \\ &= 15 + \frac{4.975}{1.05} + \frac{1}{1.05} \left[\frac{10}{1.05} + \frac{10}{(1.05)^2} + \frac{10}{(1.05)^3} + \dots \right] \\ &= 15 + \frac{4.975}{1.05} + \frac{1}{1.05} \left[\frac{10}{.05} \right] \\ &\approx 15 + 4.74 + 190.48 \\ &= 210.22. \end{aligned}$$

We've done better than MSY!

Intuitively, what's going on here is that the Bank of America is paying a higher interest rate than the Bank of Fish; we therefore make more money by

borrowing a little bit from the Bank of Fish and investing it in the Bank of America. Let's formalize this with math.

The math

Say you decide to catch M pounds of fish this year, leaving s pounds of fish in the lake. In a year, then, you'll have $\$M(1+r)$ in the Bank of America and $s+G(s)$ in the Bank of Fish. If this is the decision that maximizes your present value, you can do no better by catching more fish this year or less fish this year, and then making up for it next year. Let's see what the implications of this are.

First let's try catching less fish this year, i.e., "lending" h pounds of fish to the Bank of Fish for a year. (Note that lending to the Bank of Fish effectively means borrowing from the Bank of America: by not catching those h pounds of fish, we're reducing the amount of money we put in the Bank of America and therefore reducing the interest payment we'll get.) So: lending h pounds of fish to the Bank of Fish means that you catch $M-h$ pounds of fish this year, leaving $s+h$ pounds in the lake. In a year, you'll have $\$(M-h)(1+r)$ in the Bank of America and $s+h+G(s+h)$ in the Bank of Fish. You can then unwind the investment by catching $h+G(s+h)-G(s)$ pounds of fish and putting the proceeds in the Bank of America, leaving you with $s+G(s)$ in the Bank of Fish and

$$[(M-h)(1+r)] + [h+G(s+h)-G(s)] = M(1+r) - h \cdot r + G(s+h) - G(s)$$

in the Bank of America. Since your original choice of s was optimal, it must be the case that you have less in the Bank of America now than you did before:

$$M(1+r) - h \cdot r + G(s+h) - G(s) \leq M(1+r).$$

Cancelling terms and rearranging we get

$$\frac{G(s+h) - G(s)}{h} \leq r.$$

The left hand side here looks exactly like the definition of a derivative, and letting h go to zero we get

$$G'(s) \leq r.$$

This says that if s is the optimal number of fish to leave in the lake, then the instantaneous interest rate $G'(s)$ paid by the Bank of Fish must be no greater than the instantaneous interest rate r paid by the Bank of America. Otherwise, you could increase your present value by "lending" to the Bank of Fish.

A similar result follows when we consider "borrowing" h pounds of fish from the Bank of Fish for a year. (Here, borrowing from the Bank of Fish means investing in the Bank of America: by catching an extra h pounds of fish, we can increase the amount of money we put in the Bank of America.) The result we get here is that $G'(s) \geq r$, i.e., that the instantaneous interest rate $G'(s)$ paid by the Bank of Fish must be no less than the instantaneous interest rate r paid

by the Bank of America. Otherwise, you could increase your present value by “borrowing” from the Bank of Fish.

Combining our two answers, we see that the optimal stock size s must satisfy

$$G'(s) = r.$$

Back to our example

In the problem we started out with— $G(s) = .2s - .001s^2$, $r = .05$, initial stock of 110 pounds of fish—the condition $G'(s) = r$ comes out to

$$.2 - .002s = .05 \implies s = 75.$$

So the optimal stock size is 75 pounds of fish, meaning that you should catch $110 - 75 = 35$ pounds the first year (thereby reducing the stock size to 75) and $G(75) = 9.375$ pounds each year thereafter (thereby maintaining the stock size at 75). Your present value from this management strategy is

$$\begin{aligned} \text{PV} &= 35 + \frac{9.375}{1.05} + \frac{9.375}{(1.05)^2} + \frac{9.375}{(1.05)^3} + \dots \\ &= 35 + \frac{9.375}{.05} = 222.50 \end{aligned}$$

So 222.50 is the maximum present value you can get from your lake!

Problems

1. You’ve just inherited a plot of land, and decide to manage it so as to maximize its present value. The land has no alternative use other than growing trees, so if you don’t plant anything your present value will be zero. Assume that there is no inflation, and that the price of lumber stays constant at p per unit. The interest rate paid by the bank is $100 \cdot r\%$ per year, compounded continuously. Trees grow according to $f(t)$.
 - (a) (The basic problem) You cannot replant the trees, and there are no costs associated with planting or cutting down your trees. Write down the maximization problem and find the necessary first-order condition (NFOC) for optimal management. What is the intuition behind your result?
 - (b) (Harvesting costs) You cannot replant the trees, but cutting down the trees costs you c . Write down the maximization problem and find the necessary first-order condition for optimal management. Does harvesting costs increase or decrease your optimal rotation time compared to (a)? What is the intuition behind this result?
 - (c) (Planting costs) You cannot replant the trees, and it costs you nothing to cut them down, but in order to plant the trees you need to

spend c today. Write down the maximization problem and find the necessary first-order condition for optimal management. Does planting costs increase or decrease your optimal rotation time compared to (a)? What is the intuition behind this result? What role (if any) does c play in your decision? (Hint: What if c were \$1 billion?)

- (d) (Replanting) Now you *can* replant the trees; there are no costs associated with planting or cutting. Write down the maximization problem and find the necessary first-order condition for optimal management. Does replanting increase or decrease your optimal rotation time compared to (a)? What is the intuition behind this result?
- (e) (Challenge: all together now) Imagine that you can replant, but that there are also planting costs, or harvesting costs, or both. (In the case of planting costs, assume that you replant new trees at the same instant that you cut down the old trees.) Try to think through the intuition of what would happen to optimal rotation time *compared to (d)* in these situations. Then write down the maximization problem and (if you're brave) try to confirm your intuition.
- (f) (Alternative uses) Assume (in any of these problems) that there is an alternative use of the land that has a present value of x . What role (if any) does x play in your decision-making? (Hint: Think about your answer in (c).)

Chapter 5

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.

5.1 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%.

5.2 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 5.1a. 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 5.1b and 5.1c, both of which show average U.S. gas prices between 1918 and 2002. Figure 5.1b 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 5.1c 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 5.1c 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 5.1b and 5.1c, consider the price of a gallon of gasoline in 1925. Figure 5.1b 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 5.1c 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 5.1d 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 5.1c except for the labels on the vertical axis (1925 dollars versus year 2000 dollars).

¹This would be as simple as it sounds if people bought the same stuff every year, but it isn’t because they don’t. Take an advanced economics course to learn how to deal with complications caused by changes in the representative market basket over time.

²The 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?

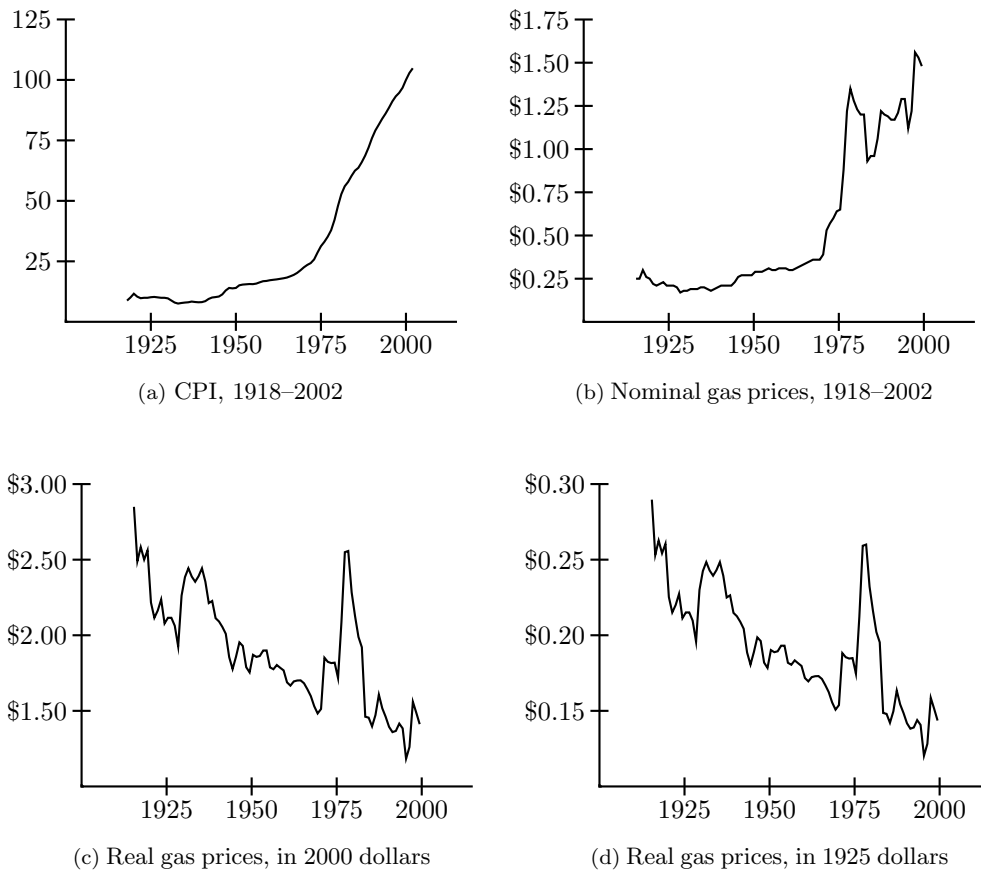


Figure 5.1: Inflation and average U.S. gas prices, 1918–2002. Figure (a) shows the Consumer Price Index (CPI) measure of inflation; note that there was a period of **deflation** (i.e., a general *decrease* in prices over time) during the Great Depression that began in 1929. Figure (b) shows the average 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.)

5.3 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 5.1'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 41) 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 42) 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 6

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*. It can also be used in contexts that don’t involve financial instruments: a driver in rush-hour traffic who switches from a slow-moving lane into a faster-moving one might reasonably be said to be engaging in arbitrage.

6.1 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.

6.2 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

¹Friedman’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.

6.3 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.

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

Part II

One v. One, One v. Many

Chapter 7

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?).

7.1 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 12 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.

7.2 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

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

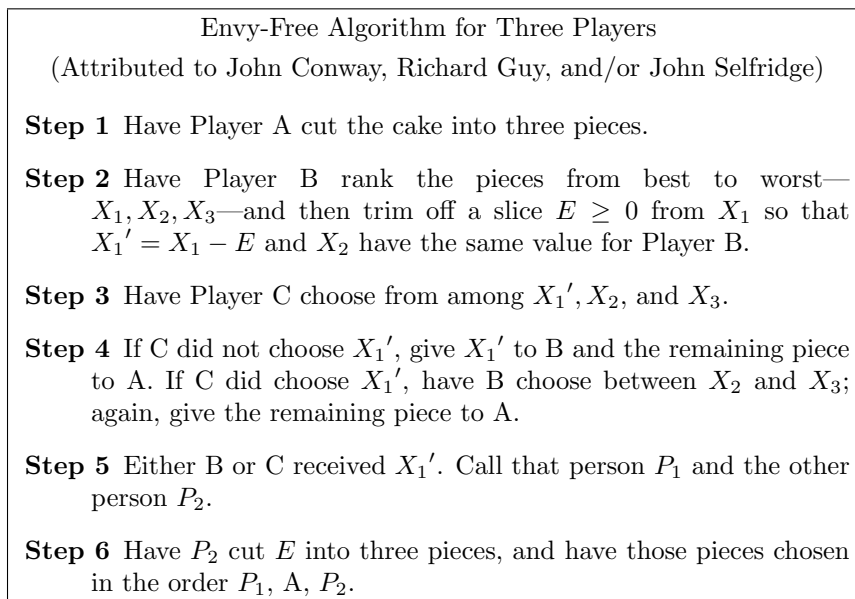


Figure 7.1: An envy-free algorithm for three players.

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 7.1 shows one, adapted from Section 1.4 of *Cake-Cutting Algorithms*.

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”?

²Robertson, 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]

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.

7.3 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.

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.

³Brams, 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].

⁴Brams, 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].

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 76; 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.

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

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.

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. Fractions are okay.) What is Alice's cost of fishing in terms of vegetables?
- (h) What is Bob's cost of fishing in terms of vegetables? What is Bob's cost of vegetables in terms of fish?
- (i) In terms of vegetables, who is the least-cost producer of fish?
- (j) In terms of fish, who is the least-cost producer of vegetables?

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

Chapter 8

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 in-

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

dividual, however, welfare economics seeks to pick the best outcome for society as a whole.

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

8.1 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

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

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

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 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.

8.2 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

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!

8.3 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 7.3 on pages 74–76.) For example, Pareto efficiency in

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

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 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

⁵These 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**.

⁶The 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).

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

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. “If situation A is Pareto efficient and situation B is Pareto inefficient, situation A must be a Pareto improvement over situation B.” Do you agree with this claim? If so, explain. If not, provide a counter-example or otherwise explain.
 4. Consider a division problem such as the division of cake or the allocation of fishing quotas.
 - (a) Economists tend to place a great deal of importance on providing opportunities to trade (e.g., allowing the buying and selling of fishing quotas). Briefly explain why this is.
 - (b) “Even if there are opportunities to trade, the initial allocation of resources (e.g., the determination of who gets the fishing quotas in an ITQ system) is important because it helps determine whether or not we reach *the* Pareto efficient allocation of resources.”
 - i. Is there such a thing as “*the* Pareto efficient allocation of resources”? Explain briefly.
 - ii. Do you agree that initial allocations are important in order to achieve Pareto efficiency, or do you think that they’re important for a different reason, or do you think that they’re not important? Support your answer with a brief explanation.
 5. Many magazines and newspapers have on-line archives containing articles from past issues. As with many information commodities, it 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 5b is unlikely to happen.

Chapter 9

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 9.1 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.

9.1 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 9.2. 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 9.1: 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 9.3. 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 9.4: 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 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




Figure 9.2: The upper subgame of the grenade game




Figure 9.3: The lower subgame of the grenade game

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 9.4: The grenade game reduced to a decision tree

Problems

1. *Challenge.* Explain (as if to a non-economist) why backward induction makes sense.
2. Analyze the sequential move game in figure 9.5 using backward induction.
 - (a) Identify (e.g., by circling) the likely outcome of this game.
 - (b) Is this outcome Pareto efficient? If it is not Pareto efficient, identify a Pareto improvement.




Figure 9.5: A random game tree

3. Analyze the sequential move game in figure 9.6 using backward induction.
 - (a) Identify (e.g., by circling) the likely outcome of this game.
 - (b) Is this outcome Pareto efficient? If it is not Pareto efficient, identify a Pareto improvement.




Figure 9.6: Another random game tree

4. Consider tipping, a social phenomenon observed in some (but not all!) countries in which restaurant patrons leave some extra money behind for their waiter or waitress. Would tipping provide much of an incentive for good service if the tips were handed over at the beginning of the meal rather than at the end? Are there any difficulties in the incentive structure when tips are left at the end of meals? Write down game trees to support your arguments.
5. (Overinvestment as a barrier to entry) Consider the following sequential move games of complete information. The games are between an incumbent monopolist (M) and a potential entrant (PE). You can answer these questions without looking at the stories, but the stories do provide some context and motivation.




Figure 9.7: Story #1

Story #1 (See figure 9.7): 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.)

Story #2 (See figure 9.8): 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.




Figure 9.8: Story #2

- (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.)
6. (The Sticks Game) The sticks game works as follows: We put n sticks on the table. Beginning with Player 1, the two players take turns removing either one or two sticks. The player who removes the last stick must pay the other player \$1.
- If there are 10 sticks on the table, which player would you rather be, and what strategy will you employ? [Hint: Use backwards induction!]
 - Challenge* If there are n sticks on the table, which player would you rather be? Can you describe a general strategy?
 - 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).
7. (The Ice Cream Pie Game, from Dixit and Nalebuff) Two players take turns making take-it-or-leave-it offers about the division of an ice cream pie. In the first round, the whole pie is available; if Player 2 accepts Player 1's proposal then the two players share the entire pie; if Player 2 rejects Player 1's proposal, half of the pie melts away, and we go to round two (in which Player 2 makes a take-it-or-leave-it offer about the division of the remaining pie). The game ends when an offer is accepted, or after the end

of the n th period (at which point Mom eats the remaining pie, meaning that the players get nothing).

- (a) Predict the outcome of the game when there are 1, 2, and 3 periods.
 - (b) Now assume that $1/3$ rd 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.
8. 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.
 9. (The Draft Game, from Brams and Taylor) Three football teams (X, Y, Z) are involved in a draft for new players. There are six players to choose from (Center, Guard, Tailback, Quarterback, Halfback, Fullback), and the draft works as follows: 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 9.9.*

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?




Figure 9.9: The draft game

- (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.
10. *Fun.* (The Hold-Up Problem) In the movie *Butch Cassidy and the Sundance Kid* (1969), Paul Newman and Robert Redford play Wild West bank robbers who are particularly fond of robbing the Union Pacific Railroad. The CEO of the railroad, Mr. E. H. Harriman, hires a “superposse” of gunslingers to bring the two robbers to justice, dead or alive. After a long (and rather boring) chase scene, Butch and Sundance manage to escape. Afterwards, Butch reads about the superposse in the newspaper and has this to say:

A set-up like that costs more than we ever took...That crazy Harriman. That's bad business. How long do you think I'd stay in operation if every time I pulled a job, it cost me money? If he'd just pay me what he's spending to make me stop robbin' him, I'd stop robbin' him. [Screaming out the door at E. H. Harriman:] Probably inherited every penny you got! Those inherited guys—what the hell do they know?

- (a) Is what Harriman is doing bad business? Explain why or why not. Your answer may depend on the assumptions you make, so explicitly state any and all assumptions. You might also want to draw a game tree, make up appropriate payoffs, and solve the game using backwards induction.
 - (b) "If he'd just pay me what he's spending to make me stop robbin' him, I'd stop robbin' him." Assume this statement is true. What does it say about the efficiency or inefficiency of the situation?
 - (c) What do you think about the argument contained in the previous quote? Can you see why this is called the "hold-up problem"?
 - (d) The hold-up problem also applies to students working jointly on projects and to firms engaged in joint ventures: after one member makes an irreversible investment, the other member may try to renegotiate the terms of the deal. Explain how contracts might help in preventing this difficulty, and why contracts wouldn't work in the case of Butch Cassidy.
11. *Fun.* The IgoUgo travel guide Andiamo provides the following advice for crossing the street in Rome: "First, stand towards a coffee bar and watch a local or two. See how they boldly walk out into traffic? Now it's your turn! Choose your moment but don't hesitate for too long. Waiting for traffic to clear will not happen. When you appear to have the most lead-time, step boldly off the curb and walk swiftly and confidently for the opposite side of the street. Do not look at the traffic careening towards you—believe it or not, they will stop for you! But do not look at them—do not make eye contact—this is an invitation for sport. Just walk briskly with your head up and your eyes on the prize- the opposite sidewalk."
- (a) "[D]o not make eye contact—this is an invitation for sport." Explain.
 - (b) Set up a game tree for this "pedestrian in Rome" game and solve it.
12. (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?
13. *Fun*. (The surprise exam paradox) Your class meets 5 days a week, and on Friday your teacher tells you that there will be a surprise exam next week, meaning (1) that there will be an exam, and (2) that it will be a surprise (i.e., you won't be able to anticipate the night before the exam that the exam will be the next day). What can you conclude about the exam? Relate this problem to the Centipede Game discussed previously.

Chapter 10

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 10.1. 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 10.1: The duopoly game

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

10.1 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.

10.2 The Prisoners’ Dilemma

Figure 10.1 is one version of an important game called the Prisoners’ Dilemma. Another version (which explains the game’s name) is shown in Figure 10.2. 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 10.2.)

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

Figure 10.2: The Prisoners’ Dilemma

Question: What is a good prediction for the Prisoners' 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 Prisoners' 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 Prisoners' 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 Prisoners' Dilemma game? To examine this possibility, we'll use the version of the Prisoners' Dilemma game shown in Figure 10.3; 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 10.3: Another version of the Prisoners' 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.

10.3 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 Prisoners' Dilemma game in Figure 10.3 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 10.4: 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 10.4 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) What outcome do you expect if this game is played twice? Explain briefly.

Chapter 11

Iterated Dominance and Nash Equilibrium

In the previous chapter we examined simultaneous move games in which each player had a dominant strategy; the Prisoner's Dilemma game was one example. In many games, however, one or more players do not have dominant strategies. This chapter explores two solution concepts that we can use to analyze such games.

The first solution concept, iterated dominance, is a **refinement** of the dominant strategies approach from the previous chapter, meaning that iterated dominance is a stronger technique that builds upon (or refines) the results of the dominant strategies approach. In other words: the idea of dominant strategies often allows us to narrow down our prediction for the outcome of a game; iterated dominance allows us to narrow down our prediction at least as far, and sometimes further.

Unfortunately, this extra strength does not come for free. While dominant strategies is a reasonably simple idea, iterated dominance is (while not exactly a Nobel-prize-winning concept) one step closer to rocket science. As such, it requires more powerful assumptions about the intellectual capabilities of the optimizing individuals who are playing the games.

The second solution concept in this chapter, Nash equilibrium, is a refinement of iterated dominance: Nash equilibrium allows us to narrow down our prediction at least as far as iterated dominance, and sometimes further. Again, this extra strength does not come for free. Nonetheless, Nash equilibrium is one of the central concepts in the study of strategic behavior—a fact which helps explain why Nash equilibrium *is* a Nobel-prize-winning concept.

11.1 Iterated Dominance

The transition from dominant strategies to iterated dominance involves two ideas. The first is this: even when a player doesn't have a dominant strategy

(i.e., a *best* strategy, regardless of what the other players do), that player might still have one strategy that dominates another (i.e., a strategy A that is *better* than strategy B, regardless of what the other players do). As suggested by the terms “best” and “better”, the difference here is between a superlative statement (e.g., “Jane is the best athlete in the class”) and a comparative statement (“Jane is a better athlete than Ted”); because comparatives are weaker statements, we can use them in situations where we might not be able to use superlatives.

For example, consider the game in Figure 11.1. First note that there are no strictly dominant strategies in this game: U is not the best strategy for Player 1 if Player 2 plays L or C, M is not the best strategy for Player 1 if Player 2 plays R, and D is not the best strategy for Player 1 if Player 2 plays L or C. Similarly, L is not the best strategy for Player 2 if Player 1 plays U or D, C is not the best strategy for Player 2 if Player 1 plays M, and R is not the best strategy for Player 2 if Player 1 plays U, M, or D.

Although there are no strictly dominant strategies, we can see that no matter what Player 1 does, Player 2 always gets a higher payoff from playing L than from playing R. We can therefore say that L **strictly dominates** R for Player 2, or that R is **strictly dominated** by L for Player 2. (Note that we cannot say that L is a **strictly dominant** strategy for Player 2—it does not dominate C—but we can say that R is a **strictly dominated** strategy for Player 2: an optimizing Player 2 would *never* play R.)

The second idea in the transition from dominant strategies to iterated dominance is similar to the backward induction idea of anticipating your opponents’ moves: players should recognize that other players have strictly dominated strategies, and should act accordingly. In our example, Player 1 should recognize that R is a strictly dominated strategy for Player 2, and therefore that there is no chance that Player 2 will play R. In effect, the game now looks like that shown in Figure 11.2 on the next page: the lines through the payoffs in the R column indicate that both players know that these payoffs have no chance of occurring because R is not a viable strategy for Player 2.

But now we see that Player 1 has an obvious strategy: given that Player 2 is never going to play R, Player 1 should always play M. Once R is out of the way, U and D are both dominated by M for Player 1: regardless of whether Player 2 plays L or C, Player 1 always gets his highest payoff by playing M. This is the

		Player 2		
		L	C	R
Player 1	U	1, 10	3, 20	40, 0
	M	10, 20	50, -10	6, 0
	D	2, 20	4, 40	10, 0

Figure 11.1: A game without dominant strategies

		Player 2		
		L	C	R
Player 1	U	1, 10	3, 20	40, 0
	M	10, 20	50, -10	6, 0
	D	2, 20	4, 40	10, 0

Figure 11.2: Eliminating R, which is strictly dominated by L for Player 2

idea of **iteration**, i.e., repetition. Combining this with the idea of dominated strategies gives us the process of **iterated dominance**: starting with the game in Figure 11.1, we look for a strictly dominated strategy; having found one (R), we eliminate it, giving us the game in Figure 11.2. We then repeat the process, looking for a strictly dominated strategy in that game; having found one (or, actually two: U and D), we eliminate them. A final iteration would yield (M, L) as a prediction for this game: knowing that Player 1 will always play M, Player 2 should always play L.

A complete example

Consider the game in Figure 11.3 below. There are no strictly dominant strategies, but there is a strictly dominated strategy: playing U is strictly dominated by D for Player 1. We can conclude that Player 1 will never play U, and so our game reduces to the matrix in Figure 11.4a on the next page.

But Player 2 should know that Player 1 will never play U, and if Player 1 never plays U then some of Player 2's strategies are strictly dominated! Namely, playing L and playing R are both strictly dominated by playing C as long as Player 1 never plays U. So we can eliminate those strategies for Player 2, yielding the matrix in Figure 11.4b. Finally, Player 1 should anticipate that Player 2 (anticipating that Player 1 will never play U) will never play L or R, and so Player 1 should conclude that M is strictly dominated by D (the matrix in Figure 11.4c). Using iterated strict dominance, then, we can predict that Player 1 will choose D and Player 2 will choose C.

		Player 2		
		L	C	R
Player 1	U	1,1	2,0	2,2
	M	0,3	1,5	4,4
	D	2,4	3,6	3,0

Figure 11.3: Iterated strict dominance example

	L	C	R		L	C	R		L	C	R
U	1/1	2/0	2/2	U	1/1	2/0	2/2	U	1/1	2/0	2/2
M	0,3	1,5	4,4	M	0/3	1,5	4/4	M	0/3	1/5	4/4
D	2,4	3,6	3,0	D	2/4	3,6	3/0	D	2/4	3,6	3/0
(a)				(b)				(c)			

Figure 11.4: Solution to iterated strict dominance example

Question: Does the order of elimination matter?

Answer: Although it is not obvious, the end result of iterated strict dominance is always the same regardless of the sequence of eliminations. In other words, if in some game you can either eliminate U for Player 1 or L for Player 2, you don't need to worry about which one to "do first": either way you'll end up at the same answer.

A side note here is that this result only holds under iterated **strict** dominance, according to which we eliminate a strategy only if there is some other strategy that yields payoffs that are *strictly higher* no matter what the other players do. If you eliminate a strategy when there is some other strategy that yields payoffs that are *higher or equal* no matter what the other players do, you are doing iterated **weak** dominance, and in this case you will *not* always get the same answer regardless of the sequence of eliminations. (For an example see problem 10.) This is a serious problem, and helps explain why we focus on iterated strict dominance.

11.2 Nash Equilibrium

Tenuous as it may seem, iterated strict dominance is not a very strong solution concept, meaning that it does not yield predictions in many games. An example is the game in Figure 11.5: there are no strictly dominant strategies and no strictly dominated strategies.

So game theorists have come up with other solution concepts. The most important one is called **Nash equilibrium** (abbreviated NE). A Nash equi-

		Player 2		
		L	C	R
Player 1	U	5,1	2,0	2,2
	M	0,4	1,5	4,5
	D	2,4	3,6	1,0

Figure 11.5: Nash equilibrium example

librium occurs when the strategies of the various players are best responses to each other. Equivalently but in other words: given the strategies of the other players, each player is acting optimally. Equivalently again: No player can gain by deviating alone, i.e., by changing his or her strategy single-handedly.

In the game in Figure 11.5, the strategies (D, C) form a Nash equilibrium: if Player 1 plays D, Player 2 gets her best payoff by playing C; and if Player 2 plays C, Player 1 gets his best payoff by playing D. So the players' strategies are best responses to each other; equivalently, no player can gain by deviating alone. (Question: Are there any other Nash equilibria in this game?)

Algorithms for Finding Nash Equilibria

The best way to identify the Nash equilibria of a game is to first identify all of the outcomes that are *not* Nash equilibria; anything left must be a Nash equilibrium. For example, consider the game in Figure 11.5. The strategy pair (U, L) is not a Nash equilibrium because Player 2 can gain by deviating alone to R; (U, C) is not a NE because Player 1 can gain by deviating alone to D (and Player 2 can gain by deviating alone to L or R); etc. If you go through the options one by one and cross out those that are *not* Nash equilibria, the remaining options *will* be Nash equilibria (See Figure 11.6a).

A shortcut (but one you should use carefully!) is to underline each player's best responses.¹ To apply this to the game in Figure 11.5, first assume that Player 2 plays L; Player 1's best response is to play U, so underline the "5" in the box corresponding to (U, L). Next assume that Player 2 plays C; Player 1's best response is to play D, so underline the "3" in the box corresponding to (D, C). Finally, assume that Player 2 plays R; Player 1's best response is to play M, so underline the "4" in the box corresponding to (M, R). Now do the same thing for Player 2: go through all of Player 1's options and underline the best response for Player 2. (Note that C and R are both best responses when Player 1 plays M!) We end up with Figure 11.6b: the only boxes with both payoffs underlined are (D, C) and (M, R), the Nash equilibria of the game.

	L	C	R
U	5 / 1	2 / 0	2 / 2
M	0 / 4	1 / 5	4,5
D	2 / 4	3,6	1 / 0

(a)

	L	C	R
U	<u>5</u> ,1	2,0	2, <u>2</u>
M	0,4	1, <u>5</u>	<u>4</u> , <u>5</u>
D	2,4	<u>3</u> , <u>6</u>	1,0

(b)

Figure 11.6: Finding Nash equilibria: (a) with strike-outs; (b) with underlinings

¹It is easy to confuse the rows and columns and end up underlining the wrong things. Always double-check your answers by confirming that no player can gain by deviating alone.

Some History

Nash equilibrium is one of the fundamental concepts of game theory. It is named after John Nash, a mathematician born in the early part of this century. He came up with his equilibrium concept while getting his Ph.D. in mathematics at Princeton, then got a professorship at MIT, then went mad (e.g., claimed that aliens were sending him coded messages on the front page of the New York Times), then spent many years in and out of various mental institutions, then slowly got on the road to recovery, then won the Nobel Prize in Economics in 1994, and now putters around Princeton playing with computers. You can read more about him in a fun book called *A Beautiful Mind* by Sylvia Nasar.²

11.3 Infinitely Repeated Games

We saw in the last chapter that there's no potential for cooperation (at least in theory) if we play the Prisoner's Dilemma game twice, or 50 times, or 50 million times. What about infinitely many times?

In order to examine this possibility, we must first figure out exactly what it means to win (or lose) this game infinitely many times. Here it helps to use the present value concepts from Chapter 1: with an interest rate of 5%, winning \$1 in each round does *not* give you infinite winnings. Rather, the present value of your winnings (using the perpetuity formula, assuming you get paid at the end of each round) is $\frac{\$1}{.05} = \20 .

So: with an interest rate of r we can ask meaningful questions about the potential for cooperation. One point that is immediately clear is that there is still plenty of potential for *non-cooperation*: the strategies of playing (D, D) forever continue to constitute a Nash equilibrium of this game.

But perhaps there are other strategies that are also Nash equilibria. Because the game is played infinitely many times, we cannot use backward induction to solve this game. Instead, we need to hunt around and look for strategies that might yield a cooperative Nash equilibrium.

One potentially attractive idea is to use a **trigger strategy**: begin by cooperating and assuming that the other player will cooperate (i.e., that both players will play C), and enforce cooperation by threatening to return to the (D, D) equilibrium. Formally, the trigger strategy for each player is as follows: In the first stage, play C. Thereafter, if (C, C) has been the result in all previous stages, play C; otherwise, play D.

We can see that the cooperative outcome (C, C) will be the outcome in each stage game if both players adopt such a trigger strategy. But do these strategies constitute a Nash equilibrium? To check this, we have to see if the strategies are best responses to each other. In other words, given that Player 1 adopts

²There is also a movie of the same name, starring Russell Crowe. Unfortunately, it takes some liberties with the truth; it also does a lousy job of describing the Nash equilibrium concept.

the trigger strategy above, is it optimal for Player 2 to adopt a similar trigger strategy, or does Player 2 have an incentive to take advantage of Player 1?

To find out, let's examine Player 2's payoffs from cooperating and from deviating:

If Player 2 cooperates, she can expect to gain \$1 at the end of each round, yielding a present value payoff of $\frac{\$1}{r}$. (If $r = .05$ this turns out to be \$20.)

If Player 2 tries to cheat Player 1 (e.g., by playing D in the first round), Player 2 can anticipate that Player 1 will play D thereafter, so the best response for Player 2 is to play D thereafter as well. So the best deviation strategy for Player 2 is to play D in the first round (yielding a payoff of \$10 since Player 1 plays C) and D thereafter (yielding a payoff of \$0 each round since Player 1 plays D also). The present value of all this is simply \$10.

We can now compare these two payoffs, and we can see that cooperating is a best response for Player 2 as long as $\frac{\$1}{r} \geq 10$. Since the game is symmetric, cooperating is a best response for Player 1 under same condition, so we have a Nash equilibrium (i.e., mutual best responses) as long as $\frac{\$1}{r} \geq 10$. Solving this yields a critical value of $r = .1$. When r is below this value (i.e., the interest rate is less than 10%), cooperation is possible. When r is above this value (i.e., the interest rate is greater than 10%), cheating is too tempting and the trigger strategies do not form a Nash equilibrium. The intuition here is quite nice: By cooperating instead of deviating, Player 2 accepts lower payoffs now (1 instead of 10) in order to benefit from higher payoffs later (1 instead of 0). Higher interest rates make the future less important, meaning that Player 2 benefits less by incurring losses today in exchange for gains tomorrow. With sufficiently high interest rates, Player 2 will take the money and run; but so will Player 1!

11.4 Mixed Strategies

Figure 11.7 shows another game, called the Battle of the Sexes. In this game, Player 1 prefers the opera, and Player 2 prefers wrestling, but what

		Player 2	
		Opera	WWF
Player 1	Opera	2,1	0,0
	WWF	0,0	1,2

Figure 11.7: The battle of the sexes

both players really want above all is to be with each other. They both choose simultaneously, though, and so cannot guarantee that they'll end up together. (Imagine, for example, that they are at different work places and can't reach each other and must simply head to one of the two events after work and wait for the other person at will-call.)

The Nash equilibriums of this game are (Opera, Opera) and (WWF, WWF). But there is another Nash equilibrium that is perhaps a little better at predicting reality: that equilibrium is for both players to play a **mixed strategy**, i.e., to choose different strategies with various probabilities. (In this case, the mixed strategy equilibrium is for Player 1 to choose opera with probability $2/3$ and WWF with probability $1/3$, and for Player 2 to choose opera with probability $1/3$ and WWF with probability $2/3$. You should be able to use what you've learned about expected value to show that these are mutual best responses.) One of the main results from game theory is that every finite game has at least one Nash equilibrium. That Nash equilibrium may only exist in mixed strategies, as in the following example.

Example: (Matching pennies, Figure 11.8). Players 1 and 2 each have a penny, and they put their pennies on a table simultaneously. If both show the same face (both heads or both tails), Player 2 must pay \$1 to Player 1; if one is heads and the other is tails, Player 1 must pay \$1 to Player 2.

Not surprisingly, the only NE in this game is for each player to play heads with probably $1/2$ and tails with probability $1/2$.

11.5 Math: Mixed Strategies

Consider the “Matching Pennies” game shown in Figure 11.8. There are no pure strategy Nash equilibria in this game, but intuitively it seems like randomizing between heads and tails (with probability 50% for each) might be a good strategy. To formalize this intuition we introduce the concept of **mixed strategy Nash equilibrium**.

In a mixed strategy Nash equilibrium, players do not have to choose just one strategy (say, Heads) and play it with probability 1. Instead, they can specify probabilities for all of their different options and then randomize (or mix) between them. To see how this might work in practice, a player who specifies Heads with probability .3 and Tails with probability .7 could put 3

		Player 2	
		Heads	Tails
Player 1	Heads	1,-1	-1,1
	Tails	-1,1	1,-1

Figure 11.8: Matching pennies

cards labeled Heads and 7 cards labeled Tails into a hat; when the times comes to actually play the game, she draws a card from the hat and plays accordingly. She may only play the game once, but her *odds* of playing Heads or Tails are .3 and .7, respectively.

Finding Mixed Strategy Nash Equilibria

To find mixed strategy Nash equilibria, we can simply associate different probabilities with the different options for each player. This gets messy for big payoff matrices, so we will restrict our attention to games (such as Matching Pennies) in which each player has only two options. In that game, let us define p to be the probability that player 1 chooses Heads and q to be the probability that player 2 chooses Heads. Since probabilities have to add up to 1, the probability that players 1 and 2 choose Tails must be $1 - p$ and $1 - q$, respectively.

Now let's write down the expected payoff for player 1 given these strategies. With probability p player 1 chooses Heads, in which case he gets +1 if player 2 chooses Heads (which happens with probability q) and -1 if player 2 chooses Tails (which happens with probability $1 - q$). With probability $1 - p$ player 1 chooses Tails, in which case he gets -1 if player 2 chooses Heads (which happens with probability q) and $+1$ if player 2 chooses Tails (which happens with probability $1 - q$). So player 1's expected value is

$$\begin{aligned} E(\pi_1) &= p[q(1) + (1 - q)(-1)] + (1 - p)[q(-1) + (1 - q)(1)] \\ &= p(2q - 1) + (1 - p)(1 - 2q). \end{aligned}$$

Similarly, player 2's expected payoff is

$$\begin{aligned} E(\pi_2) &= q[p(-1) + (1 - p)(1)] + (1 - q)[p(1) + (1 - p)(-1)] \\ &= q(1 - 2p) + (1 - q)(2p - 1). \end{aligned}$$

Now, we want to find p and q that form a Nash equilibrium, i.e., that are mutual best responses. To do this, we take partial derivatives and set them equal to zero. Here's why:

First, player 1 wants to choose p to maximize $E(\pi_1) = p(2q - 1) + (1 - p)(1 - 2q)$. One possibility is that a maximizing value of p is a corner solution, i.e., $p = 0$ or $p = 1$. These are player 1's pure strategy options: $p = 1$ means that player 1 always plays Heads, and $p = 0$ means that player 1 always plays Tails.

The other possibility is that there is an interior maximum, i.e., a maximum value of p with $0 < p < 1$. In this case, the partial derivative of $E(\pi_1)$ with respect to p must be zero:

$$\frac{\partial E(\pi_1)}{\partial p} = 0 \implies 2q - 1 - (1 - 2q) = 0 \implies 4q = 2 \implies q = \frac{1}{2}.$$

This tells us that *any* interior value of p is a candidate maximum as long as $q = \frac{1}{2}$. Mathematically, this makes sense because if $q = \frac{1}{2}$ then player 1's

expected payoff (no matter what his choice of p) is always

$$E(\pi_1) = p(2q - 1) + (1 - p)(1 - 2q) = p(0) + (1 - p)(0) = 0.$$

Intuitively, what is happening is that player 2 is randomly choosing between Heads and Tails. As player 1, any strategy you follow is a best response. If you always play Heads, you will get an expected payoff of 0; if you always play Tails, you will get an expected payoff of 0; if you play heads with probability .5 or .3, you will get an expected payoff of 0.

Our conclusion regarding player 1's strategy, then, is this: *If* player 2 chooses $q = \frac{1}{2}$, i.e., randomizes between Heads and Tails, then any choice of p is a best response for player 1. But *if* player 2 chooses $q \neq \frac{1}{2}$, then player 1's best response is a pure strategy: if player 2 chooses $q > \frac{1}{2}$ then player 1's best response is to always play Heads; if player 2 chooses $q < \frac{1}{2}$ then player 1's best response is to always play Tails.

We can now do the math for player 2 and come up with a similar conclusion. Player 2's expected payoff is $E(\pi_2) = q(1 - 2p) + (1 - q)(2p - 1)$. Any value of q that maximizes this is either a corner solution (i.e., one of the pure strategies $q = 1$ or $q = 0$) or an interior solution with $0 < q < 1$, in which case

$$\frac{\partial E(\pi_2)}{\partial q} = 0 \implies 1 - 2p - (2p - 1) = 0 \implies 4p = 2 \implies p = \frac{1}{2}.$$

So if player 1 chooses $p = \frac{1}{2}$ then any choice of q is a best response for player 2. But if player 1 chooses $p \neq \frac{1}{2}$, then player 2's best response is a pure strategy: if player 1 chooses $p > \frac{1}{2}$ then player 2's best response is to always play Tails; if player 1 chooses $p < \frac{1}{2}$ then player 2's best response is to always play Heads.

Now we can put our results together to find the Nash equilibrium in this game. If player 1's choice of p is a best response to player 2's choice of q then either $p = 1$ or $p = 0$ or $q = \frac{1}{2}$ (in which case any p is a best response). And if player 2's choice of q is a best response to player 1's choice of p then either $q = 1$ or $q = 0$ or $p = \frac{1}{2}$ (in which case any q is a best response).

Three choices for player 1 and three choices for player 2 combine to give us nine candidate Nash equilibria:

Four pure strategy candidates: $(p = 1, q = 1), (p = 1, q = 0), (p = 0, q = 1), (p = 0, q = 0)$.

One mixed strategy candidate: $(0 < p < 1, 0 < q < 1)$.

Four pure/mixed combinations: $(p = 1, 0 < q < 1), (p = 0, 0 < q < 1), (0 < p < 1, q = 1), (0 < p < 1, q = 0)$.

We can see from the payoff matrix that the four pure strategy candidates are not mutual best responses, i.e., are not Nash equilibria. And we can quickly see that the four pure/mixed combinations are also not best responses; for example, $(p = 1, 0 < q < 1)$ is not a Nash equilibrium because if player 1 chooses $p = 1$ then player 2's best response is to choose $q = 0$, not $0 < q < 1$.

But the mixed strategy candidate does yield a Nash equilibrium: player 1's choice of $0 < p < 1$ is a best response as long as $q = \frac{1}{2}$. And player 2's choice of $0 < q < 1$ is a best response as long as $p = \frac{1}{2}$. So the players' strategies are mutual best responses if $p = q = \frac{1}{2}$. This is the mixed strategy Nash equilibrium of this game.

Another Example

Consider the “Battle of the Sexes” game shown in Figure 11.7 and duplicated below. Again, let p be the probability that player 1 chooses Opera and q be the probability that player 2 chooses Opera (so that $1 - p$ and $1 - q$ are the respective probabilities that players 1 and 2 will choose WWF). Then player 1's expected payoff is

$$\begin{aligned} E(\pi_1) &= p[q(2) + (1 - q)(0)] + (1 - p)[q(0) + (1 - q)(1)] \\ &= 2pq + (1 - p)(1 - q). \end{aligned}$$

Similarly, player 2's expected payoff is

$$\begin{aligned} E(\pi_2) &= q[p(1) + (1 - p)(0)] + (1 - q)[p(0) + (1 - p)(2)] \\ &= pq + (1 - q)(2)(1 - p). \end{aligned}$$

Now, we want to find p and q that form a Nash equilibrium, i.e., that are mutual best responses. To do this, we take partial derivatives and set them equal to zero.

So: player 1 wants to choose p to maximize $E(\pi_1) = 2pq + (1 - p)(1 - q)$. Any value of p that maximizes this is either a corner solution (i.e., one of the pure strategies $p = 1$ or $p = 0$) or an interior solution with $0 < p < 1$, in which case the partial derivative of $E(\pi_1)$ with respect to p must be zero:

$$\frac{\partial E(\pi_1)}{\partial p} = 0 \implies 2q - (1 - q) = 0 \implies 3q = 1 \implies q = \frac{1}{3}.$$

This tells us that *any* interior value of p is a candidate maximum as long as $q = \frac{1}{3}$. Mathematically, this makes sense because if $q = \frac{1}{3}$ then player 1's expected payoff (no matter what his choice of p) is always

$$E(\pi_1) = 2pq + (1 - p)(1 - q) = \frac{2}{3}p + \frac{2}{3}(1 - p) = \frac{2}{3}.$$

		Player 2	
		Opera	WWF
Player 1	Opera	2,1	0,0
	WWF	0,0	1,2

Figure 11.9: The battle of the sexes

Our conclusion regarding player 1's strategy, then, is this: *If* player 2 chooses $q = \frac{1}{3}$, then any choice of p is a best response for player 1. But *if* player 2 chooses $q \neq \frac{1}{3}$, then player 1's best response is a pure strategy: if player 2 chooses $q > \frac{1}{3}$ then player 1's best response is to always play Opera; if player 2 chooses $q < \frac{1}{3}$ then player 1's best response is to always play WWF.

We can now do the math for player 2 and come up with a similar conclusion. Player 2's expected payoff is $E(\pi_2) = pq + (1-q)(2)(1-p)$. Any value of q that maximizes this is either a corner solution (i.e., one of the pure strategies $q = 1$ or $q = 0$) or an interior solution with $0 < q < 1$, in which case

$$\frac{\partial E(\pi_2)}{\partial q} = 0 \implies p - 2(1-p) = 0 \implies 3p = 2 \implies p = \frac{2}{3}.$$

So if player 1 chooses $p = \frac{2}{3}$ then any choice of q is a best response for player 2. But if player 1 chooses $p \neq \frac{2}{3}$, then player 2's best response is a pure strategy: if player 1 chooses $p > \frac{2}{3}$ then player 2's best response is to always play Opera; if player 1 chooses $p < \frac{2}{3}$ then player 2's best response is to always play WWF.

Now we can put our results together to find the Nash equilibrium in this game. If player 1's choice of p is a best response to player 2's choice of q then either $p = 1$ or $p = 0$ or $q = \frac{1}{3}$ (in which case any p is a best response). And if player 2's choice of q is a best response to player 1's choice of p then either $q = 1$ or $q = 0$ or $p = \frac{2}{3}$ (in which case any q is a best response).

Three choices for player 1 and three choices for player 2 combine to give us nine candidate Nash equilibria:

Four pure strategy candidates: $(p = 1, q = 1), (p = 1, q = 0), (p = 0, q = 1), (p = 0, q = 0)$.

One mixed strategy candidate: $(0 < p < 1, 0 < q < 1)$.

Four pure/mixed combinations: $(p = 1, 0 < q < 1), (p = 0, 0 < q < 1), (0 < p < 1, q = 1), (0 < p < 1, q = 0)$.

We can see from the payoff matrix that there are two Nash equilibria among the four pure strategy candidates: $(p = 1, q = 1)$ and $(p = 0, q = 0)$. The other two are not Nash equilibria. We can also see that the four pure/mixed combinations are not best responses; for example, $(p = 1, 0 < q < 1)$ is not a Nash equilibrium because if player 1 chooses $p = 1$ then player 2's best response is to choose $q = 1$, not $0 < q < 1$.

But the mixed strategy candidate does yield a Nash equilibrium: player 1's choice of $0 < p < 1$ is a best response as long as $q = \frac{1}{3}$. And player 2's choice of $0 < q < 1$ is a best response as long as $p = \frac{2}{3}$. So the players' strategies are mutual best responses if $p = \frac{2}{3}$ and $q = \frac{1}{3}$. This is the mixed strategy Nash equilibrium of this game.

So this game has three Nash equilibria: two in pure strategies and one in mixed strategies.

Problems

1. *Challenge.* Explain (as if to a non-economist) why iterated dominance make sense.
2. *Super Challenge.* Explain (as if to a non-economist) why Nash equilibrium makes sense.
3. Show that there are no strictly dominant strategies in the game in Figure 11.3 on page 107.
4. Analyze games (a) through (e) on the following page(s). First see how far you can get using iterated dominance. Then find the Nash equilibrium(s). If you can identify a unique outcome, determine whether it is Pareto efficient. If it is not, identify a Pareto improvement.
5. The game “Rock, Paper, Scissors” works as follows: You and your opponent simultaneously choose rock, paper, or scissors. If you pick the same one (e.g., if you both pick rock), you both get zero. Otherwise, rock beats scissors, scissors beats paper, and paper beats rock, and the loser must pay the winner \$1.
 - (a) Write down the payoff matrix for this game.
 - (b) Does iterated dominance help you solve this game?
 - (c) *Calculus/Challenge.* Can you find any *mixed strategy* Nash equilibria?
6. *Challenge.* Prove that the pure strategy Nash equilibrium solutions are a subset of the iterated dominance solutions, i.e., that iterated dominance never eliminates any pure strategy Nash equilibrium solutions.
7. Rewrite Story #1 from Overinvestment Game (from problem 3 in Chapter 8) as a simultaneous move game and identify the (pure strategy) Nash equilibria. Does your answer suggest anything about the relationship between backward induction and Nash equilibrium?
8. *Challenge.* Prove that backward induction solutions are a subset of Nash equilibrium solutions, i.e., that any backward induction solution is also a Nash equilibrium solution. (Note: Backward induction is in fact a refinement of Nash equilibrium called “subgame perfect Nash equilibrium”.)
9. *Fun/Challenge.* Section 11.3 describes a trigger strategy for yielding cooperating in the infinitely repeated Prisoner’s Dilemma game shown in figure 10.3. Can you think of another strategy that yields even higher payoffs for the players? Can you show that it’s a Nash equilibrium?

(a)

		Player 2		
		L	C	R
Player 1	U	0,3	2,1	5,0
	M	4,8	3,2	8,3
	D	3,7	6,3	6,8

(b)

		Player 2		
		L	C	R
Player 1	U	-1,4	7,3	5,2
	M	2,0	5,-1	6,2
	D	1,2	1,0	1,0

(c)

		Player 2		
		L	C	R
Player 1	U	1,0	7,3	2,1
	M	1,0	1,2	6,2
	D	1,2	1,-3	1,0

(d)

		Player 2		
		L	C	R
Player 1	U	3,-1	5,4	3,2
	M	-2,5	1,3	2,1
	D	3,3	3,6	3,0

(e)

		Player 2		
		L	C	R
Player 1	U	3,-1	1,0	-1,-1
	M	1,-5	6,3	-7,-5
	D	-8,-10	-1,-3	-1,-1

10. *Challenge.* The end of the section on iterated dominance mentioned the dangers of iterated weak dominance, namely that different sequences of elimination can yield different predictions for the outcome of a game. Show this using the game in Figure 11.10. (Hint: Note that U is weakly dominated by M for Player 1 and that M is weakly dominated by D for Player 1.)

		Player 2	
		L	R
Player 1	U	50, 10	6, 20
	M	50, 10	8, 9
	D	60, 15	8, 15

Figure 11.10: The dangers of iterated weak dominance

Calculus Problems

- C-1. Find all Nash equilibria (pure and mixed) in the game shown in figure 11.11. (Use p as the probability that Player 1 plays U and q as the probability that Player 2 plays L.)

		Player 2	
		L	R
Player 1	U	1, 3	0, 0
	D	0, 0	3, 1

Figure 11.11: A game with a mixed strategy equilibrium

- C-2. Find all Nash equilibria (pure and mixed) in the game shown in figure 11.12. (Use p as the probability that Player 1 plays U and q as the probability that Player 2 plays L.)

		Player 2	
		L	R
Player 1	U	0, 0	-1, 5
	D	-2, 1	1, -2

Figure 11.12: Another game with a mixed strategy equilibrium

Chapter 12

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?

Answer²: 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.)

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

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

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

12.1 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](http://www.vba.nl).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

³<http://www.vba.nl>

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

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!

12.2 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

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 12.1: eBay's Proxy Bidding feature

with the highest sit-down price would win and would pay an amount essentially equal to the *second-highest* sit-down price. (If the highest sit-down price was \$100 and the second-highest was \$20, everybody except for the highest bidder would drop out at a price of about \$20, at which point the auction would be over.⁴) For further supporting evidence, Figure 12.1 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

⁴The 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>

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

We can now return to one of our earlier questions:

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 at 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**.

12.3 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 12.2 and the accompanying text, both excerpted from a working paper⁷ on the 1996-1997 FCC auction:

[Figure 12.2] 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.

⁶A 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.

⁷Peter 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 12.2: 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 $\text{Value}(\text{Winning})$ to denote the value of winning the auction. Assume that the value of losing the auction is zero.
 - (c) Write down an expression for the expected value of bidding $\$x$ in a first-price sealed bid auction. Assume that your gain or "profit" from winning an auction is the difference between your true value for the item and the price you actually have to pay for it. Can you use this expected value expression to highlight the issues faced by a bidder in such an auction? For example, can you show mathematically why bidders should shade their bids?

- (d) Write down an expression for the expected value of bidding x in a second-price sealed bid auction. (Again, assume that your gain or “profit” from winning an auction is the difference between your true value for the item and the price you actually have to pay for it.) Can you use this expected value expression to highlight the issues faced by a bidder in such an auction? For example, can you show mathematically why bidders should bid their true value?

Chapter 13

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 2¹

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

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

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

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

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

caught by a fisherman trying to recover lost equipment off the coast of Georgia. The fishery rapidly expanded, with catch 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.

13.1 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.

13.2 A Brief Look at Government Intervention

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

forever. (See Figure 3.3 on page 36.) 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

²For obvious reasons, the establishment of Exclusive Economic Zones greatly displeased foreigners who were used to fishing in these areas. Iceland's efforts led to three "Cod Wars" with Great Britain, the most serious of which (the Third Cod War, in 1975–76) featured a number of boat ramming and the deployment of part of the British Navy.

as soon as the season began, regardless of the weather. This threatened the safety of vessels and the crews aboard them.

Wasted fish Fishing in stormy weather also led to lost gear: in 1990 lost gear killed an estimated 2 million pounds of halibut in the Alaska fishery. A Pareto improvement would be to end the race for fish, or at least ground all vessels during bad weather. In fisheries that used nets, another problem was compaction: so many fish were being caught in each haul that the fish in the bottom of the nets would get squashed. Again, a Pareto improvement would be to end the race for fish; fishing boats could then afford to pull their nets up more often, reducing compaction.

Overcapitalization in the fishing industry There was a tremendous overinvestment in fishing capital (e.g., fishing boats, automatic hook baiters and other fancy fishing gear, etc.). After the TAC limit triggered the close of the fishing season, the boats would return to dock, often to sit dormant until the next season. A Pareto improvement would be to end the race for fish; given a longer fishing season, the same number of fish could be caught with less capital.

Overcapitalization in the processing industry With so many fish being caught in such short time frames, the race for fish led to overinvestment in fish processing. As with fishing boats, factories for processing and freezing fish would sit unused for much of the year, springing to life only during the short fishing season. Again, ending the race for fish would lead to a Pareto improvement: given a longer fishing season, the same number of fish could be processed with less capital.

Frozen fish Consumers prefer fresh fish, but short fishing seasons meant that they had to settle for frozen fish during much of the year.

Put it all together and the result was a **dissipation of resource rents**. In English, this means that human beings were not receiving the maximum benefit from this natural resource. Since consumers are willing to pay a lot of money for fish that nature produces “for free”, the potential exists for somebody to get a great deal: consumers should get highly valued fish for low prices, or fishermen should make large profits, or both. In other words, the potential exists for somebody to capture the resource rents associated with the fishery. But the race for fish squandered much of this potential bounty. Even though consumers placed much higher values on fresh fish, they had to settle for frozen fish for much of the year. Fishermen weren’t getting ahead either: they had to spend so much money on capital (bigger fishing boats with fancy technology) that they didn’t make much profit. In short, everybody lost.

It is important to note that the race for fish is not inevitable. The Coase Theorem says that there are always incentives for individuals to bargain their way out of Pareto inefficient situations, and in some fisheries this happened. Most notably, factory trawlers in the Pacific whiting fishery avoided the race for fish by forming a co-operative that imposed limits on each vessel and put

observers on each boat to make sure the limits were obeyed. Although participation in the co-operative was voluntary, each owner knew that breaking the rules would result in an end of the co-operative and a return of the (unprofitable) race for fish the following season. In a repeated game, this **trigger strategy** gives each firm an incentive to participate in the co-operative.

In short, the factory trawlers avoided the race for fish by colluding to fix harvest amounts. Its success, however, depended on the unique situation in the whiting fishery: with only three firms in the market and a government-mandated moratorium on entry by other firms, the conditions for successful collusion were perfect. As in normal markets, however, collusion between fishermen becomes more difficult as the number of fishermen increases, and is extremely difficult when there is free entry into the fishery. These impediments to bargaining help explain why fishermen in most fisheries were collectively unable to stop the race for fish.

13.3 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 7, 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 14

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-

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

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.

14.1 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 6 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**.

14.2 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 15

Supply and Demand Basics

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

– Shel Silverstein, “Mr. Smeds and Mr. Spats”
A Light in the Attic (1981)

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 15.1a 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 19.)

The figure also shows an area representing their product, pq , which is **total revenue** and also **total expenditure**. (If the market price p is \$2.00 per unit and the quantity transacted q is 1,000 units, then the total amount of money that changes hands—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




Figure 15.1: (a) the observables, p and q , and $pq = \text{total revenue (and/or total expenditure)}$; (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.

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 15.1b.

15.1 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 15.2.) When the dot moves around, it is because these curves are moving.

The **market supply curve** (see Figure 15.2a) is a graphical presentation of *hypothetical* questions like these: *If* the market price were \$6 per unit, how many units of this good would sellers want to sell? *If* the market price were \$7 per unit, how many units of this good would sellers want to sell? Put all these hypothetical questions together, and you get the market supply curve, which answers the following question: *If* the market price were p per unit, how many units of this good would sellers want to sell?

Similarly, the **market demand curve** (see Figure 15.2b) is a graphical presentation of *hypothetical* questions like these: *If* the market price were \$6 per unit, how many units of this good would buyers want to buy? *If* the market price were \$7 per unit, how many units of this good would buyers want to buy? Put all these hypothetical questions together, and you get the market demand curve, which answers the following question: *If* the market price were p per unit, how many units of this good would buyers want to buy?

Note from Figure 15.2 that supply and demand curves move in opposite




Figure 15.2: Supply and demand

directions. The supply curve is *upward sloping*, indicating that as the price per unit rises sellers want to sell more units (or, equivalently, that as the price falls sellers want to sell fewer units). The demand curve is *downward sloping*, indicating that as the price per unit rises buyers want to buy fewer units (or, equivalently, that as the price falls buyers want to buy more units).¹ We will examine these slopes in more detail in Chapter 18.

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 15.2c. 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

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

15.2 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 15.3:

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 15.3a.)

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 15.3b.)

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 15.3c.)

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 15.3d.)

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 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




Figure 15.3: 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.

transacted Q on the x -axis and the market price P on the y -axis. Chapter 19 returns to this issue.

15.3 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:

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

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

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

When you answered the previous question, you probably assumed that that oil would become available immediately, i.e., that oil companies could immediately begin extracting and selling that oil. (I wanted you to assume that, so do not go back and rethink your answer above!) It turns out that life is more complicated than that: it takes time to build pipelines and to drill oil wells into pristine Arctic wilderness, so any oil that comes from ANWR will not reach the market for something like 5 years. This fact became a source of contention during the presidential campaign, with Al Gore arguing that opening ANWR would have no effect on current gasoline prices because of this 5-year time lag, and George W. Bush arguing. . . well, I don't remember what his argument was, but it probably had something to do with 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 16

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 16.1.

16.1 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 16.1: 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 16.1 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 16.1, 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 16.2.

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 16.2 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 19.

16.2 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 16.2: 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 16.1 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 16.1, 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 16.3.

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




Figure 16.3: Effect of a \$.40 per gallon tax on the buyers

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 16.3 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 19.

16.3 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.

¹Note, 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.

16.4 *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

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

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

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

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

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

Similarly, if there's a tax on the sellers of 20% and the market price is p , what sellers really get is $100 - 20 = 80\%$ of p , i.e., $.8p$; if the before-tax supply curve is $q = 6000 + 800p$, the after-tax supply curve will be

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

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 16.4 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 16.4: 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?

Algebra Problems

- A-1. Consider a market with demand curve $q = 16 - 10p$ and supply curve $q = -8 + 20p$. (Here q is in millions of pounds per day and p is in dollars per pound.

- (a) Determine the market equilibrium price and quantity and the total revenue in this market.
- (b) Calculate the price elasticity of demand and the price elasticity of supply at the market equilibrium.
- (c) Imagine that the government imposes a \$.60 per-unit tax on the buyers. Write down the new market supply and demand curves, and find the new market equilibrium price and quantity. How much of the tax burden is borne by the buyers, and how much by the sellers? Calculate the ratio of these tax burdens, and compare with the ratio of the elasticities calculated above.
- (d) Now imagine that the government instead decides to impose a \$.60 per-unit tax on the sellers. How will this change things? Write down the new market supply and demand curves, find the new market equilibrium price and quantity, and compare with your answer from above (where the tax is on the buyer).
- (e) Now imagine that the government instead decides to impose a sales tax of 50% on the sellers. How will this change things? Write down the new market supply and demand curves and find the new market equilibrium price and quantity.
- (f) Now imagine that the government instead decides to impose a 100% sales tax on the buyers. How will this change things? Write down the new market supply and demand curves, find the new market equilibrium price and quantity, and compare with your answer from above (where the tax is on the sellers).

A-2. Consider a market with demand curve $q = 500 - 20p$ and supply curve $q = 50 + 25p$.

- (a) What is the original market equilibrium price and quantity?
- (b) What do the demand and supply curves look like with a \$2 per-unit tax on the buyers?
- (c) With a \$2 per-unit tax on the sellers?
- (d) With a 20% sales tax on the buyers?
- (e) With a 20% sales tax on the sellers?

A-3. Consider a world with market demand curve $q = 50 - 6p$ and market supply curve $q = 20p - 28$.

- (a) What is the market equilibrium price and quantity?
- (b) How would the equations for the supply and demand curves change if the government imposed a tax of \$.50 per unit on the buyers?
- (c) How would the equations for the supply and demand curves change if the government imposed a sales tax of 20% on the sellers?

Chapter 17

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 17.1a 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 17.1b shows a “steep” demand curve, one which appears to be very insensitive to price




Figure 17.1: (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 17.1a: buyers appear to be very sensitive to price changes. But graph the demand curve for gasoline in pennies and it will look like Figure 17.1b: 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.

17.1 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 17.2: 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 17.2. (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 them¹:

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

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.

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

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

17.2 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 17.2 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 17.3a 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 17.3b shows a supply curve that is perfectly




Figure 17.3: (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 6.) 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 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).

17.3 Math: Elasticities and Calculus

Without calculus, the price elasticity of demand at point A is defined as

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

The ratio $\frac{\Delta q}{\Delta p}$ approximates the slope of the demand curve, so we can use calculus to take the limit as Δp and Δq approach zero. The result is a more robust definition of the price elasticity of demand at point A :

$$\varepsilon(A) = \frac{dq}{dp} \cdot \frac{p_A}{q_A}.$$

For example, if the demand curve is given by $q = 50 - 2p$, the slope of the demand curve is $\frac{dq}{dp} = -2$. Then the price elasticity of demand at the point $p = 10, q = 30$ is $\varepsilon = -2 \frac{10}{30} = -\frac{2}{3}$ and the price elasticity of demand at the point $p = 20, q = 10$ is $\varepsilon = -2 \frac{20}{10} = -4$.

Similarly, if the demand curve is given by $q = 1000 - 2p^2$, the slope of the demand curve is $\frac{dq}{dp} = -4p$. Then the price elasticity of demand at the point $p = 10, q = 800$ is $\varepsilon = -4(10) \frac{10}{800} = -\frac{1}{2}$ and the price elasticity of demand at the point $p = 20, q = 200$ is $\varepsilon = -4(20) \frac{20}{200} = -8$.

We can use calculus to get a more definitive result about the relationship between total revenue and the price elasticity of demand. Since total revenue is $TR = pq$ and the demand curve tells us that q is a function of p , we can use the product rule to calculate the effect on total revenue of an increase in price:

$$\frac{dTR}{dp} = \frac{d(pq)}{dp} = p \left[\frac{d}{dp}(q) \right] + q \left[\frac{d}{dp}(p) \right] = p \frac{dq}{dp} + q(1) = q \left(\frac{p}{q} \frac{dq}{dp} + 1 \right) = q(\varepsilon + 1).$$

Since q is always positive, the sign of $\frac{dTR}{dp}$ depends on the sign of $\varepsilon + 1$. If demand is elastic, i.e., if the price elasticity of demand is $\varepsilon < -1$, then a price

increase will *reduce* total revenue; if demand is inelastic, i.e., if the price elasticity of demand is $\varepsilon > -1$, then a price increase will *increase* total revenue.

Other elasticities can be similarly defined in terms of derivatives. For example, the **cross-price elasticity** measures the impact of a change in the price of good 2 on the quantity demanded of good 1:

$$\varepsilon_{12}(A) = \frac{\% \text{ change in } q_1}{\% \text{ change in } p_2} = \frac{\frac{\Delta q_1}{q_{1A}}}{\frac{\Delta p_2}{p_{2A}}} = \frac{\Delta q_1}{\Delta p_2} \cdot \frac{p_{2A}}{q_{1A}} = \frac{dq_1}{dp_2} \cdot \frac{p_{2A}}{q_{1A}}.$$

For example, if the demand curve for good 1 is given by $q_1 = 50 - 2p_1 - .5p_2$, the partial derivative of the demand curve with respect to p_2 is $\frac{\partial q_1}{\partial p_2} = -.5$. Then the cross-price elasticity at the point $p_1 = 8, p_2 = 10, q_1 = 50 - 2(8) - .5(10) = 29$ is $\varepsilon = -.5 \frac{10}{29} = -\frac{5}{29}$.

Cross-price elasticities turn out to be useful for nailing down a concrete definition for substitutes and complements. Good 2 is a **substitute** for good 1 if an increase in the price of good 2 *increases* demand for good 1, i.e., if the cross-price elasticity is positive. Similarly, good 2 is a **complement** for good 1 if an increase in the price of good 2 *decreases* demand for good 1, i.e., if the cross-price elasticity is negative.

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. Do you think demand and supply curves are more elastic (i.e., more responsive to price changes) in the short run or in the long run? Why? How does your result translate graphically in terms of the steepness or flatness of the various curves?
3. Figure 17.4 shows a hypothetical demand curve for oranges.




Figure 17.4: 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 point X? At point Y? At point 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 16. 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 18

Supply and Demand Details

The great thing about supply and demand is that the basic ideas are relatively easy and extremely powerful. The previous chapters analyze a wide range of situations using only three concepts: market supply curves, market demand curves, and market equilibrium.

This chapter looks at some details of supply and demand. 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?

18.1 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 price of \$6, I would want to sell 4”) gives us *one point* on that firm’s individual supply curve.




Figure 18.1: Individual supply and demand curves

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 18.1) look just like market supply curves (see Figure 15.2c). 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.

18.2 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 18.2 and Figure 18.2), and over all the individual buyers to get the market demand curve (Table 18.3 and Figure 18.3). Finally, we can put it all together by combining the market supply curve and the market demand curve (Table 18.4 and Figure 18.4).

18.3 *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

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

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

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

18.4 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 18.2: 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 18.3: 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 18.4: Supply and demand together




Figure 18.2: Aggregating supply




Figure 18.3: Aggregating demand




Figure 18.4: 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.

18.5 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 17.3.) As on the demand

¹The 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 2.)

18.6 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

²One 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 17.3b. 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. Consider a world with 1,000 buyers: 500 of them have an individual demand curve of $q = 20 - 2p$, and 500 of them have an individual demand curve of $q = 10 - 5p$. There are also 500 sellers, each with an individual supply curve of $q = 6p - 10$. Determine the market demand curve and the market supply curve.
2. *Challenge* Show that it is theoretically impossible for a profit-maximizing firm to have a downward-sloping supply curve. To put this into mathematical terms, consider high and low prices (p^H and p^L , with $p^H > p^L$) and high and low quantities (q^H and q^L , with $q^H > q^L$). With a downward-sloping supply curve, q^H would be the profit-maximizing quantity at p^L and q^L would be the profit-maximizing quantity at p^H . Your job is to show that this is not possible. (*Hint*: Let $C(q)$ be the cost of producing q units of output, so that 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 .)
3. Consider a world with 300 consumers, each with demand curve $q = 25 - 2p$, and 500 suppliers, each with supply curve $q = 5 + 3p$.
 - (a) Calculate the equations for the market demand curve and the market supply curve.
 - (b) Determine the market equilibrium price and quantity and the total revenue in this market.
 - (c) Calculate the price elasticity of demand and the price elasticity of supply at the market equilibrium. Then calculate their ratio.
 - (d) Imagine that the government imposes a \$1 per-unit tax on the buyers. Write down the new market supply and demand curves, and find the new market equilibrium price and quantity. How much of the

tax burden is borne by the buyers, and how much by the sellers? Calculate the ratio of these tax burdens, and compare with the ratio of the elasticities calculated above.

- (e) Now imagine that the government instead decides to impose a \$1 per-unit tax on the sellers. How will this change things? Write down the new market supply and demand curves, find the new market equilibrium price and quantity, and compare with your answer from above (where the tax is on the buyer).
- (f) Now imagine that the government instead decides to impose a 50% sales tax on the sellers. Find the new market equilibrium price and quantity.
- (g) Finally, imagine that the government instead decides to impose a 100% sales tax on the buyers. Find the new market equilibrium price and quantity, and compare with your answer from above (where the tax is on the seller).

Chapter 19

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**.

19.1 Reinterpreting the Supply Curve

Consider, for example, a firm whose supply curve indicates that (1) at any price above \$5 per unit the firm wants to sell at least 10 units; and (2) at any price below \$5 per unit the firm wants to sell at most 9 units. We can now ask the following question: What is the firm’s *marginal cost* of producing the 10th unit? (**Marginal cost** is the cost of producing an additional unit. The marginal cost of the 10th unit is the cost differential between producing 10 units and producing only 9 units.)

The answer follows from profit maximization: the firm’s marginal cost of producing the 10th unit must be equal to \$5. Since the firm wants to sell at least 10 units at any price above \$5 per unit (say, \$5.10), the marginal cost of




Figure 19.1: Every point on an individual (or market) supply curve is also a point on an individual (or social) marginal cost curve

the 10th unit cannot be greater than \$5. (If the marginal cost of the 10th unit was, say, \$6, the firm wouldn't be maximizing profits by selling 10 units at a market price of \$5.10 per unit; it would make \$.90 more profit by selling only 9 units instead of 10.) And since the firm doesn't want to sell 10 units at any price below \$5 per unit (say, \$4.90), the marginal cost of the 10th unit 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 a market price of \$4.90 per unit; it would make \$.90 more profit 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 per unit the firm wants to supply 10 units. From the perspective of the **marginal cost curve**, it indicates that the marginal cost of the 10th unit of production is \$5. This is an example of a more general result: *every point on an individual supply curve is also a point on that individual's marginal cost curve.*¹

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

The close connection between supply curves and marginal cost curves gives us two ways to think about the graph in Figure 19.1. 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

¹This 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.




Figure 19.2: Producer surplus = total revenue – total variable cost

suppliers together) of producing that q th 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 per unit, the firms in the market are only going to get to keep \$3 per unit, so at a price of \$4 with a \$1 tax the sellers should be willing to sell exactly what they were willing to sell at \$3 without the tax. But 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!

19.2 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 19.2a), 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 per unit and the firm produces three units, its total revenue will be $\$3.50 \cdot 3 = \10.50 , and its total variable cost of production

²If 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).

will 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 6.)

19.3 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 18 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 any price below \$5 per unit she wants to buy at least 10 units; and (2) at any price above \$5 per unit she wants to buy at most 9 units. We can now ask the following question: What is this person's *marginal benefit* from obtaining the 10th unit? (**Marginal benefit** is the benefit from obtaining an additional unit. The marginal benefit of the 10th unit is the difference in benefits between having 10 units and having only 9 units.)

The answer follows from individual optimization: the individual's marginal benefit from obtaining the 10th unit must be equal to \$5. Since she wants to buy at least 10 units at any price below \$5 per unit (say, \$4.90 per unit), the marginal benefit of the 10th unit cannot be less than \$5. (If the marginal benefit of the 10th unit was, say, \$4, she wouldn't be optimizing by buying 10 units at a market price of \$4.90; she would be better off by buying only 9 units instead of 10.) And since she wants to buy at most 9 units at any price above \$5 per unit (say, \$5.10), the marginal benefit of the 10th unit 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 a market price of \$5.10; she would 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 per unit this individual wants to buy 10 units. From the perspective of the **marginal benefit curve**, it indicates that the marginal benefit of obtaining the 10th unit of production is \$5. This is an example of a more general result: *every point on an individual demand curve is also a point on that individual's marginal benefit curve.*

The same result applies to the market as a whole: every point on the **market demand curve** is also a point on the **social marginal benefit curve**. If the

³Again, the complications stem from the income effect.




Figure 19.3: Every point on an individual (or market) demand curve is also a point on an individual (or social) marginal benefit curve

market wants to buy 1 million units at any price below \$5 per unit but fewer than 1 million units at any price above \$5 per unit, the marginal benefit to society of obtaining the 1 millionth unit must be \$5.

The close connection between demand curves and marginal benefit curves gives us two ways to think about the graph in Figure 19.3. 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 q th 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 per unit, the buyers in the market are going to end up paying \$5, so at a price of \$4 with a \$1 tax the buyers should be willing to buy exactly what they were willing to buy at \$5 without the tax. But 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!

19.4 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 19.4a), the total benefit can be represented by the area under the demand curve (a.k.a., the marginal benefit




Figure 19.4: Consumer surplus = total benefit – total expenditure

curve).⁴

If the market price is \$3.50 per unit 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 19.5. 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 $CS = TB - TE$. Total variable cost is the area under the supply curve (a.k.a. the marginal cost curve), so producer surplus is $PS = TR - TVC$. Together producer surplus and consumer surplus constitute the **gains from trade**.

19.5 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

⁴Again, 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).

⁵Just 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.

⁶To 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 19.5: 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 20

Math: Deriving Supply and Demand Curves

At a basic level, individual supply and demand curves come from individual optimization: if at price p an individual or firm is willing to buy or sell q units of some good, it must be because buying or selling q units of that good is optimal for that individual or firm. But we can go into more depth here by actually deriving individual demand curves from utility maximization and individual supply curves from profit maximization. A useful first step is to examine the cost minimization problem.

20.1 Cost Minimization

For individuals, the cost minimization problem is to achieve a specified level of utility (say, $U = 10$) at least cost. For firms, the cost minimization problem is to produce a specified amount of output (say, $Y = 10$) at least cost. *These problems are identical: if you like, you can think of the individual as a firm whose “product” is utility, or of the firm as an individual whose “utility” depends on output.* We will reinforce this connection by using examples with similar notation: the individual we will consider gets utility from drinking lattes (L) and eating cake (K); the firm we will consider produces output from inputs of labor (L) and capital (K).

Utility Functions and Indifference Curves

Part I of this course introduced **utility functions** and **indifference curves**: if the individual’s utility function is $U(L, K) = L^{\frac{1}{2}} K^{\frac{1}{2}}$ (an example of a **Cobb-Douglas utility function**¹), then the indifference curve corresponding to a

¹The general form of a Cobb-Douglas utility function is $U = L^{\alpha} K^{\beta}$ where α and β are positive constants.

utility level of, say, 2 is the set of all consumption bundles that provide the individual with a utility of 2.

In our example, the indifference curve corresponding to a utility level of 2 contains the points $(L = 1, K = 4)$, $(L = 4, K = 1)$, and $(L = 2, K = 2)$. The equation for this indifference curve is $L^{\frac{1}{2}}K^{\frac{1}{2}} = 2$, which we can rewrite as $LK = 4$ or $K = 4L^{-1}$. The slope of this indifference curve, $\frac{dK}{dL} = -4L^{-2}$, measures the **marginal rate of substitution (MRS)** between lattes and cake: an individual with a utility level of 2 who currently has L lattes and K pieces of cake would be willing to trade up to $4L^{-2}$ pieces of cake in order to gain an extra latte. Such a substitution would leave the individual on the same indifference curve, and therefore with the same utility.

An important result that will be useful later is that the slope of the indifference curve (i.e., the marginal rate of substitution) can also be written as

$$\text{MRS} = -\frac{\frac{\partial U}{\partial L}}{\frac{\partial U}{\partial K}}.$$

Here the numerator is the **marginal utility of lattes** ($MU_L = \frac{\partial U}{\partial L}$), the extra utility the individual would get from an additional latte. The denominator is the **marginal utility of cake** ($MU_K = \frac{\partial U}{\partial K}$), the extra utility the individual would get from an additional piece of cake. Intuitively, the slope of the indifference curve tells us the maximum amount of cake the individual is willing to give up in order to receive one more latte. Since one more latte gives the individual MU_L extra utility, the amount of cake the individual should be willing to give up in order to get an additional latte is K such that $MU_K \cdot K = MU_L$, i.e., $K = \frac{MU_L}{MU_K}$. (For example, if the marginal utility of lattes is 3 and the marginal utility of cake is 1, the individual should be willing to give up 3 pieces of cake to get one more latte.) It follows that the slope of the indifference curve is

$$\text{MRS} = -\frac{MU_L}{MU_K} = -\frac{\frac{\partial U}{\partial L}}{\frac{\partial U}{\partial K}}.$$

Production Functions and Isoquants

Firms have structures analogous to utility functions and indifference curves; these are called **production functions** and **isoquants**. Given inputs of labor (L) and capital (K), the production function $f(L, K)$ describes the quantity of output that can be produced from these inputs. If the firm's production function is $Y = L^{\frac{1}{2}}K^{\frac{1}{2}}$ (an example of a **Cobb-Douglas production function**), then the **isoquant** corresponding to an output level of, say, 2 is the set of all input bundles that the firm can use to produce 2 units of output.

In our example, the isoquant corresponding to an output level of 2 contains the points $(L = 1, K = 4)$, $(L = 4, K = 1)$, and $(L = 2, K = 2)$. The equation

for this isoquant is $L^{\frac{1}{2}}K^{\frac{1}{2}} = 2$, which we can rewrite as $LK = 4$ or $K = 4L^{-1}$. The slope of this isoquant, $\frac{dK}{dL} = -4L^{-2}$, measures the **marginal rate of technical substitution (MRTS)** between labor and capital: a firm with an output target of 2 which currently has L units of labor and K units of capital would be willing to trade up to $4L^{-2}$ units of capital in order to gain an extra unit of labor. Such a substitution would leave the firm on the same isoquant, and therefore with the same output.

An important result that will be useful later is that the slope of the isoquant can also be written as

$$\text{MRTS} = -\frac{\frac{\partial f}{\partial L}}{\frac{\partial f}{\partial K}}.$$

Here the numerator is the **marginal product of labor** ($MP_L = \frac{\partial f}{\partial L}$), the extra output the firm would get from an additional unit of labor. The denominator is the **marginal product of capital** ($MP_K = \frac{\partial f}{\partial K}$), the extra output the firm would get from an additional unit of capital. Intuitively, the slope of the isoquant tells us the maximum amount of capital the firm is willing to give up in order to get one more unit of labor. Since one unit of labor allows the firm to produce MP_L extra units of output, the amount of capital the firm should be willing to give up in order to get an additional unit of labor is K such that $MP_K \cdot K = MP_L$, i.e., $K = \frac{MP_L}{MP_K}$. (For example, if the marginal product of labor is 3 and the marginal product of capital is 1, the firm should be willing to give up 3 units of capital to get one more unit of labor.) It follows that the slope of the isoquant is

$$\text{MRTS} = -\frac{MP_L}{MP_K} = -\frac{\frac{\partial f}{\partial L}}{\frac{\partial f}{\partial K}}.$$

The Cost Function

If lattes and cake (or labor and capital) have unit prices of p_L and p_K , respectively, then the total cost of purchasing L units of one and K units of the other is

$$C(L, K) = p_L L + p_K K.$$

The cost minimization problem for the individual is to choose L and K to minimize the cost necessary to reach a specified utility level (say, $U = 2$). The cost minimization problem for the firm is to choose L and K to minimize the cost necessary to reach a specified output level (say, $Y = 2$). Mathematically, the individual wants to choose L and K to minimize $p_L L + p_K K$ subject to the constraint $U(L, K) = 2$; the firm wants to choose L and K to minimize $p_L L + p_K K$ subject to the constraint $f(L, K) = 2$.

To solve this problem, we need to find the values of our choice variables (L and K) that minimize cost. Since two equations in two unknowns generally

yields a unique solution, our approach will be to find two relevant equations involving L and K ; solving these simultaneously will give us the answer to the cost minimization problem.

One equation involving L and K is clear from the set-up of the problem. For the individual, we must have $U(L, K) = 2$, i.e., $L^{\frac{1}{2}}K^{\frac{1}{2}} = 2$. For the firm, we must have $f(L, K) = 2$, i.e., $L^{\frac{1}{2}}K^{\frac{1}{2}} = 2$.

Our second constraint is a necessary first-order condition (NFOC) that looks like

$$\frac{\frac{\partial U}{\partial L}}{p_L} = \frac{\frac{\partial U}{\partial K}}{p_K} \quad \text{or} \quad \frac{\frac{\partial f}{\partial L}}{p_L} = \frac{\frac{\partial f}{\partial K}}{p_K}.$$

At the end of this section we will provide three explanations for this NFOC. First, however, we show how to combine the NFOC with the utility (or production) constraint to solve the cost minimization problem.

Solving the Individual's Cost Minimization Problem

Consider an individual with utility function $U = L^{\frac{1}{2}}K^{\frac{1}{2}}$. Assume that the prices of lattes and cake are $p_L = 1$ and $p_K = 2$. What is the minimum cost necessary to reach a utility level of 2?

Well, we know that the solution must satisfy the constraint $L^{\frac{1}{2}}K^{\frac{1}{2}} = 2$, i.e., $LK = 4$. Next, we consider our mysterious NFOC. The partial derivative of utility with respect to L is $\frac{\partial U}{\partial L} = \frac{1}{2}L^{-\frac{1}{2}}K^{\frac{1}{2}}$; the partial derivative of utility with respect to K is $\frac{\partial U}{\partial K} = \frac{1}{2}L^{\frac{1}{2}}K^{-\frac{1}{2}}$. Our NFOC is therefore

$$\frac{\frac{\partial U}{\partial L}}{p_L} = \frac{\frac{\partial U}{\partial K}}{p_K} \implies \frac{\frac{1}{2}L^{-\frac{1}{2}}K^{\frac{1}{2}}}{1} = \frac{\frac{1}{2}L^{\frac{1}{2}}K^{-\frac{1}{2}}}{2} \implies \frac{1}{2}L^{-\frac{1}{2}}K^{\frac{1}{2}} = \frac{1}{4}L^{\frac{1}{2}}K^{-\frac{1}{2}}.$$

Multiplying through by $4L^{\frac{1}{2}}K^{\frac{1}{2}}$ we get

$$2K = L.$$

In other words, the cost-minimizing solution is a consumption bundle with twice as many lattes as pieces of cake.

We can now combine our two equations to find the answer. We know (from the NFOC) that $2K = L$ and (from the utility constraint) that $LK = 4$. Solving simultaneously we get $(2K)K = 4 \implies 2K^2 = 4 \implies K = \sqrt{2}$. It follows from either of our two equations that the optimal choice of lattes is $L = 2\sqrt{2}$. So the cost minimizing consumption bundle that achieves a utility level of 2 is $(L, K) = (2\sqrt{2}, \sqrt{2})$, and the minimum cost necessary to reach that utility level is

$$C(L, K) = p_LL + p_KK = (1)2\sqrt{2} + (2)\sqrt{2} = 4\sqrt{2}.$$

Solving the Firm's Cost Minimization Problem

Now consider a firm with production function $Y = L^{\frac{1}{2}}K^{\frac{1}{2}}$. The prices of capital and labor are $p_L = 1$ and $p_K = 2$. What is the minimum cost necessary to produce q units of output?

Well, we know that the solution must satisfy the constraint $L^{\frac{1}{2}}K^{\frac{1}{2}} = q$, i.e., $LK = q^2$. Next, we consider our mysterious NFOC. The partial derivative of the production function with respect to L is $\frac{\partial f}{\partial L} = \frac{1}{2}L^{-\frac{1}{2}}K^{\frac{1}{2}}$; the partial derivative of the production function with respect to K is $\frac{\partial f}{\partial K} = \frac{1}{2}L^{\frac{1}{2}}K^{-\frac{1}{2}}$. Our NFOC is therefore

$$\frac{\frac{\partial f}{\partial L}}{p_L} = \frac{\frac{\partial f}{\partial K}}{p_K} \implies \frac{\frac{1}{2}L^{-\frac{1}{2}}K^{\frac{1}{2}}}{1} = \frac{\frac{1}{2}L^{\frac{1}{2}}K^{-\frac{1}{2}}}{2} \implies \frac{1}{2}L^{-\frac{1}{2}}K^{\frac{1}{2}} = \frac{1}{4}L^{\frac{1}{2}}K^{-\frac{1}{2}}.$$

Multiplying through by $4L^{\frac{1}{2}}K^{\frac{1}{2}}$ we get

$$2K = L.$$

In other words, the cost-minimizing solution is an input mix with twice as many units of labor as capital.

We can now combine our two equations to find the answer. We know (from the NFOC) that $2K = L$ and (from the utility constraint) that $LK = q^2$. Solving simultaneously we get $(2K)K = q^2 \implies 2K^2 = q^2 \implies K = \frac{q}{\sqrt{2}}$. It follows from either of our two equations that the optimal choice of labor is $L = q\sqrt{2}$. So the cost minimizing consumption bundle that achieves an output level of q is $(L, K) = \left(q\sqrt{2}, \frac{q}{\sqrt{2}}\right)$, and the minimum cost necessary to reach that output level is

$$C(L, K) = p_LL + p_KK \implies C(q) = (1)q\sqrt{2} + 2\frac{q}{\sqrt{2}} = 2q\sqrt{2}.$$

The function $C(q)$ is the firm's **cost function**: specify how much output you want the firm to produce and $C(q)$ tells you the minimum cost necessary to produce that amount of output. Note that we have transformed the cost function from one involving L and K to one involving q ; this will prove useful in deriving supply curves.

Now that we've seen how to use the mysterious NFOC, let's see why it makes sense. The remainder of this section provides three explanations—one intuitive, one graphical, and one mathematical—for our NFOC,

$$\frac{\frac{\partial U}{\partial L}}{p_L} = \frac{\frac{\partial U}{\partial K}}{p_K} \quad \text{or} \quad \frac{\frac{\partial f}{\partial L}}{p_L} = \frac{\frac{\partial f}{\partial K}}{p_K}.$$

An Intuitive Explanation for the NFOC

The first explanation is an intuitive idea called the **last dollar rule**. If our cost-minimizing individual is really minimizing costs, shifting one dollar of spending

from cake to lattes cannot increase the individual's utility level; similarly, shifting one dollar of spending from lattes to cake cannot increase the individual's utility level. The individual should therefore be indifferent between spending his "last dollar" on lattes or on cake.

To translate this into mathematics, consider shifting one dollar of spending from *cake* to *lattes*. Such a shift would allow the individual to spend one more dollar on lattes, i.e., to buy $\frac{1}{p_L}$ more lattes; this would increase his utility by $\frac{\partial U}{\partial L} \cdot \frac{1}{p_L}$. (Recall that $\frac{\partial U}{\partial L}$ is the marginal utility of lattes.) But this shift would require him to spend one less dollar on cake, i.e., to buy $\frac{1}{p_K}$ fewer pieces of cake; this would reduce his utility by $\frac{\partial U}{\partial K} \cdot \frac{1}{p_K}$. (Recall that $\frac{\partial U}{\partial K}$ is the marginal utility of cake.) Taken as a whole, this shift cannot increase the individual's utility level, so we must have

$$\frac{\partial U}{\partial L} \cdot \frac{1}{p_L} - \frac{\partial U}{\partial K} \cdot \frac{1}{p_K} \leq 0 \implies \frac{\partial U}{\partial L} \frac{1}{p_L} \leq \frac{\partial U}{\partial K} \frac{1}{p_K}.$$

Now consider shifting one dollar of spending from *lattes* to *cake*. Such a shift would allow the individual to spend one more dollar on cake, i.e., to buy $\frac{1}{p_K}$ more pieces of cake; this would increase his utility by $\frac{\partial U}{\partial K} \frac{1}{p_K}$. But this shift would require him to spend one less dollar on lattes, i.e., to buy $\frac{1}{p_L}$ fewer lattes; this would reduce his utility by $\frac{\partial U}{\partial L} \frac{1}{p_L}$. Overall, this shift cannot increase the individual's utility level, so we must have

$$\frac{\partial U}{\partial K} \cdot \frac{1}{p_K} - \frac{\partial U}{\partial L} \cdot \frac{1}{p_L} \leq 0 \implies \frac{\partial U}{\partial K} \frac{1}{p_K} \leq \frac{\partial U}{\partial L} \frac{1}{p_L}.$$

Looking at the last two equations, we see that

$$\frac{\partial U}{\partial L} \cdot \frac{1}{p_L} \leq \frac{\partial U}{\partial K} \cdot \frac{1}{p_K} \quad \text{and} \quad \frac{\partial U}{\partial K} \cdot \frac{1}{p_K} \leq \frac{\partial U}{\partial L} \cdot \frac{1}{p_L}.$$

The only way both of these equations can hold is if

$$\frac{\partial U}{\partial L} \cdot \frac{1}{p_L} = \frac{\partial U}{\partial K} \cdot \frac{1}{p_K} \quad \text{i.e.,} \quad \frac{\frac{\partial U}{\partial L}}{p_L} = \frac{\frac{\partial U}{\partial K}}{p_K}.$$

So if the individual is minimizing cost, this equation must hold.

The identical logic works for firms. If the firm is minimizing costs, shifting one dollar of spending from capital to labor cannot increase the firm's output; similarly, shifting one dollar of spending from labor to capital cannot increase the firm's output. The firm should therefore be indifferent between spending its "last dollar" on labor or on capital.

Mathematically, we end up with

$$\frac{\frac{\partial f}{\partial L}}{p_L} = \frac{\frac{\partial f}{\partial K}}{p_K}.$$

With one extra dollar, the firm could hire $\frac{1}{p_L}$ extra units of labor; the extra output the firm could produce is therefore $\frac{\partial f}{\partial L} \cdot \frac{1}{p_L}$. (Recall that $\frac{\partial f}{\partial L}$ is the marginal product of labor, i.e., the extra output the firm could produce with one extra unit of labor.) Similarly, spending an extra dollar on capital would allow the firm to hire $\frac{1}{p_K}$ extra units of capital; the extra output the firm could produce is therefore $\frac{\partial f}{\partial K} \cdot \frac{1}{p_K}$. (Recall that $\frac{\partial f}{\partial K}$ is the marginal product of capital, i.e., the extra output the firm could produce with one extra unit of capital.) If the firm is minimizing cost, it must be equating these two fractions.

A Graphical Explanation for the NFOC

The second explanation for the NFOC is graphical. Recall from Part I that an individual's **budget constraint** is the set of all consumption bundles (L, K) that an individual can purchase with a given budget. The line $p_L L + p_K K = 10$ is the budget constraint corresponding to a budget of \$10; we can rewrite this as $K = \frac{10 - p_L L}{p_K}$. The slope of the budget constraint, $\frac{dK}{dL} = -\frac{p_L}{p_K}$, measures the **marginal rate of transformation** between lattes and cake. In order to afford an extra latte, the individual needs to give up $\frac{p_L}{p_K}$ pieces of cake in order to stay within his budget. (For example, if lattes cost \$1 and cake costs \$.50 per piece, he would have to give up 2 pieces of cake to afford one extra latte.)

Firms have structures analogous to budget constraints called **isocosts**: the set of all input bundles (L, K) that the firm can purchase with a given budget. The line $p_L L + p_K K = 10$ is the isocost corresponding to a budget of \$10; we can rewrite this as $K = \frac{10 - p_L L}{p_K}$. The slope of the isocost, $\frac{dK}{dL} = -\frac{p_L}{p_K}$, measures the **marginal rate of technical transformation** between labor and capital. In order to afford an extra unit of labor, the firm needs to give up $\frac{p_L}{p_K}$ units of capital in order to stay within its budget. (For example, if labor costs \$1 per unit and capital costs \$.50 per unit, the firm would have to give up 2 units of capital to afford one extra unit of labor.)

Graphically, the cost minimization problem is for the individual to find the lowest budget constraint that intersects a specified indifference curve (or, equivalently, for the firm to find the lowest isocost that intersects a specified isoquant). We can see from Figure 20.1 that the solution occurs at a point where the budget constraint is tangent to the indifference curve (or, equivalently, where the isocost is tangent to the isoquant). At this point of tangency, the slope of the




Figure 20.1: Minimizing costs subject to a utility (or output) constraint

indifference curve must equal the slope of the budget constraint:

$$-\frac{\frac{\partial U}{\partial L}}{\frac{\partial U}{\partial K}} = -\frac{p_L}{p_K} \implies \frac{\frac{\partial U}{\partial L}}{p_L} = \frac{\frac{\partial U}{\partial K}}{p_K}.$$

Equivalently, in the case of firms we have that the slope of the isoquant must equal the slope of the isocost:

$$-\frac{\frac{\partial f}{\partial L}}{\frac{\partial f}{\partial K}} = -\frac{p_L}{p_K} \implies \frac{\frac{\partial f}{\partial L}}{p_L} = \frac{\frac{\partial f}{\partial K}}{p_K}.$$

A Mathematical Explanation for the NFOC

The third and final explanation for the NFOC comes from brute force mathematics. The individual's problem is to choose L and K to minimize costs $p_L L + p_K K$ subject to a utility constraint $U(L, K) = \bar{U}$. It turns out that we can solve this problem by writing down the **Lagrangian**

$$\mathcal{L} = p_L L + p_K K + \lambda[\bar{U} - U(L, K)].$$

(The Greek letter λ —pronounced “lambda”—is called the Lagrange multiplier; it has important economic meanings that you can learn more about in upper-level classes.) Magically, the necessary first-order conditions (NFOCs) for the individual turn out to be

$$\frac{\partial \mathcal{L}}{\partial L} = 0, \quad \frac{\partial \mathcal{L}}{\partial K} = 0, \quad \text{and} \quad \frac{\partial \mathcal{L}}{\partial \lambda} = 0,$$

i.e.,

$$p_L - \lambda \frac{\partial U}{\partial L} = 0, \quad p_K - \lambda \frac{\partial U}{\partial K} = 0, \quad \text{and} \quad \bar{U} - U(K, L) = 0.$$

Solving the first two for λ we get

$$\lambda = \frac{p_L}{\frac{\partial U}{\partial L}} \quad \text{and} \quad \lambda = \frac{p_K}{\frac{\partial U}{\partial K}}.$$

Setting these equal to each other and rearranging yields the mysterious NFOC,

$$\frac{\frac{\partial U}{\partial L}}{p_L} = \frac{\frac{\partial U}{\partial K}}{p_K}.$$

The mathematics of the firm's problem is identical: choose L and K to minimize costs $p_L L + p_K K$ subject to a production constraint $f(L, K) = \bar{Y}$. The Lagrangian is

$$\mathcal{L} = p_L L + p_K K + \lambda[\bar{Y} - f(L, K)]$$

and the necessary first-order conditions (NFOCs) are

$$\frac{\partial \mathcal{L}}{\partial L} = 0, \quad \frac{\partial \mathcal{L}}{\partial K} = 0, \quad \text{and} \quad \frac{\partial \mathcal{L}}{\partial \lambda} = 0,$$

i.e.,

$$p_L - \lambda \frac{\partial f}{\partial L} = 0, \quad p_K - \lambda \frac{\partial f}{\partial K} = 0, \quad \text{and} \quad \bar{Y} - f(L, K) = 0.$$

Solving the first two for λ we get

$$\lambda = \frac{p_L}{\frac{\partial f}{\partial L}} \quad \text{and} \quad \lambda = \frac{p_K}{\frac{\partial f}{\partial K}}.$$

Setting these equal to each other and rearranging yields the firm's NFOC,

$$\frac{\frac{\partial f}{\partial L}}{p_L} = \frac{\frac{\partial f}{\partial K}}{p_K}.$$

20.2 Supply Curves

The firm's ultimate job is not to minimize costs but to maximize profits. An examination of profit maximization allows us to derive **supply curves**, which show how a change in the price of the firm's output (p) affect the firm's choice of output (q), *holding other prices constant*. We will also be able to derive the firm's **factor demand curves**, which show how a change in the price of one of the firm's inputs (e.g., p_L , the price of labor) affects the firm's choice of how many units of that input to purchase, *holding other prices constant*.

The caveat "holding other prices constant" arises because of graphical and mental limitations. Supply and demand graphs are only two-dimensional, so we

cannot use them to examine how multiple price changes affect the firm's output or input choices; this would require three- or more-dimensional graphs. In fact, what is really going on is this: the firm's profit-maximizing level of output is a function of *all* of the input and output prices. If there are 5 inputs and 1 output, there are 6 prices to consider. Seeing how the profit-maximizing level of output changes as these prices change would require a *seven*-dimensional graph (six dimensions for the prices, and one dimension for the level of output). Unfortunately, we have a hard time visualizing seven-dimensional graphs. To solve this problem, we look at two-dimensional *cross-sections* of this seven-dimensional graph. By changing the output price while holding all of the input prices constant, we are able to see (in two dimensions) how the output price affects the firm's choice of output; in other words, we are able to see the firm's supply curve. Similarly, by changing the price of labor while holding all other prices constant, we are able to see how the price of labor affects the firm's profit-maximizing choice of labor; this gives us the firm's factor demand curve for labor.

An important implication here is that the levels at which we hold other prices constant is important: the firm's supply curve with input prices of $p_L = p_K = 2$ will be different than its supply curve with input prices of $p_L = p_K = 4$.

Choosing Output to Maximize Profits

To derive the firm's supply curve, we hold the input prices constant at specified levels. We can therefore use the firm's cost function $C(q)$ (derived previously) to express the profit maximization problem as the problem of choosing output q to maximize profits,

$$\pi(q) = pq - C(q).$$

To maximize this we take a derivative with respect to our choice variable and set it equal to zero. We get an NFOC of

$$\frac{d\pi}{dq} = 0 \implies p - C'(q) = 0 \implies p = C'(q).$$

The right hand side here is the marginal cost of producing an additional unit of output. The left hand side is the marginal benefit of producing an additional unit of output, namely the market price p . The NFOC therefore says that a profit-maximizing firm produces until the marginal cost of production is equal to the output price; this is true as long as the firm's profit-maximization problem has an interior solution. It follows that **the marginal cost curve is the supply curve** (as long as the firm's profit-maximization problem has an interior solution²).

²At low prices, the firm's optimal choice of output may be the corner solution $q = 0$. In these cases the marginal cost curve and the supply curve will not be the same.

Choosing Inputs to Maximize Profits

To get the factor demand curves, we express the firm's objective in terms of its input choices: it wants to choose inputs L and K to maximize profits

$$\pi(L, K) = pf(L, K) - p_L L - p_K K.$$

To maximize this we take partial derivatives with respect to our choice variables and set them equal to zero. For labor we get an NFOC of

$$\frac{\partial \pi}{\partial L} = 0 \implies p \frac{\partial f}{\partial L} - p_L = 0 \implies p \frac{\partial f}{\partial L} = p_L.$$

The right hand side here is the marginal cost of purchasing an additional unit of labor, namely its market price p_L . The left hand side is the marginal benefit of purchasing an additional unit of labor, namely the market price of the additional output that can be produced with that extra labor: p is the market price of output, and $\frac{\partial f}{\partial L}$ is the marginal product of labor, i.e., the extra output that can be produced with one more unit of labor. Multiplying these yields the **value of the marginal product of labor** (VMP_L). The NFOC therefore says that a profit-maximizing firm purchases labor until the value of the marginal product of labor VMP_L is equal to the cost of labor p_L . If $VMP_L > p_L$ then the firm can increase profits by purchasing additional units of labor; if $VMP_L < p_L$ then the firm can increase profits by purchasing fewer units of labor.

This NFOC turns out to yield the firm's factor demand curve for labor. Given a fixed input of capital and a fixed output price p , the NFOC tells us how much labor the firm wants to hire as a function of the price of labor p_L .

An identical analysis shows that a profit-maximizing firm must purchase capital until the value of the marginal product of capital VMP_K is equal to the cost of capital p_K :

$$\frac{\partial \pi}{\partial K} = 0 \implies p \frac{\partial f}{\partial K} - p_K = 0 \implies p \frac{\partial f}{\partial K} = p_K.$$

This NFOC gives us the firm's factor demand curve for capital.

20.3 Demand Curves

The individual's ultimate job is not to minimize costs subject to a utility constraint but to maximize utility subject to a budget constraint. An examination of utility maximization allows us to derive demand curves, which show how a change in the price of some good (e.g., p_L , the price of lattes) affects the individual's consumption choice for that good, *holding other prices constant and holding the individual's budget constant*. We can also derive the individual's **Engel curves**, which show how a change in the individual's *budget* affects the individual's consumption choices for, say, lattes, *holding all prices constant*.

As in the case of firms, the caveat about holding things constant arises because our graphs (and our minds) work best in only two dimensions. If the




Figure 20.2: Maximizing utility subject to a budget constraint of 6

individual's utility level depends on five different consumption goods and his budget, his optimal choice of lattes (L) would be a seven-dimensional graph: five dimensions for the prices of the different goods, one dimension for the budget, and one dimension for his optimal choice of L . Since we can't visualize seven dimensions, we hold five of these dimensions constant and get a two-dimensional cross-section. Holding the budget and the other four prices constant allows us to see how the price of lattes affects the individual's optimal choice of lattes; this gives us the individual's demand curve for lattes. Holding all of the prices constant allows us to see how the individual's budget affects his choice of lattes; this gives us the individual's Engel curves for lattes. As in the case of firms, the levels at which we hold various things constant is important: the individual's demand for lattes with a budget of \$20 and a price of cake of \$1 per piece will be different than his demand for lattes with a budget of \$30 and a price of cake of \$2 per piece.

Marshallian (Money Held Constant) Demand Curves

Consider the problem of choosing consumption amounts L and K to maximize utility $U(L, K)$ subject to a budget constraint $p_L L + p_K K = \bar{C}$. Graphically, this problem (with $p_L = 1$, $p_K = 2$, and $\bar{C} = 6$) and its solution are shown in Figure 20.2.

As with cost minimization, our solution method will be to find two equations involving L and K and then solve them simultaneously. The first equation comes from the statement of the problem: L and K must satisfy the budget constraint, $p_L L + p_K K = \bar{C}$.

The second equation turns out to be the same mysterious NFOC we found in the case of cost minimization:

$$\frac{\frac{\partial U}{\partial L}}{p_L} = \frac{\frac{\partial U}{\partial K}}{p_K}.$$

Intuitively, this is because the utility-maximizing individual must still follow the last dollar rule: if he is choosing optimally, he should be indifferent between spending his last dollar on lattes or on cake. Graphically, we can see from Figure 20.2 that the individual's task is to choose the highest indifference curve that lies on the relevant budget constraint; the optimum occurs at a point of tangency between the indifference curve and the budget constraint, meaning that the slope of the indifference curve must equal the slope of the budget constraint. Mathematically, we can write down the relevant Lagrangian,

$$\mathcal{L} = U(L, K) - \lambda[\bar{C} - p_L L - p_K K].$$

The NFOCs include $\frac{\partial \mathcal{L}}{\partial L} = 0$ and $\frac{\partial \mathcal{L}}{\partial K} = 0$, i.e., $\frac{\partial U}{\partial L} - \lambda p_L = 0$ and $\frac{\partial U}{\partial K} - \lambda p_K = 0$. Solving these for λ and setting them equal to each other yields the desired NFOC. Combining the budget constraint with the NFOC allow us to solve for the individual's **Marshallian demand curves**, Engel curves, &etc.

For example, consider an individual with a budget of \$6 and a utility function of $U = L^{\frac{1}{2}} K^{\frac{1}{2}}$. The price of cake is $p_K = 2$. What is this individual's Marshallian demand curve for lattes?

To solve this problem, we substitute the budget ($\bar{C} = \$6$) and the price of cake (\$2 per piece) into the budget constraint to get $p_L L + 2K = 6$. Next, we substitute the price of cake into the NFOC to get

$$\frac{\frac{\partial U}{\partial L}}{p_L} = \frac{\frac{\partial U}{\partial K}}{p_K} \implies \frac{\frac{1}{2} L^{-\frac{1}{2}} K^{\frac{1}{2}}}{p_L} = \frac{\frac{1}{2} L^{\frac{1}{2}} K^{-\frac{1}{2}}}{2} \implies \frac{1}{2p_L} L^{-\frac{1}{2}} K^{\frac{1}{2}} = \frac{1}{4} L^{\frac{1}{2}} K^{-\frac{1}{2}}.$$

Multiplying through by $4p_L L^{\frac{1}{2}} K^{\frac{1}{2}}$ we get

$$2K = p_L L.$$

We now have two equations ($p_L L + 2K = 6$ and $2K = p_L L$) in three unknowns (L , K , and p_L). We can eliminate K by using the second equation to substitute for $2K$ in the first equation: $p_L L + p_L L = 6$. This simplifies to $2p_L L = 6$ or $p_L L = 3$, which we can rewrite as

$$L = \frac{3}{p_L}.$$

This is the Marshallian demand curve for lattes when we fix the budget at 6 and the price of cake at \$2 per piece. If the price of lattes is $p_L = 1$, the optimal choice is $L = 3$; if the price of lattes is $p_L = 3$, the optimal choice is $L = 1$.




Figure 20.3: Minimizing costs subject to a utility constraint

Hicksian (Utility Held Constant) Demand Curves

For reasons to be discussed in the next section, it turns out that we should also consider the demand curves that come from the cost minimization problem discussed previously: choose consumption amounts L and K to minimize costs $C(L, K) = p_L L + p_K K$ subject to a utility constraint $U(L, K) = \bar{U}$. Graphically, this problem (with $p_L = 1$, $p_K = 2$, and $\bar{U} = 2$) and its solution are shown in Figure 20.3.

To solve this problem, we combine our utility constraint, $U(L, K) = \bar{U}$, with the usual NFOC,

$$\frac{\frac{\partial U}{\partial L}}{p_L} = \frac{\frac{\partial U}{\partial K}}{p_K}.$$

Combining the budget constraint with the NFOC allow us to solve for the individual's **Hicksian demand curves**.

For example, consider an individual with a utility constraint of $\bar{U} = 2$ and a utility function of $U = L^{\frac{1}{2}} K^{\frac{1}{2}}$. The price of cake is $p_K = 2$. What is this individual's Hicksian demand curve for lattes?

To solve this problem, we substitute the utility level ($\bar{U} = 2$) into the utility constraint to get $L^{\frac{1}{2}} K^{\frac{1}{2}} = 2$. Next, we substitute the price of cake (\$2 per piece) into the NFOC to get

$$\frac{\frac{\partial U}{\partial L}}{p_L} = \frac{\frac{\partial U}{\partial K}}{p_K} \implies \frac{\frac{1}{2} L^{-\frac{1}{2}} K^{\frac{1}{2}}}{p_L} = \frac{\frac{1}{2} L^{\frac{1}{2}} K^{-\frac{1}{2}}}{2} \implies \frac{1}{2p_L} L^{-\frac{1}{2}} K^{\frac{1}{2}} = \frac{1}{4} L^{\frac{1}{2}} K^{-\frac{1}{2}}.$$

Multiplying through by $4p_L L^{\frac{1}{2}} K^{\frac{1}{2}}$ we get

$$2K = p_L L.$$

We now have two equations ($L^{\frac{1}{2}} K^{\frac{1}{2}} = 2$ and $2K = p_L L$) in three unknowns (L , K , and p_L). We can eliminate K by using the second equation to substitute for K in the first equation:

$$L^{\frac{1}{2}} \left(\frac{p_L L}{2} \right)^{\frac{1}{2}} = 2 \implies L \left(\frac{p_L}{2} \right)^{\frac{1}{2}} = 2 \implies L\sqrt{p_L} = 2\sqrt{2}.$$

We can rewrite this as

$$L = \frac{2\sqrt{2}}{\sqrt{p_L}}$$

This is the Hicksian demand curve for lattes when we fix utility at $U = 2$ and the price of cake at \$2 per piece. If the price of lattes is $p_L = 1$, the optimal choice is $L = 2\sqrt{2}$; if the price of lattes is $p_L = 2$, the optimal choice is $L = 2$.

Marshall v. Hicks

Why do we use both Marshallian and Hicksian demand curves? Well, the advantage of the Marshallian demand curve is that it matches real life: people actually do have budget constraints, and they have to do the best they can subject to those constraints. The advantage of the Hicksian demand curve is that it makes for nice economic theory. For example, we can prove that Hicksian demand curves are downward sloping: people buy less as prices rise. And the area underneath a Hicksian demand curve represents a useful concept called **consumer surplus**. Unfortunately, Marshallian demand curves do not have nice theoretical properties; Marshallian demand curves do not have to be downward sloping, and the area underneath a Marshallian demand curve *does not* represent consumer surplus (or anything else interesting).

We are therefore in the uncomfortable situation of getting pulled in two different directions: reality pulls us toward Marshallian demand curves and theory pulls us toward Hicksian demand curves. There is, however, a silver lining in this cloud: comparing Marshallian and Hicksian demand curves provides insight into both of them.

Consider what happens if, say, the price of lattes increases. In the Marshallian approach, your purchasing power declines, and you will end up on a lower indifference curve (i.e., with a lower utility level). In the Hicksian approach, your budget adjusts upward to *compensate* you for the increased price of lattes, and you end up on the same indifference curve (i.e., with the same utility level). For this reason the Hicksian demand curve is also called the **compensated demand curve**, while the Marshallian demand curve is called the **uncompensated demand curve**.³

³Of course, the Hicksian compensation scheme can work against you, too. If the price of lattes *decreases* then the Marshallian approach puts you on a higher indifference curve, i.e., with a higher utility level. The Hicksian approach reduces your budget so that you end up on the same indifference curve, i.e., with the same utility level.

An obvious question here is: What exactly is the Hicksian (or compensated) demand curve compensating you *for*? This answer is, *for changes in the purchasing power of your budget*. The Hicksian approach increases or decreases your income so that you're always on the same indifference curve, regardless of price changes. The Hicksian demand curve therefore represents a pure **substitution effect**: when the price of lattes increases in the Hicksian world, you buy fewer lattes—not because your budget is smaller, but because lattes now look more expensive compared to cake. This change in *relative* prices induces you to buy fewer lattes and more cake, i.e., to substitute out of lattes and into cake.

The substitution effect is also present in the Marshallian approach. When the price of lattes increases, you buy fewer lattes because of the change in relative prices. But there is an additional factor: the increased price of lattes reduces your *purchasing power*. This is called an **income effect** because a loss of purchasing power is essentially a loss of income. Similarly, a decrease in the price of lattes increases your purchasing power; the income effect here is essentially a gain in income.

The Marshallian demand curve therefore incorporates both income and substitution effects; the Hicksian demand curve only incorporates the substitution effect. If you take more economics, you can prove this mathematically using the Slutsky equation, which relates these two types of demand curves.

For our purposes, it is more important to do a qualitative comparison. We can prove that Hicksian demand curves are downward sloping (i.e., that people buy fewer lattes when the price of lattes goes up) because the substitution effect always acts in the *opposite* direction of the price change: when prices go up the substitution effect induces people to buy less, and when prices go down the substitution effect induces people to buy more.

The income effect is more difficult. For **normal goods**, an increase in income leads you to buy more: your purchasing power rises, so you buy more vacations and more new computers &etc. But for **inferior goods**, an increase in your income leads you to buy *less*: your purchasing power rises, so you buy less Ramen and less used underwear &etc. In a parallel fashion, a reduction in your income leads you to buy more inferior goods: you now have less money, so you buy *more* Ramen and more used underwear &etc.

In sum, while the substitution effect always acts in the opposite direction of the price change, the income effect can go in either direction: increases or decreases in income can either increase or decrease your optimal consumption of some item, depending on whether the item is a normal good or an inferior good. In some instances (rare in practice but possible in theory), the income effect for an inferior good can overwhelm the substitution effect. Consider, for example, an individual who consumes Ramen (an inferior good) and steak (a normal good). If the price of Ramen goes up, the substitution effect will lead the individual to substitute out of Ramen and into steak, i.e., to consume less Ramen. But the increased price of Ramen will also reduce the individual's purchasing power, and this income effect will lead him to consume *more* Ramen. If the income effect is strong enough, the overall result may be that an increase in the price of Ramen leads the individual to consume *more* Ramen. In other

words, we have an *upward sloping* Marshallian demand curve! A good with an upward sloping Marshallian demand curve is called a **Giffen good**.

Problems

- C-1. Explain the importance of taking derivatives and setting them equal to zero.
- C-2. Use intuition, graphs, or math to explain the “mysterious NFOC” (a.k.a. the last dollar rule).
- C-3. Consider an individual with a utility function of $U(L, K) = L^2 K$.
- What is the equation for the indifference curve corresponding to a utility level of 9000?
 - What is the slope $\frac{dK}{dL}$ of the indifference curve, i.e., the marginal rate of substitution?
 - Explain intuitively what the marginal rate of substitution measures.
 - Assume that the prices of L and K are $p_L = 2$ and $p_K = 3$. Write down the problem for minimizing cost subject to the constraint that utility must equal 9000. Clearly specify the objective function and the choice variables.
 - Explain how to go about solving this problem.
 - Solve this problem, i.e., find the minimum cost required to reach a utility level of 9000. What are the optimal choices of L and K ?
- C-4. Consider the production function $f(L, K) = L^{\frac{1}{4}} K^{\frac{1}{2}}$.
- What is the equation for the isoquant corresponding to an output level of q ?
 - What is the slope $\frac{dK}{dL}$ of the isoquant, i.e., the marginal rate of technical substitution?
 - Explain intuitively what the marginal rate of technical substitution measures.
 - Assume that the prices of L and K are $p_L = 2$ and $p_K = 2$. Write down the problem for minimizing cost subject to the constraint that output must equal q . Clearly specify the objective function and the choice variables.
 - Explain how to go about solving this problem.
 - Solve this problem, i.e., find the minimum cost $C(q)$ required to reach an output level of q . What are the optimal choices of L and K ? (Note: these will be functions of q . You may wish to do the next problem first if you’re getting confused by all the variables.)

- (g) What is the minimum cost required to reach an output level of 10? What are the optimal choices of L and K ?
- C-5. Consider the production function $f(L, K) = L^{\frac{1}{4}}K^{\frac{1}{2}}$. Assume (as above) that the firm's input prices are $p_L = p_K = 2$; also assume that the price of the firm's output is p .
- Write down the problem of choosing output q to maximize profits. Use the cost function $C(q) \approx 3.78q^{\frac{4}{3}}$ (which you derived above) to represent costs. Clearly specify the objective function and the choice variables.
 - Explain how to go about solving this problem. Also explain why it's kosher to substitute $C(q)$ into the profit function, i.e., explain why cost-minimization is a necessary condition for profit-maximization.
 - Solve this problem to derive the firm's supply curve. Use approximations where helpful.
 - If the price of output is $p = 16$, how much will the firm produce? What will its profits be?
- C-6. The previous problems have dealt with **long run** cost curves and supply curves, meaning that the firm has complete control over all of its inputs. In the **short run**, however, the firm cannot change its capital stock—it can choose how much labor to hire, but it can't build any new factories. In this problem we will examine short run cost curves and short run supply curves.
- Assume that the firm's production function is $f(L, K) = L^{\frac{1}{4}}K^{\frac{1}{2}}$, and that capital is fixed at $K = 4$. What is the equation for the isoquant corresponding to an output level of q ?
 - Assume further that the prices of L and K are $p_L = 2$ and $p_K = 2$. Write down the problem for minimizing cost subject to the constraint that output must equal q . Clearly specify the objective function and the choice variables.
 - In a previous problem you provided an intuitive explanation for the marginal rate of technical substitution. Given that capital is fixed at $K = 4$, what is the relevance (if any) of this concept in the present problem?
 - How will the price of capital p_K affect the firm's behavior?
 - Solve this problem, i.e., find the minimum cost $C(q)$ required to reach an output level of q . What is the optimal choice of L ?
 - Write down the profit maximization problem, using the function $C(q)$ you found above. Calculate the firm's short run supply curve.
- C-7. Consider a firm with production function $f(L, K) = L^{\frac{1}{2}}K^{\frac{1}{2}}$ and input prices of $p_L = 1$ and $p_K = 2$.

- (a) Calculate the supply curve for this firm. (Note: the related cost-minimization problem was done in the text, with an answer of $C(q) = 2q\sqrt{2}$.)
- (b) How much will the firm supply at a price of $p = 2$? *Hint: Think about corner solutions!*
- (c) How much will the firm supply at a price of $p = 4$?
- (d) Show that this firm's production function exhibits **constant returns to scale**, i.e., that doubling inputs doubles output, i.e., $f(2L, 2K) = 2f(L, K)$.
- (e) Does the idea of constant returns to scale help explain the firm's behavior when $p = 4$? *Hint: Think about this in the context of the objective function.* If it does help you explain the firm's behavior, you may find value in knowing that **increasing returns to scale** occurs when doubling inputs more than doubles outputs, i.e., $f(2L, 2K) > 2f(L, K)$, and that **decreasing returns to scale** occurs when doubling inputs less than doubles outputs, i.e., $f(2L, 2K) < 2f(L, K)$. An industry with constant or increasing returns to scale can often lead to monopolization of production by a single company.

C-8. Consider an individual with a utility function of $U(L, K) = L^2K$ and a budget of M .

- (a) Write down the utility-maximization problem. Clearly specify the objective function and the choice variables. Use p_L and p_K as the prices of L and K .
- (b) Explain how to go about solving this problem if you were given values of M , p_L , and p_K .
- (c) Assume the prices of L and K are $p_L = 2$ and $p_K = 3$, and that the individual's budget is $M = 90$. Find the maximum level of utility this individual can reach with a budget of 90. What are the utility-maximizing choices of L and K ?
- (d) Calculate the individual's Marshallian demand curve for lattes⁴ if the price of K is $p_K = 3$ and the budget is $M = 90$. Use p_L for the price of lattes. What is the slope of this demand curve when $p_L = 2$?
- (e) Calculate the individual's Marshallian demand curve for cake if the price of L is $p_L = 2$ and the budget is $M = 90$. Use p_K for the price of cake. What is the slope of this demand curve when $p_K = 3$?

⁴The **Marshallian demand curve** for lattes shows the relationship between the price of lattes and the quantity of lattes demanded, holding all other prices *and the individual's budget* constant. Marshallian demand curves are also called "money-held constant" demand curves, and the repeated "m" makes for a useful mnemonic.

- (f) Calculate the individual's Engel curves for lattes and cake⁵ if the price of lattes is $p_L = 2$ and the price of cake is $p_K = 3$. Use M for the individual's budget.
- C-9. Consider the same individual as above, with utility function $U(L, K) = L^2K$. This time assume that the individual has a utility constraint of $U = 9000$.
- Write down the cost-minimization problem. Clearly specify the objective function and the choice variables. Use p_L and p_K as the prices of L and K .
 - Explain how to go about solving this problem if you were given values of p_L , and p_K .
 - Assume the prices of L and K are $p_L = 2$ and $p_K = 3$. Find the minimum budget required to reach a utility level of 9000. What are the cost-minimizing choices of L and K ?
 - Calculate the individual's Hicksian demand curve for lattes⁶ if the price of K is $p_K = 3$. Use p_L for the price of lattes. Calculate the slope of this demand curve when $p_L = 2$.
 - Calculate the individual's Hicksian demand curve for cake if the price of L is $p_L = 2$. Use p_K for the price of cake. Calculate the slope of this demand curve when $p_K = 3$.
- C-10. The previous problem (about Hicksian demand curves) asked you to minimize cost subject to a utility constraint of $U = 9000$; you should have gotten a minimum budget of $M = 90$. The problem before that (about Marshallian demand curves) asked you to maximize utility subject to a budget constraint of $M = 90$; you should have gotten a maximum utility of $U = 9000$. Explain why these two problems give symmetric answers.
- C-11. This question examines the relationship between “willingness to pay” and the area under demand curves.
- Consider (as above) an individual with utility $U(L, K) = L^2K$ who is facing prices of $p_L = 2$ and $p_K = 3$. We know from above that this individual can reach a utility level of 9000 at a minimum cost of \$90 by choosing $L = 30$, $K = 10$. Write down a *formal question* that corresponds to this informal question: what is this individual's willingness-to-pay for 10 more lattes?

⁵The **Engel curve** for lattes shows the relationship between the individual's *income* and the quantity of lattes demanded, holding all prices constant.

⁶The **Hicksian demand curve** for lattes shows the relationship between the price of lattes and the quantity of lattes demanded, holding all other prices *and the individual's utility level* constant. Hicksian demand curves are also called “utility-held constant” demand curves.

- (b) Solve this problem by answering the following questions: (1) If we give this individual 30 lattes, how much money do they have to spend on cake to reach a utility of 9000? (2) If we give this individual 40 lattes, how much money do they have to spend on cake to reach a utility of 9000? (3) Subtract. [Note: You can also do this problem by assuming that this individual *buys* the first 30 lattes at a price of \$2 each. The answers to (1) and (2) will be \$60 higher, but the answer to (3) will be the same.]
- (c) Recall from question C-9d above that the Hicksian demand curve for lattes is a rearrangement of this equation: $L^3 = 54000p_L^{-1}$. Instead of rearranging this to get the demand curve $L = (54000p_L^{-1})^{\frac{1}{3}} = (30)2^{\frac{1}{3}}p_L^{-\frac{1}{3}}$, rearrange it to get an **inverse Hicksian demand curve**, i.e., an expression for p_L as a function of L
- (d) Draw a graph of this inverse demand curve. Shade in the area between $L = 30$ and $L = 40$.
- (e) Calculate the area under the inverse Hicksian demand curve between $L = 30$ and $L = 40$. Compare with your willingness-to-pay calculation above.
- (f) Can you think of a *formal question* that is answered by the area under a *Marshallian* demand curve?

Chapter 21

Transition: Welfare Economics

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

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

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

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

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

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

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

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

As long as markets are complete and competitive, then, the First Welfare Theorem says that competition will result in a Pareto efficient outcome. Recall from Chapter 7, 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 7 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.

21.1 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
MONEY IN			(billions)
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
MONEY IN			(millions)
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 (millions)
MONEY IN			
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.

Appendix B

Math: Monopoly and Oligopoly

B.1 Monopoly: The Vanilla Version

The vanilla version of monopoly features a monopolist who must charge everyone the same price p . In this case the monopolist looks at her demand curve (say, $q = 20 - p$) and chooses p and q to maximize profits, which are revenues pq minus costs (say, $C(q) = 4q$). Of course, the monopolist cannot choose just any pair (p, q) or else she would choose $p = q = \infty$; her choice is *constrained* by consumer demand, i.e., by the demand curve $q = 20 - p$. In other words, the monopolist's true problem is to choose p and q to maximize

$$\pi = pq - C(q)$$

subject to the constraint

$$q = D(p).$$

In our example, the monopolist's problem is to choose p and q to maximize $\pi = pq - 4q$ subject to $q = 20 - p$.

There are a number of ways to solve this problem (and we will return to constrained maximization later in the quarter), but the easiest way for now is to solve the constraint for the **inverse demand curve** $p(q) = D^{-1}(q)$ (in our example, $p(q) = 20 - q$) and then eliminate the constraint by substituting this value of p into the profit function. In other words, we reduce the monopolist's problem to that of choosing q to maximize

$$\pi = p(q) \cdot q - C(q).$$

In our example, the monopolist's problem is to choose q to maximize $\pi = (20 - q) \cdot q - 4q$.

Now we are on familiar ground: we take a derivative with respect to q and set it equal to zero to find candidate maxima. The derivative in our example is

$$\frac{d\pi}{dq} = 0 \implies (20 - q) + (-1)(q) - 4 = 0 \implies 2q = 16 \implies q = 8.$$

So in our example, the profit-maximizing choice is $q = 8$, which implies $p = 20 - q = 12$ and $\pi = pq - C(q) = (12)(8) - 4(8) = 64$.

In general,

$$\frac{d\pi}{dq} = 0 \implies p(q) + p'(q) \cdot q - C'(q) = 0. \quad (\text{B.1})$$

There are two interpretations of this. The standard interpretation is to rewrite this as

$$p(q) + p'(q) \cdot q = C'(q).$$

The left hand side here is $\frac{d(pq)}{dq}$, i.e., “marginal revenue”: producing and selling an extra unit brings you a gain of $p(q)$, but in order to sell that extra unit you need to lower the price on *all* units by $p'(q)$, so there is a cost of $p'(q)$ per unit, or $p'(q) \cdot q$ in total. So your marginal revenue from producing another unit is $p(q) + p'(q) \cdot q$. The right hand side of this equation is marginal cost, so this equation tells the monopolist to produce until marginal revenue (the extra revenue you get from producing an additional unit) equals marginal cost (the extra cost of producing an additional unit.) This is a necessary first-order condition for profit-maximization. The various curves are shown in Figure B.1 on page 227.

A second interpretation (which is totally non-standard) is to rewrite Equation B.1 as

$$p(q) = C'(q) - p'(q) \cdot q.$$

The story behind this interpretation is that the monopolist has already sold a bunch of units at a high price, but has sold them with guaranteed refunds: if she lowers her price, she guarantees to make up the difference to her previous customers. So: the left hand side here is the gain she can get $p(q)$ she can get from selling an additional unit. The right hand side here is her cost: the cost of producing that extra unit ($C'(q)$), plus the cost of refunding money to her previous customers; the price difference is $p'(q)$ and she has q previous customers, so her refund cost is $-p'(q) \cdot q$. (Note that $p'(q)$ is negative, so $-p'(q) \cdot q$ is positive.) The various curves here are shown in Figure B.2 on page 227.

B.2 Monopoly: The Chocolate-with-Sprinkles Version

As we can see from above, the difficulty facing the monopolist is that she has to choose between selling only a few units for a high price or lots of units for a




Figure B.1: The standard approach: curves representing demand, marginal revenue, and marginal cost




Figure B.2: A non-standard approach: curves representing demand, marginal social cost, and marginal monopoly cost

low price. Wouldn't it be great if she could sell to some people for a high price and to other people for a lower price?

Well, yes, it would be for her (the monopolist). This is the case of the **price discriminating monopoly**: a monopoly that can charge different customers different prices. The various types of price discrimination are discussed in the text (on p. 76), and can be read about in more detail in books like *Information Rules* by Carl Shapiro and Hal Varian. Also discussed in the text are limitations on price discrimination such as resale: if customers can resell products (e.g., on eBay), monopolists will have a hard time selling the same product for different prices.

The mathematics for modeling price discrimination are (slightly) beyond the level of this class, but you can study it in an advanced class on game theory (or in this class for an extra credit project). These models generally feature different types of customers—e.g., business customers who are willing to pay a lot of money for a high-speed printer and family customers who are willing to pay less money for a low-speed printer. The trick is for the monopolist to design a mechanism that uses these differences to get as much money as possible from her customers. For example, IBM approached its printer problem by producing a high-speed printer that it sold for a high price and a low-speed printer that it sold for a low price; a computer magazine analysis revealed that the two printers were identical in all respects except one: the low-speed, slow-price printer had a chip in it that made it stop for a few seconds at the end of each line! (The story and others like it can be found in Shapiro and Varian's book.)

The monopolist's situation fits under the rubric of **principal-agent problems**. These feature a principal (e.g., the monopolist, or an employer) and one or more agents (e.g., customers, or employees). Both the principal and the agent(s) have objective functions, but in general these objective functions are different, e.g., the employer wants the employee to work hard but the employee wants to slack off as much as possible without getting fired. The trick is for the principal to design an incentive scheme that will induce the desired response from the agent(s). Such schemes are constrained by the desires of the agent(s): if the employer asks for too much hard work, the employee will quit and find another job; if the monopolist charges too much for its high-speed printer, the business customers will buy the low-speed printer instead, or won't buy any printer at all. These problems are (IMHO) lots of fun, and have really neat implications that are sometimes visible in the real world.

B.3 Duopoly

Returning to a world without price discrimination, we can now consider what happens in a market with two producers (i.e., a **duopoly**). We will study three different scenarios: collusion (in which the firms work together), Cournot competition (in which the firms compete on the basis of quantity), and Stackleberg competition (in which one firm produces and then the other firm responds.)

The situation we will consider is the following: Coke and Pepsi are engaged

in monopolistic competition in the diet soft drink industry. Both firms have costs $C(q) = 5q$, where q is measured in millions of gallons of syrup. Coke can set any price p_C for each gallon of its syrup, and Pepsi can similarly choose p_P , but both are constrained by their customers' demand curves. Since the two products are substitutes, their demand curves are interrelated, as indicated by the inverse demand curves

$$\begin{aligned} p_C &= 20 - 2q_C - q_P \\ p_P &= 20 - 2q_P - q_C. \end{aligned}$$

These inverse demand curves indicate that the maximum price that Coke can charge is strongly influenced by the amount of Diet Coke it produces, and less strongly influenced by the amount of Diet Pepsi that Pepsi produces. The same is true for Pepsi: the more Diet Pepsi they produce, the less they can charge; and the more Diet Coke Coke produces, the less Pepsi can charge.

Each company wants to maximize profits subject to the constraints of their respective demand curves. As in the monopoly situation above, we can substitute for the per-gallon prices p_C and p_P from the inverse demand curves to get the following objective functions and choice variables:

Coke wants to choose q_C to maximize $\pi_C = p_C(q_C, q_P) \cdot q_C - C(q_C)$.

Pepsi wants to choose q_P to maximize $\pi_P = p_P(q_C, q_P) \cdot q_P - C(q_P)$.

These objective functions simplify as

$$\begin{aligned} \pi_C &= (20 - 2q_C - q_P) \cdot q_C - 5q_C = (15 - q_P)q_C - 2q_C^2 \\ \pi_P &= (20 - 2q_P - q_C) \cdot q_P - 5q_P = (15 - q_C)q_P - 2q_P^2. \end{aligned}$$

Collusion

In the collusive situation the companies work together to maximize their *joint* profits. In other words, we imagine a situation in which a single firm (say, Microsoft) buys both Coke *and* Pepsi, somehow miraculously avoiding anti-trust scrutiny. In this case the owner has two choice variables (q_C and q_P) and chooses them to maximize joint profits,

$$\pi_C + \pi_P = (15 - q_P)q_C - 2q_C^2 + (15 - q_C)q_P - 2q_P^2.$$

Our necessary first-order condition (NFOC) for an interior maximum is that the partial derivatives with respect to the choice variables must all equal zero, so we must have

$$\frac{\partial(\pi_C + \pi_P)}{\partial q_C} = 0 \implies 15 - q_P - 4q_C - q_P = 0 \implies 4q_C = 15 - 2q_P$$

and

$$\frac{\partial(\pi_C + \pi_P)}{\partial q_P} = 0 \implies -q_C + 15 - q_C - 4q_P = 0 \implies 4q_P = 15 - 2q_C.$$

We now have a system of two simultaneous equations ($4q_C = 15 - 2q_P$ and $4q_P = 15 - 2q_C$) in two unknown (q_C and q_P). We can solve these, e.g., by solving the first equation for q_C and plugging into the second:

$$4q_P = 15 - 2q_C \implies 4q_P = 15 - 2\left(\frac{15 - 2q_P}{4}\right) \implies 8q_P = 30 - 15 + 2q_P.$$

This simplifies to $q_P = 2.5$, and we can plug this back in to either of our NFOCs to get $q_C = 2.5$.

So the output choices that maximize joint profits are $q_C = q_P = 2.5$, and we can plug these numbers into the joint profit function to get

$$\pi_C + \pi_P = (15 - q_P)q_C - 2q_C^2 + (15 - q_C)q_P - 2q_P^2 = 37.5.$$

So if the two firms collude then they will each produce 2.5 million gallons of syrup and will have a combined profit of \$37.5 million.

We can also use the inverse demand curves to find the prices they will charge, e.g., $p_C = 20 - 2q_C - q_P = 12.5$ and $p_P = 20 - 2q_P - q_C = 12.5$; so each firm will charge \$12.50 for each gallon of syrup. To check our answers, we can recalculate profits (hopefully yielding 37.5) using

$$\pi_C + \pi_P = p_C \cdot q_C - 2q_C^2 + p_P \cdot q_P - 2q_P^2.$$

Cournot (Quantity) Competition

Our next scenario features what is called **Cournot competition**: the two firms simultaneously choose quantities q_C and q_P , the prices follow from the inverse demand curves, and profits are determined accordingly.

In this case there is *no* joint profit maximization. Instead, Coke chooses its quantity q_C to maximize its profits,

$$\pi_C = (20 - 2q_C - q_P) \cdot q_C - 5q_C = (15 - q_P)q_C - 2q_C^2$$

and Pepsi chooses its quantity q_P to maximize its profits,

$$\pi_P = (20 - 2q_P - q_C) \cdot q_P - 5q_P = (15 - q_C)q_P - 2q_P^2.$$

The necessary first-order condition (NFOC) for an interior maximum to Coke's problem is that the derivative with respect to its choice variable must be zero, i.e.,

$$\frac{d\pi_C}{dq_C} = 0 \implies 15 - q_P - 4q_C = 0 \implies q_C = 3.75 - .25q_P.$$

This is the **best response function** for Coke. *If* Pepsi produces q_P , *then* Coke's profit-maximizing choice of q_C is $q_C = 3.75 - .25q_P$: if Pepsi produces $q_P = 2$ then Coke's best response is $q_C = 3.25$; if Pepsi produces $q_P = 4$ then Coke's best response is $q_C = 2.75$; if Pepsi produces $q_P = 8$ then Coke's best response is $q_C = 1.75$; and if Pepsi produces $q_P = 16$ then Coke's best response

is $q_C = -.25$, suggesting that we have a corner solution of $q_C = 0$. (If Pepsi produces $q_P = 16$ then Coke's inverse demand curve shows that Coke cannot charge a price higher than 4, in which case it cannot cover its production costs and should therefore produce 0.)

Similarly, the necessary first-order condition (NFOC) for an interior maximum to Pepsi's problem is that the derivative with respect to its choice variable must be zero, i.e.,

$$\frac{d\pi_P}{dq_P} = 0 \implies 15 - q_C - 4q_P = 0 \implies q_P = 3.75 - .25q_C.$$

This is the **best response function** for Pepsi. *If* Coke produces q_C , *then* Pepsi's profit-maximizing choice of q_P is $q_P = 3.75 - .25q_C$.

The Cournot solution to this game is to find mutual best responses, i.e., a Nash equilibrium solution. (Since Cournot solved this problem before Nash was born, it doesn't seem fair to call it the Nash equilibrium solution; it is, however, sometimes called the Nash-Cournot solution.)

To find the Cournot solution we simultaneously solve our two best response functions, $q_C = 3.75 - .25q_P$ and $q_P = 3.75 - .25q_C$. We can do this by plugging q_C into the second function and solving:

$$q_P = 3.75 - .25(3.75 - .25q_P) \implies 4q_P = 15 - (3.75 - .25q_P) \implies 3.75q_P = 11.25$$

This simplifies to $q_P = 3$; plugging into Coke's best response function we get $q_C = 3$.

So the outputs $q_C = 3$ and $q_P = 3$ are mutual best responses: if Coke produces 3 million gallons of Diet Coke, Pepsi's best response (i.e., its profit-maximizing choice) is to produce 3 million gallons of Diet Pepsi. And if Pepsi produces 3 million gallons of Diet Pepsi, Coke's best response is to produce 3 million gallons of Diet Coke.

We have therefore found the Cournot solution to this game: $q_C = q_P = 3$. We can now plug this answer into the various profit functions to get

$$\pi_C = (15 - q_P)q_C - 2q_C^2 = 18$$

and

$$\pi_P = (15 - q_C)q_P - 2q_P^2 = 18.$$

We can also solve for their respective prices from the inverse demand curves to get $p_C = p_P = 11$.

How does this compare with the cooperative outcome? Well, cooperating yields joint profits of \$37.5 million; under Cournot competition, the joint profits are only \$18 + \$18 = \$36 million. So the firms could do better by cooperating, meaning that firms would be likely to collude (e.g., by fixing prices) if it weren't illegal and if it were easy for them to communicate and write enforceable contracts.

Stackleberg (Leader-Follower) Competition

Our final scenario features what is called **Stackleberg leader-follower competition**: one firm (say, Coke) gets to play first, by choosing q_C ; the second firm (Pepsi) then observes q_C and subsequently chooses q_P . The prices follow from the inverse demand curves, and profits are determined accordingly.

This game is a lot like the Cournot game in that the choice variables are the same and the objective functions look the same. The key difference is that the two firms are choosing sequentially rather than simultaneously. We must therefore apply the tool of backward induction, starting with the second stage of the game (Pepsi's choice).

Pepsi chooses second, so when it chooses q_P it *knows* what Coke has decided to produce. It therefore takes q_C as given and chooses its quantity q_P to maximize its profits,

$$\pi_P = (20 - 2q_P - q_C) \cdot q_P - 5q_P = (15 - q_C)q_P - 2q_P^2.$$

The necessary first-order condition (NFOC) for an interior maximum to Pepsi's problem is that the derivative with respect to its choice variable must be zero, i.e.,

$$\frac{d\pi_P}{dq_P} = 0 \implies 15 - q_C - 4q_P = 0 \implies q_P = 3.75 - .25q_C.$$

This is the **best response function** for Pepsi. *If* Coke produces q_C , *then* Pepsi's profit-maximizing choice of q_P is $q_P = 3.75 - .25q_C$.

So far this is exactly the same as the Cournot solution. The difference is that Coke should *anticipate* that Pepsi will produce $q_P = 3.75 - .25q_C$, so Coke needs to take this into account when it determines q_C .

Coke's problem, then, is to choose q_C to maximize its profits,

$$\pi_C = (20 - 2q_C - q_P) \cdot q_C - 5q_C = (15 - q_P)q_C - 2q_C^2$$

subject to the constraint that Pepsi will produce $q_P = 3.75 - .25q_C$.

To solve this problem, we can substitute the constraint into the objective function to get

$$\begin{aligned} \pi_C &= (15 - q_P)q_C - 2q_C^2 \\ &= [15 - (3.75 - .25q_C)]q_C - 2q_C^2 \\ &= (11.25 + .25q_C)q_C - 2q_C^2 \\ &= 11.25q_C - 1.75q_C^2. \end{aligned}$$

The necessary first-order condition (NFOC) for an interior maximum to Coke's problem is that the derivative with respect to its choice variable must be zero, i.e.,

$$\frac{d\pi_C}{dq_C} = 0 \implies 11.25 - 3.5q_C = 0 \implies q_C = \frac{11.25}{3.5} = \frac{45}{14} \approx 3.214.$$

So Coke's profit-maximizing strategy is to choose $q_C = 3.214$. This choice of q_C will induce Pepsi to produce $q_P = 3.75 - .25q_C \approx 2.946$ in accordance with its best-response functions. The resulting profit levels are then given by

$$\pi_C = 11.25q_C - 1.75q_C^2 \approx 18.08$$

and

$$\pi_P = (15 - q_C)q_P - 2q_P^2 \approx 17.36.$$

Two things are worth noting here. First, joint profits are about \$35.4 million, even lower than under Cournot competition. Second, there is a **first mover advantage**! first mover advantage duopoly! first mover advantage: Coke's profits are higher than Pepsi's. This advantage comes from Coke's ability to commit. **Commitment** allows Coke to muscle its way into a greater market share and higher profits.

B.4 The Transition to Perfect Competition

What happens if we take a monopoly model and add more firms? Well, if we add one firm we get duopoly. If we add additional firms we get **oligopoly**. The interesting thing about oligopoly is that it bridges the gap between monopoly (a market with just one firm) and competitive markets (a market with a large number of small firms).

To see how this works, consider a simple monopoly model: only one firm produces soda, the demand curve for soda is $q = 10 - p$, and the cost of producing soda is $C(q) = 2q$. The monopolist chooses p and q to maximize $\pi = pq - C(q) = pq - 2q$ subject to the constraint $q = 10 - p$. Equivalently, the monopolist chooses q to maximize $\pi = (10 - q)q - 2q = 8q - q^2$. Setting a derivative equal to zero gives us $8 - 2q = 0$, meaning that the monopolist will choose $q = 4$, $p = 10 - q = 6$, and get profits of $\pi = pq - C(q) = 16$.

Next: imagine that a second firm, also with costs $C(q) = 2q$, enters the market and engages in Cournot competition with the first firm. The demand curve is still $q = 10 - p$, but now $q = q_1 + q_2$, i.e., the market output is the sum of each firm's output. We can transform the demand curve $q_1 + q_2 = 10 - p$ into the inverse demand curve, $p = 10 - q_1 - q_2$.

Now, firm 1 chooses q_1 to maximize $\pi_1 = pq_1 - C(q_1) = (10 - q_1 - q_2)q_1 - 2q_1 = (8 - q_2)q_1 - q_1^2$. Setting a derivative equal to zero gives us $8 - q_2 - 2q_1 = 0$, which rearranges as $q_1 = 4 - .5q_2$. This is the best response function for firm 1: given q_2 , it specifies the choice of q_1 that maximizes firm 1's profits.

Since the problem is symmetric, firm 2's best response function is $q_2 = 4 - .5q_1$. Solving these simultaneous to find the Cournot solution yields

$$q_1 = 4 - .5(4 - .5q_1) \implies .75q_1 = 2 \implies q_1 = \frac{8}{3} \approx 2.67.$$

We get the same result for q_2 , so the market price will be $10 - q_1 - q_2 = \frac{14}{3} \approx 4.67$.

Each firm will earn profits of

$$pq_i - 2q_i = \frac{14}{3} \cdot \frac{8}{3} - 2 \cdot \frac{8}{3} = \frac{64}{9} \approx 7.11,$$

so industry profits will be about $2(7.11) = 14.22$.

Next: imagine that a third firm, also with costs $C(q) = 2q$, enters the market and engages in Cournot competition with the first two firms. The demand curve is still $q = 10 - p$, but now $q = q_1 + q_2 + q_3$. We can transform the demand curve $q_1 + q_2 + q_3 = 10 - p$ into the inverse demand curve, $p = 10 - q_1 - q_2 - q_3$.

Now, firm 1 chooses q_1 to maximize $\pi_1 = pq_1 - C(q_1) = (10 - q_1 - q_2 - q_3)q_1 - 2q_1 = (8 - q_2 - q_3)q_1 - q_1^2$. Setting a derivative equal to zero gives us $8 - q_2 - q_3 - 2q_1 = 0$, which rearranges as $q_1 = 4 - .5(q_2 + q_3)$. This is the best response function for firm 1: given q_2 and q_3 , it specifies the choice of q_1 that maximizes firm 1's profits.

Since the problem is symmetric, firms 2 and 3 have similar best response functions, $q_2 = 4 - .5(q_1 + q_3)$ and $q_3 = 4 - .5(q_1 + q_2)$, respectively. Solving these three best response functions simultaneous yields $q_1 = q_2 = q_3 = 2$. (Note: the brute force method of solving these equations simultaneous gets a bit messy. An easier way is to notice that the solution must be symmetric, i.e., with $q_1 = q_2 = q_3$; substituting for q_2 and q_3 in firm 1's best response function yields $q_1 = 4 - .5(2q_1)$, which quickly leads to the solution.)

So with three firms we have each firm producing $q_i = 2$; output price will therefore be $p = 10 - q_1 - q_2 - q_3 = 4$, and each firm will have profits of $\pi_i = pq_i - C(q_i) = 4(2) - 2(2) = 4$. Industry profits will therefore be $3(4) = 12$.

If we step back, we can see a general trend: the total amount produced by the industry is increasing (from 4 to $2(2.67) = 5.34$ to $3(2) = 6$), the price is dropping (from 6 to 4.67 to 4), and total industry profits are decreasing (from 16 to 14.22 to 12).

We can generalize this by considering a market with n identical firms, each with costs $C(q) = 2q$. Facing a demand curve of $q_1 + q_2 + \dots + q_n = 10 - p$, we can get an inverse demand curve of $p = 10 - q_1 - q_2 - \dots - q_n$. Firm 1, for example, will then choose q_1 to maximize

$$\pi_1 = pq_1 - C(q_1) = (10 - q_1 - q_2 - \dots - q_n)q_1 - 2q_1 = (8 - q_2 - q_3 - \dots - q_n)q_1 - q_1^2.$$

Setting a derivative equal to zero gives us a necessary first order condition (NFOC) of $8 - q_2 - q_3 - \dots - q_n - 2q_1 = 0$. The other firms will have symmetric solutions, so in the end we must have $q_1 = q_2 = \dots = q_n$. Substituting these into the NFOC gives us

$$8 - (n - 1)q_1 - 2q_1 = 0 \implies 8 - (n + 1)q_1 = 0 \implies q_1 = \frac{8}{n + 1}.$$

Each firm produces the same amount of output, so total industry output will be $\frac{8n}{n + 1}$. As n gets larger and larger, approaching infinity, total industry output therefore approaches 8. The market price for any given n is $p = 10 - \frac{8n}{n + 1}$, so as

n approaches infinity the market price approaches $10 - 8 = 2$. Since production costs are $C(q) = 2q$, this means that firm profits approach zero as the number of firms gets larger and larger. We will see in the next part of the course that the industry is converging to the competitive market outcome! So competitive markets can be thought of as a limiting case of oligopoly as the number of firms gets very large. . . .

Problems

1. Consider two firms, each with costs $C(q) = 3q^2$. They produce identical products, and the market demand curve is given by $q_1 + q_2 = 10 - p$. Find each firm's output and the market price under (1) collusion, (2) Cournot competition, and (3) Stackleberg leader-follower competition. Also find total industry profits under collusion.
2. The same problem, only now Firm 1 has costs of $C(q_1) = 4q_1^2$. (Firm 2's costs remain at $C(q_2) = 3q_2^2$.)
3. The same problem, only now Firm 1 has costs of $C(q_1) = 3q_1^2$ and Firm 2 has costs of $C(q_2) = 4q_2^2$.
4. The same problem, with each firm having costs of $C(q) = 3q^2$, only now there is monopolistic competition, so that the inverse demand curve for Firm 1's output is $p_1 = 10 - 2q_1 - q_2$ and the inverse demand curve for Firm 2's output is $p_2 = 10 - q_1 - 2q_2$.

Appendix C

Answer Key for Part I: The Optimizing Individual

Chapter 1: Decision Theory

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

We can also see this mathematically: the amount spent on the rights to the Games is a constant term in the profit function, e.g., $\pi = f(h) - c$ where $f(h)$ are advertising revenues from showing the games for h hours and c is the amount spent on the rights to the games. To maximize this, we take a derivative with respect to h and set it equal to zero. Note that the constant term c drops out when we take a derivative, and therefore has no impact on the optimal choice of h .

2. A billionaire would probably take this bet because this is a terrific (although risky) investment. A penniless person would probably not take the bet because the risk of ending up with zero is too great. So sunk costs cannot be entirely ignored in decision-making; rather, the point is that it is not possible to base decisions *solely* on sunk costs.
3. The law does not deter Alice from committing additional crimes because she's already facing the death penalty if she's caught. The law does deter Betty, because she hasn't killed anybody yet.
4. You should not believe the spokesman's explanation because the R&D expenditure is a sunk cost. If it spent twice as much or half as much to discover the drug, it should still charge the same price, because that's the price that maximizes profit. The only time that R&D costs affect the company's behavior is *before* they're sunk: when the company is thinking

about spending money on R&D, it has to determine whether or not it's going to be profitable to make that investment given their estimate of how much they'll be able to charge for the pill. Once they do the R&D, however, it's a sunk cost and will no longer influence their profit-maximizing decisions.

Calculus Problems

C-1. Microeconomics is about the actions and interactions of optimizing agents (e.g., profit-maximizing firms, utility-maximizing consumers). For differentiable functions with interior maxima or minima, the way to find those interior maxima or minima is to take a derivative and set it equal to zero. This gives you *candidate values* for maxima or minima; the reason is that slopes (i.e., derivatives) are equal to zero at the top of a hill (a maximum) or at the bottom of a valley (a minimum).

C-2.

$$\begin{aligned}\frac{d}{dx}[(c \cdot f(x))] &= \lim_{h \rightarrow 0} \frac{c \cdot f(x+h) - c \cdot f(x)}{h} \\ &= c \cdot \lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h} = c \cdot \frac{d}{dx}[f(x)].\end{aligned}$$

C-3. $\frac{d}{dx}(x^2) = \frac{d}{dx}(x \cdot x) = x \cdot \frac{d}{dx}(x) + x \cdot \frac{d}{dx}(x) = 2x.$

For higher powers, use induction: We've just proved that our rule ($\frac{d}{dx}(x^n) = nx^{n-1}$) is true for $n = 2$. So now we assume that it's true for n ($\frac{d}{dx}(x^n) = nx^{n-1}$) and need to show that it's true for x^{n+1} . But we can just use the same trick again:

$$\frac{d}{dx}(x^{n+1}) = \frac{d}{dx}(x \cdot x^n) = x \cdot \frac{d}{dx}(x^n) + x^n \cdot \frac{d}{dx}(x) = xnx^{n-1} + x^n = (n+1)x^n.$$

C-4.

$$\begin{aligned}\frac{d}{dx} \left[\frac{f(x)}{g(x)} \right] &= \frac{d}{dx} [f(x) \cdot (g(x))^{-1}] \\ &= f(x) \cdot \frac{d}{dx} [(g(x))^{-1}] + (g(x))^{-1} \cdot \frac{d}{dx} [f(x)] \\ &= f(x) \cdot (-1)[g(x)]^{-2} \frac{d}{dx} [g(x)] + (g(x))^{-1} \cdot \frac{d}{dx} [f(x)] \\ &= \frac{-f(x) \cdot g'(x)}{[g(x)]^2} + \frac{f'(x)}{g(x)} \cdot \frac{g(x)}{g(x)} \\ &= \frac{g(x) \cdot f'(x) - f(x) \cdot g'(x)}{[g(x)]^2}.\end{aligned}$$

- C-5. (a) We have $f'(x) = 2x$ and $f''(x) = 2$. Candidate solutions for interior maxima or minima are where $f'(x) = 0$. The only candidate is $x = 0$, which turns out to be a minimum value of the function $f(x)$. Note that the sign of the second derivative, $f''(0) = 2 > 0$, identifies this as a minimum; this is also clear from a graph of $f(x)$.
- (b) We have $f'(x) = 2(x^2 + 2) \cdot \frac{d}{dx}(x^2 + 2) = 4x(x^2 + 2) = 4x^3 + 8x$ and $f''(x) = 12x^2 + 8$. Candidate solutions for interior maxima or minima are where $f'(x) = 0$. The only candidate is $x = 0$, which turns out to be a minimum value of the function $f(x)$. (We are not interested in imaginary roots such as $i\sqrt{2}$.) Note that the sign of the second derivative, $f''(0) = 8 > 0$, identifies this as a minimum.
- (c) We have $f'(x) = \frac{1}{2} \cdot (x^2 + 2)^{-\frac{1}{2}} \cdot \frac{d}{dx}(x^2 + 2) = x \cdot (x^2 + 2)^{-\frac{1}{2}}$ and

$$\begin{aligned} f''(x) &= x \cdot \frac{d}{dx} \left[(x^2 + 2)^{-\frac{1}{2}} \right] + \left[(x^2 + 2)^{-\frac{1}{2}} \right] \cdot \frac{d}{dx}(x) \\ &= x \cdot \left(-\frac{1}{2} \right) \cdot (x^2 + 2)^{-\frac{3}{2}} \frac{d}{dx}(x^2 + 2) + \left[(x^2 + 2)^{-\frac{1}{2}} \right] \cdot 1 \\ &= -x^2 \cdot (x^2 + 2)^{-\frac{3}{2}} + (x^2 + 2)^{-\frac{1}{2}} \end{aligned}$$

Candidate solutions for interior maxima or minima are where $f'(x) = 0$. The only candidate is $x = 0$, which turns out to be a minimum value of the function $f(x)$. Note that the sign of the second derivative, $f''(0) = \frac{1}{\sqrt{2}} > 0$, identifies this as a minimum.

- (d) We have

$$\begin{aligned} f'(x) &= -x \cdot \frac{d}{dx} \left[(x^2 + 2)^{\frac{1}{2}} \right] + \left[(x^2 + 2)^{\frac{1}{2}} \right] \cdot \frac{d}{dx}(-x) \\ &= -x \cdot \left(\frac{1}{2} \right) \cdot (x^2 + 2)^{-\frac{1}{2}} \frac{d}{dx}(x^2 + 2) + \left[(x^2 + 2)^{\frac{1}{2}} \right] \cdot (-1) \\ &= -x^2 \cdot (x^2 + 2)^{-\frac{1}{2}} - (x^2 + 2)^{\frac{1}{2}} \\ \text{and } f''(x) &= \frac{d}{dx} \left[-x^2 \cdot (x^2 + 2)^{-\frac{1}{2}} \right] - \frac{d}{dx} \left[(x^2 + 2)^{\frac{1}{2}} \right] \\ &= -x^2 \cdot \frac{d}{dx} \left[(x^2 + 2)^{-\frac{1}{2}} \right] + (x^2 + 2)^{-\frac{1}{2}} \cdot \frac{d}{dx}(-x^2) \\ &\quad - \frac{1}{2} \left[(x^2 + 2)^{-\frac{1}{2}} \right] \frac{d}{dx}(x^2 + 2) \\ &= -x^2 \cdot \left(-\frac{1}{2} \right) \left[(x^2 + 2)^{-\frac{3}{2}} \right] \frac{d}{dx}(x^2 + 2) - 2x(x^2 + 2)^{-\frac{1}{2}} \\ &\quad - x(x^2 + 2)^{-\frac{1}{2}} \\ &= x^3(x^2 + 2)^{-\frac{3}{2}} - 3x(x^2 + 2)^{-\frac{1}{2}} \end{aligned}$$

Candidate solutions for interior maxima or minima are where $f'(x) = 0$. Multiplying both sides by $(x^2 + 2)^{\frac{1}{2}}$ we get $-x^2 - (x^2 + 2) = 0$,

which simplifies to $x^2 = -1$. Since this equation has no solutions, $f(x)$ has no interior maxima or minima.

(e) We have

$$\begin{aligned}
 f'(x) &= \frac{1}{(x^2 + 2)^{\frac{1}{2}}} \cdot \frac{d}{dx} \left[(x^2 + 2)^{\frac{1}{2}} \right] \\
 &= (x^2 + 2)^{-\frac{1}{2}} \cdot x \cdot (x^2 + 2)^{-\frac{1}{2}} \text{ (from (c))} \\
 &= x \cdot (x^2 + 2)^{-1} \\
 \text{and } f''(x) &= x \cdot \frac{d}{dx} [(x^2 + 2)^{-1}] + [(x^2 + 2)^{-1}] \frac{d}{dx}(x) \\
 &= x(-1)(x^2 + 2)^{-2} \frac{d}{dx}(x^2 + 2) + (x^2 + 2)^{-1} \\
 &= -2x^2(x^2 + 2)^{-2} + (x^2 + 2)^{-1}
 \end{aligned}$$

Candidate solutions for interior maxima or minima are where $f'(x) = 0$. The only candidate is $x = 0$, which turns out to be a minimum value of the function $f(x)$. Note that the sign of the second derivative, $f''(0) = \frac{1}{2} > 0$, identifies this as a minimum.

C-6. (a) We have

$$\begin{aligned}
 \frac{\partial}{\partial x} &= \frac{\partial}{\partial x}(x^2y) + \frac{\partial}{\partial x}(-3x) + \frac{\partial}{\partial x}(2y) \\
 &= y \cdot \frac{\partial}{\partial x}(x^2) - 3 \cdot \frac{\partial}{\partial x}(x) + 0 \\
 &= 2xy - 3 \\
 \text{and } \frac{\partial}{\partial y} &= \frac{\partial}{\partial y}(x^2y) + \frac{\partial}{\partial y}(-3x) + \frac{\partial}{\partial y}(2y) \\
 &= x^2 \cdot \frac{\partial}{\partial y}(y) + 0 + 2 \cdot \frac{\partial}{\partial y}(y) \\
 &= x^2 + 2
 \end{aligned}$$

(b) We have $\frac{\partial}{\partial x} = e^{xy} \cdot \frac{\partial}{\partial x}(xy) = ye^{xy}$. Since $f(x, y)$ is symmetric in x and y , we must also have $\frac{\partial}{\partial y} = xe^{xy}$.

(c) We have $\frac{\partial}{\partial x} = y^2 \cdot \frac{\partial}{\partial x}(e^x) - 0 = y^2e^x$ and $\frac{\partial}{\partial x} = e^x \cdot \frac{\partial}{\partial x}(y^2) - 2 = 2ye^x - 2$.

C-7. (a) The slope of the demand curve is $\frac{dq}{dp} = -2 < 0$, which makes sense because as the price goes up consumers should want to buy less.

(b) The inverse demand curve is $p = 10 - .5q$.

(c) The slope of the inverse demand curve is $-.5$, which is the inverse of the slope of the demand curve.

- C-8. (a) The inverse demand curve is $p = 20 - q$, and substituting into the profit function yields $\pi = (20 - q)q - 2q - F = 18q - q^2 - F$. Taking a derivative and setting it equal to zero gives us our candidate solution for an interior maximum: $\pi' = 0 \implies 18 - 2q = 0 \implies q^* = 9$. Substituting this back into the inverse demand curve yields $p^* = 11$, so that the profit-maximizing profit level is $\pi^* = 11 \cdot 9 - 2 \cdot 9 - F = 81 - F$.
- (b) We can see from above that the monopolist will choose not to enter for $F > 81$: zero profits are better than negative profits. Note that we would get a **corner solution** to the maximization problem in this case. The answer $q^* = p^* = 0$ does *not* show up as one of our candidate interior solutions.

C-9. In a competitive market the firm is assumed to be so small that its choice of q doesn't affect the market price p . So the firm treats p like a constant and takes a derivative of the profit function to find candidate (interior) maxima: $\pi' = 0 \implies p - C'(q) = 0 \implies p = C'(q^*)$. This says that the firm should produce until price equals marginal cost. Give $C(q)$ as above, we get $p = q + 2 \implies q^* = p - 2$.

C-10. If the monopolist chooses to produce q , the inverse demand curve establishes the maximum price as $p(q)$. Substituting this into the profit function gives $\pi(p, q) = pq - C(q) \implies \pi(q) = p(q) \cdot q - C(q)$. Taking a derivative and setting it equal to zero to find candidate (interior) maxima yields $p(q) \cdot 1 + p'(q) \cdot q - C'(q) = 0$. Given the specific demand and cost curves above, we get $(20 - q) + (-1)q - (q + 2) = 0 \implies q^* = 6$.

C-11. First simplify a bit: $U = p_B^{-\frac{1}{2}} \cdot (100Z - p_Z Z^2)^{\frac{1}{2}}$. Then take a derivative with respect to Z :

$$\begin{aligned} \frac{d}{dZ} U &= p_B^{-\frac{1}{2}} \cdot \frac{d}{dZ} \left[(100Z - p_Z Z^2)^{\frac{1}{2}} \right] \\ &= p_B^{-\frac{1}{2}} \cdot \frac{1}{2} (100Z - p_Z Z^2)^{-\frac{1}{2}} \frac{d}{dZ} (100Z - p_Z Z^2) \\ &= p_B^{-\frac{1}{2}} \cdot \frac{1}{2} (100Z - p_Z Z^2)^{-\frac{1}{2}} \cdot (100 - 2p_Z Z). \end{aligned}$$

Setting this derivative equal to zero we find that the only candidate solution is $100 - 2p_Z Z = 0 \implies Z^* = \frac{50}{p_Z}$.

- C-12. (a) Firm 1 chooses q_1 to maximize $\pi_1 = (20 - q_1 - q_2)q_1 - 2q_1 = 18q_1 - q_1^2 - q_1 q_2$.
- (b) Symmetrically, Firm 2 chooses q_2 to maximize $\pi_2 = (20 - q_1 - q_2)q_2 - 2q_2 = 18q_2 - q_2^2 - q_1 q_2$.
- (c) The difference between the two maximization problems is that q_1 is a constant in Firm 2's problem (because Firm 1 has already chosen q_1 by the time Firm 2 gets to pick q_2) but q_2 is *not* a constant in

Firm 1's problem (because Firm 1's choice of q_1 will probably affect Firm 2's choice of q_2).

- (d) Take a partial derivative of Firm 2's profit function with respect to its choice variable and set it equal to zero: $\frac{\partial}{\partial q_2} \pi_2 = 18 - 2q_2 - q_1 = 0 \implies q_2^* = 9 - .5q_1$.
- (e) Firm 1 needs to anticipate how its choice of q_1 will affect Firm 2's behavior in order to choose q_1 optimally.
- (f) Plug Firm 2's best response function into Firm 1's objective function, take a derivative, and set it equal to zero to find Firm 1's profit-maximizing choice of q_1 :

$$\begin{aligned} \frac{\partial}{\partial q_1} \pi_1 &= \frac{\partial}{\partial q_1} [18q_1 - q_1^2 - q_1(9 - .5q_1)] \\ &= 18 - 2q_1 - 9 + q_1 \\ &= 9 - q_1 \implies q_1^* = 9. \end{aligned}$$

- (g) Firm 1's optimal choice is $q_1^* = 9$, which means that Firm 2's optimal choice is $q_2^* = 9 - .5(9) = 4.5$, which means that $p = 20 - 9 - 4.5 = 6.5$, which means that $\pi_1^* = 6.5 \cdot 9 - 2 \cdot 9 = 40.5$ and $\pi_2^* = 6.5 \cdot 4.5 - 2 \cdot 4.5 = 20.25$. Firm 1 comes out ahead, so there is a first mover advantage here.

Chapter 2: Optimization and Risk

1. The expected value is $\frac{1}{6}(1) + \frac{1}{6}(2) + \frac{1}{6}(3) + \frac{1}{6}(4) + \frac{1}{6}(5) + \frac{1}{6}(6) = \frac{21}{6}$.
2. The expected value is $\frac{1}{3}(99) + \frac{2}{3}(-33) = 11$.
3. You should switch: your odds of winning will increase from $\frac{1}{3}$ to $\frac{2}{3}$. A more extreme example may help provide some intuition behind this result: assume that there are 100 doors, only one of which leads to a car; after you pick a door, Monty opens up 98 of the other doors to reveal goats and then offers you the opportunity to switch to the remaining unopened door. Doing so will increase your odds of winning from $\frac{1}{100}$ to $\frac{99}{100}$.
Still confused? Go [online](http://cartalk.cars.com/About/Monty/)¹ for a fun computer simulation of the Monty Hall problem (with accompanying discussion).
4. (a) The expected value of guessing randomly is $\frac{1}{5}(1) + \frac{4}{5}(0) = \frac{1}{5}$.
(b) If an incorrect answer is worth x , the expected value from guessing randomly is $\frac{1}{5}(1) + \frac{4}{5}(x) = \frac{1+4x}{5}$. If the teacher wants this expected value to equal zero, she must set $x = -\frac{1}{4}$.

¹<http://cartalk.cars.com/About/Monty/>

- (c) Since this makes random guessing a fair bet, it will discourage risk averse students but not risk loving students.
- (d) If you can't eliminate any answers, the expected value of guessing randomly is $\frac{1}{5}(1) + \frac{4}{5}(-\frac{1}{3}) = -\frac{1}{15}$. If you can eliminate one answer, you have a 1 in 4 chance of getting the right answer if you guess randomly, so your expected value if you can eliminate one answer is $\frac{1}{4}(1) + \frac{3}{4}(-\frac{1}{3}) = 0$. If you can eliminate two answers, you have a 1 in 3 chance of getting the right answer if you guess randomly, so your expected value if you can eliminate two answers is $\frac{1}{3}(1) + \frac{2}{3}(-\frac{1}{3}) = \frac{1}{9}$. So you need to eliminate at least two answers in order to make random guessing yield an expected value greater than zero.
5. No, they are not particularly risky. This is because of the **law of large numbers**, discussed in Section 2.1. The individual bettor plays roulette only a few times, and so faces a lot of risk. The casino plays roulette thousands of times each day, and so has a very good idea of what the overall outcome will be; since each \$1 wager has an expected payoff of only \$.95, it can expect to gain about \$.05 for every dollar wagered.
- Similarly, although insurance companies have no idea whether an individual driver is going to get into an accident this year, or whether an individual person is going to die this year, or whether an individual home is going to burn to the ground this year, the law of large numbers usually gives them a very good idea of the percentage of accidents or the percentage of deaths or the percentage of fires to expect from the hundreds of thousands of cars, lives, and homes they cover.
- So casinos or insurance companies are not necessarily any riskier as business endeavors than, say, running a photocopy shop.
6. Your expected value from Option 1 is $.8(40) + .2(-20) = 28$. Your expected value from Option 2 is $.8(-5) + .2(100) = 16$. Your expected value from Option 3 is 0. So Option 1 maximizes your expected value.
7. This is a difficult problem. For more information on it, read Raiffa's book or do some research on Bayes's Rule.
8. (a) Your expected value from bidding b in either type of auction is

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

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

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

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

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

- (c) Chapter 12 discusses this in more detail.
9. (a) The expected cost is $(.80)(0) + (.20)(\$20) = \4 .
- (b) The city could double the amount of the ticket from \$20 to \$40. This would mean that the expected value of risking it is still \$4: $(.90)(0) + (.10)(\$40) = \4 .

Chapter 3: Optimization over Time

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

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

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

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

3. The perpetuity formula says that the present value of a perpetuity paying x every year is $\frac{x}{s}$. This is like living off the interest because if you put $\frac{x}{s}$ in the bank, every year you will get interest of $r \cdot \frac{x}{s} = x$, so with this principal you can finance annual consumption of x forever.
4. (a) Plug \$1 million, 5%, and 20 years into the annuity formula to get about \$12.5 million as the present value of the annuity.
- (b) Plug \$1 million and 5% into the perpetuity formula to get \$20 million as the present value of the perpetuity. Note that the extra payments you get—\$1 million annually beginning in year 21—are only worth about \$7.5 million in present value terms!

- (c) Increasing r will make the lump sum payment more attractive, and decreasing r will make the annual payments more attractive. Trial and error yields $s \approx .075$ as the interest rate that makes the two payoffs equal in present value terms.
 - (d) The hard way to do this is to just calculate the present value of each payment and then add them all together. Easy way #1 is to realize that the difference between the end-of-year payments and the beginning-of-year payments is just an extra payment at the beginning of the first year and a lost payment at the end of the 20th year. The present value of \$1 million today is \$1 million, and the present value of \$1 million at the end of 20 years is \$380,000. Their difference is \$620,000, so adding this to the answer from (a) yields \$13.08 million. Easy way #2 is to see that the answer from (a) is the right answer from the perspective of one year ago, so using the future value of a lump sum formula to push this answer one year into the future gives us $\$12.46(1.05) = \13.08 million.
5. (a) Note: The elements of this problem featuring logarithms are not fair game for the exam. Having said that: Solving $2x = x(1.10)^t$ using logarithms yields $t = 7.27$, pretty close to the Rule of 72 estimate of 7.2.
- (b) The Rule of 72 predicts 14.4 years at 5% and .72 years at 100%. The correct answer at 100% is 1 year (obviously, since the interest rate is 100%), so the Rule of 72 is not such a good approximation. The correct answer at 5% comes from solving $2x = x(1.05)^t$ using logarithms. We get $t \approx 14.2$, which is quite accurate.
 - (c) They suggest that the Rule of 72 works well for small interest rates, but not for large ones.
6. (a) For Investment #1, use the annuity formula to get a present value of about \$772 at a 5% interest rate. For Investment #2, the brute force way to calculate the present value is to calculate the present value of each of the 6 lump sum payments and then add them up to get about \$835. A more elegant way is to note that Investment #2 is equivalent to a ten-year annuity *minus* a four-year annuity. You can therefore use the annuity formula to calculate the present value of the ten-year annuity (\$1,544) and the four-year annuity (\$709). Subtracting one from the other gives a present value for Investment #2 of \$835.
- Investment #2 is the better option at a 5% interest rate.
- (b) Following the same process described above, Investments #1 and #2 have present values of \$502 and \$433, respectively. So Investment #1
 - (c) Higher interest rates favor Investment #1, in which you essentially forfeit money in the distant future in order to get more money in the

immediate future. Since higher interest rates make the future less important, they make Investment #1 more attractive.

- (d) There are many real-life decisions with similar features; for example, Investment #1 could be going to work right out of high school, and Investment #2 could be going to college for 4 years first to increase your earnings potential.
7. (a) The king was supposed to deliver $2^{63} \approx 1,000,000,000,000,000,000,000$ grains of rice for the last square.
- (b) The pond was half-covered on the 29th day.
- (c) The answer is the height of a single sheet of paper multiplied by $2^{39} \approx 1,000,000,000,000$. If 1,000 pages makes an inch, then this gives us 1,000,000,000 inches, or about 83 million feet, or about 16,000 miles.

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. (a) Using the annuity formula we get a present value of about \$6 trillion.
- (b) The expected damages are $\frac{1}{3}(6) + \frac{1}{3}(3) + \frac{1}{3}(0) \approx \3 trillion.
- (c) Plug \$3 trillion into the present value of a lump sum formula to get a present value of \$59 billion.
- (d) Using the present value of a lump sum formula, we get \$9,130.
9. To maximize your present value you need to compare the return you'll get from "investing in the fish" (or the trees, or the oil) with the return you'll get from investing in the bank. Investing in the bank means catching the fish, cutting down the trees, or selling the oil and putting the proceeds in the bank. Investing in the fish means letting the fish grow and reproduce so there will be more fish next year; investing in the trees means letting the trees grow so there will be more lumber next year; investing in the oil means keeping the oil in the hopes that the price will go up next year.
10. (a) Plug \$15 million, 200 years, and 2% into the future value of a lump sum formula to get a current bank account balance of \$787 million.
- (b) Plug \$15 million, 200 years, and 8% into the future value of a lump sum formula to get a current bank account balance of \$72.6 trillion

11. (a) The approximate annual interest rate is 6%.
- (b) Use the future value of a lump sum formula to calculate how much money we'll have at the end of 12 months if we put \$100 in the bank at a monthly interest rate of 0.5%, i.e. $s = 0.005$: $FV = \$100(1.005)^{12} \approx 106.17$. So a monthly interest rate of 0.5% actually corresponds to an annual interest rate of 6.17%.
- (c) The approximate monthly interest rate is 0.5%.
- (d) Use the future value of a lump sum formula to set $x(1+s)^{12} = x(1.06)$ and solve for s using logarithms. This yields $s \approx .00487$, i.e., an interest rate of about 0.487%.

12. To calculate the present value of the yearly option, we need to split the payment into three parts: the \$500,000 you get immediately (which has a present value of \$500,000), the 28-year annuity (which, using the annuity formula, has a present value of about \$3.35 million), and the payment in year 29 (which, using the present value of a lump sum formula, has a present value of about \$460,400). We add these up to get a present value of about \$4.31 million.

To calculate the present value of the weekly option, we first need to calculate a weekly interest rate that is in accordance with a yearly interest rate of 6%. A decent approximation is to divide the yearly interest rate by the number of weeks in a year (52), yielding a weekly interest rate of about 0.115%, i.e., $s \approx 0.00115$.² Using this value of s in the annuity formula along with \$5,000 and 2,000 weeks gives us a present value of about \$3.91 million for the weekly option.

Finally, the present value of receiving \$3.5 million immediately is \$3.5 million. So the best option is the yearly payments.

13. (a) Plug \$100 and 5% into the perpetuity formula to get a present value of \$2000.
- (b) Plug \$84 and 5% into the perpetuity formula to get a present value of \$1680. Adding this to the \$400 you get from catching 400 fish today and you get a present value of \$2080, which is higher than the present value of the maximum sustainable yield policy.
14. (a) No, this is a sunk cost.
- (b) Use the present value of a lump sum formula to get a present value of $\frac{\$500}{1.05} \approx \476.19 .
- (c) Use the present value of a perpetuity formula to get a present value of $\frac{\$20}{.05} = \400 . So it's better to rent.

²If you don't like approximations, here's how to use logarithms to get the true value. If you put x in the bank at a weekly interest rate of $100 \cdot s\%$, after one year (52 weeks), you'll have $x(1+s)^{52}$ in the bank. We want this to equal $x(1.06)$, since this follows from a yearly interest rate of 6%. Setting $x(1+s)^{52} = x(1.06)$ and solving for s using logarithms gives us $s \approx .00112$, i.e., an interest rate of about 0.112%.

- (d) “Jack, I disagree with you. Instead of paying \$450 today to buy a steam-cleaner, I’d rather put that \$450 in the bank and ‘live off the interest’. At the end of every year I’d have \$22.50 in interest, which would pay for the annual rental of a steam-cleaner *and* leave me with \$2.50 left over for wild parties.” (Alternately, you could put \$50 towards a wild party today and put the remaining \$400 in the bank; the interest payments would then be \$20 per year, exactly enough to rent a steam-cleaner.)

Chapter 4: *Math*: Trees and Fish

1. These are all answered (to the best of my abilities) in the text.

Chapter 5: More Optimization over Time

1. The approximation is $14.4\% - 8\% = 6.4\%$. The actual answer comes from the formula, and gives a result of 5.9% .
2. This is explained in the text.
3.
 - (a) Use the nominal interest rate and the future value formula to get a bank account balance of about \$259.37.
 - (b) Use the inflation rate and the future value formula to get an apple fritter price of about \$1.63.
 - (c) Today you have \$100 and fritters cost \$1, so you can buy $x = 100$ of them. In ten years you’ll have \$259.37 and fritters will cost \$1.63, so you’ll be able to buy about $y = 159$ of them. So we can calculate $z \approx 59$.
 - (d) The rule of thumb approximation says that the real interest rate should be about $10\% - 5\% = 5\%$. The actual value is $\frac{1+.1}{1+.05} - 1 \approx .048$, i.e., 4.8% .
 - (e) If you put \$100 in the bank at this interest rate, after 10 years you’d have about \$159. So you get $z = 59$ as your gain in purchasing power.
4. **Use real when your payments are inflation-adjusted.**
 - (a) Use the annuity formula and the real interest rate (about $6 - 4 = 2\%$) to get a present value of about \$470.
 - (b) Use the perpetuity formula and the real interest rate (about $6 - 4 = 2\%$) to get a present value of about \$5,000.

Use real to calculate future purchasing power.

You’re not figuring correctly because you’re forgetting that prices are going to rise. Yes, you’ll have 10 times more money, but you won’t be

able to buy 10 times more stuff. Using the real interest rate and the future value formula, we get a future value of \$11 million, or about 2.2 times more purchasing power.

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

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

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

Chapter 6: Transition: Arbitrage

1. This is answered (to the best of my abilities) in the text.
2. If you lend the company \$100, the expected value of your investment is $.90(120) + .10(0) = \$108$, meaning that your expected rate of return is 8%. If government bonds have a 3% expected rate of return, the risk premium associated with the junk bonds is $8\% - 3\% = 5\%$.
3. (a) The *actual* rates of return turned out to be different, but it could still be true that at any point in time over the last twenty years the *expected* rates of return were the same. Economic reasoning suggests that they should have been; otherwise investors would not have been acting optimally. In terms of its effect on this logic, the fact that the actual rates of return turned out to be different is no more cause for concern than the fact that some lottery tickets turn out to be winners and some turn out to be losers.
- (b) Again, actual rates of return can be different even though expected rates of return are the same at any moment in time. For example, there may be unexpected advances (or setbacks) in the development of electric cars or other alternatives to gasoline-powered cars.

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

- (c) Companies that issue junk bonds are financially troubled; if they go bankrupt, you lose both the interest and the principal. So even though junk bonds may offer 20% interest, their expected rate of return is much less than 20%, and therefore much closer to the expected rate of return of U.S. Treasury bonds. Also, the U.S. Treasury isn't likely to go bankrupt, meaning that two assets don't have comparable risk. So it is likely that there is a risk premium associated with the junk bonds.
- (d) The answer here is that the rate of inflation is much higher in Turkey than in the U.S. So even though the two banks pay different nominal interest rates, their expected real interest rates may be equal.

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

Appendix D

Answer Key for Part II: One v. One, One v. Many

Chapter 7: Cake-Cutting

1. Arguably, differences in opinion make fair division easier, not harder. For example, if one child only likes vanilla and the other child only like chocolate, then cake division is, well, a piece of you-know-what.
2. This question is not fair game for exam, but is solvable via induction.
3. The Coase Theorem says that people who are free to trade have a strong incentive to trade until they exhaust all possible gains from trade, i.e., until they complete all possible Pareto improvements and therefore reach a Pareto efficient allocation of resources. The implication for the cake-cutting problem is that a “mom” whose sole concern is efficiency can divide the cake up however she wants—as long as the children can trade, they should be able to reach a Pareto efficient allocation regardless of the starting point. For example, if you give the chocolate piece to the kid who loves vanilla and the vanilla piece to the kid who loves chocolate, they can just trade pieces and will end up at a Pareto efficient allocation.
4. To show that envy-free implies proportional, we will show that not proportional implies not envy-free. If a cake division is not proportional, one child gets less than $\frac{1}{n}$ th of the cake (in her estimation). This means that (according to her estimation) the other $(n - 1)$ children get more than $\frac{n-1}{n}$ th of the cake. Therefore at least one of the other children must have a piece bigger than $\frac{1}{n}$ th, meaning that the cake division is not envy-free.
5. (a) Alice would spend 4 hours gathering veggies and 2 hours fishing, providing her with 4 veggies and 4 fish. Bob would do exactly the opposite (4 hours fishing, 2 hours gathering veggies) and would also end up with 4 of each.

- (b) If they specialize, Alice spends 6 hours fishing, so she gets 12 fish; Bob spends 6 hours hunting, so he gets 12 veggies. Then they split the results, so each gets 6 fish and 6 veggies, a clear Pareto improvement over the no-trade situation.
- (c) Yes.
- (d) They would allocate their time as before, but now Alice would get 12 fish and 12 veggies and Bob would get 4 fish and 4 veggies.⁷
- (e) If they specialize as described in the problem, Alice ends up with 18 fish and 9 veggies, and Bob ends up with 1 fish and 10 veggies. After they trade, Alice ends up with 13 fish and 13 veggies, and Bob ends up with 6 fish and 6 veggies, another Pareto improvement!
- (f) Yes.
- (g) Alice must give up 2 fish to get one vegetable, and must give up 0.5 veggies to get one fish.
- (h) Bob must give up 0.5 fish to get one vegetable, and 2 veggies to get one fish.
- (i) Alice.
- (j) Bob. When they concentrate on the items for which they are the least-cost producer, they can both benefit from trade even though Alice has an absolute advantage over Bob in both fishing and gathering veggies. This is the concept of **comparative advantage**.

Chapter 8: Economics and Social Welfare

1.
 - (a) Pareto inefficient means that it's possible to make someone better off without making anyone else worse off; in other words, there's a "free lunch". Example: It is Pareto inefficient to give Tom all the chicken and Mary all the veggies because Tom's a vegetarian and Mary loves chicken.
 - (b) A Pareto improvement is a reallocation of resources that makes one person better off without making anyone else worse off. Example: giving Tom the veggies and Mary the chicken is a Pareto improvement over giving Tom the chicken and Mary the veggies.
 - (c) Pareto efficient means that there is no "free lunch", i.e., it's not possible to make someone better off without making anyone else worse off. Example: Giving Tom the veggies and Mary the chicken is a Pareto efficient allocation of resources.
2.
 - (a) A Pareto efficient allocation of resources may not be good because of equity concerns or other considerations. For example, it would be Pareto efficient for Bill Gates to own everything (or for one kid to get the whole cake), but we might not find these to be very appealing resource allocations.

- (b) A Pareto inefficient allocation is in some meaningful sense bad because it's possible to make someone better off without making anybody else worse off, so why not do it?
- 3. The claim that any Pareto efficient allocation is a Pareto improvement over any Pareto inefficient allocation is not true. For example, giving one child the whole cake is a Pareto efficient allocation, and giving each child one-third of the cake and throwing the remaining third away is Pareto inefficient, but the former is not a Pareto improvement over the latter.
- 4. Consider a division problem such as the division of cake or the allocation of fishing quotas.
 - (a) When people trade they bring about Pareto improvements—why would any individual engage in a trade unless it made him or her better off? Pareto improvements are a good thing in and of themselves, and if you get enough of them then you end up with a Pareto efficient allocation of resources.
 - (b)
 - i. No. There are multiple Pareto efficient allocations.
 - ii. Initial allocations are a matter of equity; economists tend to focus on efficiency. As long as there are opportunities to trade, a Pareto efficient outcome will result *regardless of the initial allocation*.
- 5. (a) No. One Pareto improvement would be for the publisher to provide me with the article for free: I'm better off, and they're not any worse off. Another Pareto improvement is for me to pay \$.25 for the article: I'm not any worse off, and the publisher is better off by \$.25.
- (b) Yes. Perfect price discrimination by a monopolist is Pareto efficient because it's not possible to make any of the customers better off without making the monopolist worse off.
- (c) The monopolist would need detailed information about all its different customers, e.g., the maximum amount each customer is willing to pay for each different article; such information is not readily available. There might also be a problem with resale, i.e., with some individuals paying others to purchase articles for them.

Chapter 9: Sequential Move Games

- 1. This is explained (to the best of my abilities) in the text. The basic idea is that you need to anticipate your rival's response.
- 2. (a) Backward induction predicts an outcome of (3, 3).
(b) No; a Pareto improvement is (6, 6).
- 3. (a) Backward induction predicts an outcome of (8, 3).

- (b) Yes, it is Pareto efficient.
- 4. Tipping at the beginning of the meal is problematic because then the waitress has no incentive to provide good service. (The tip is already sunk.) Tipping at the end of the meal is problematic because then the customer has no incentive to provide the tip. (The service is already sunk.)
- 5. (a) Backward induction predicts an outcome of (M: 35, PE: 5).
(b) Yes.
(a) Backward induction predicts an outcome of (M: 70, PE: 0).
(b) No; a Pareto improvement is (M: 100, PE: 0).
- 6. (a) If there are 10 sticks on the table, you should be player 2. Whenever your opponent takes 1 stick, you take 2; when he takes 2 sticks, you take 1. So you can force your opponent to move with 7 sticks, then 4 sticks, then 1 stick—so you win!
(b) Hint: The above answer suggests a general strategy to follow.
(c) I only have a partial answer, so let me know if you solve all or part of this!
- 7. (a) With one period, Player 1 offers Player 2 a sliver, and Player 2 accepts. With two periods, Player 1 offers Player 2 half the cake, and Player 2 accepts. (Both know that if Player 2 refuses, half the cake melts, Player 2 will offer Player 1 a sliver of the remaining half, and Player 1 will accept.) With three periods, Player 1 offers Player 2 one-quarter of the cake, and Player 2 accepts. (Both know that if Player 2 refuses, she'll have to offer Player 1 at least half of the remaining half, meaning that she'll get at most one-quarter.)
(b) With one period, Player 1 offers Player 2 a sliver, and Player 2 accepts. With two periods, Player 1 offers Player 2 two-thirds of the cake, and Player 2 accepts. (Both know that if Player 2 refuses, one-third of the cake melts, Player 2 will offer Player 1 a sliver of the remaining two-thirds, and Player 1 will accept.) With three periods, Player 1 offers Player 2 two-ninths of the cake, and Player 2 accepts. (Both know that if Player 2 refuses, she'll have to offer Player 1 at least two-thirds of the remaining two-thirds, meaning that she'll get at most two-ninths.)
(c) This is the magic of the Coase Theorem. It is in neither player's interest to let the cake melt away, so they have a strong incentive to figure things out at the beginning and bring about a Pareto efficient outcome. You can see the same phenomenon at work in labor disputes and lawsuits, many of which get settled before the parties really begin to hurt each other.

8. I'd be happy to look over your work if you do this.
9. (a) The naive outcome is for X to choose C, Y to choose H, and Z to choose T, producing the “naive outcome” at the top of the game tree.
- (b) No. If X and Y choose C and H, Z will choose F because this produces a better outcome for Z: FT is better than TQ! (But now backward induction kicks in: Y anticipates this, and so Y will choose G instead of H—GH is better than HQ. But X anticipates this, and so knows that a choice of C will result in CQ. X then uses backward induction to solve the bottom half of the tree—Z will choose F in the top part and H in the lower part, so Y will choose H because HG is better than FG—and determine that a choice of T will result in TC. Since X prefers TC to CQ, X chooses T in the first round, leading Y to choose H and Z to choose F.
- (c) Backward induction leads to a result of (TC, HG, FQ).
- (d) This is not Pareto efficient: the “naive” strategies produce better outcomes for all three teams!
- (e) Statement #1 is false because each team's choice in the first round will have strategic implications for its options in the second round. Statement #2 is true because each team's choice in the second round has no further ramifications; since there are no more rounds, in the second round each team faces a simple decision tree.
- (f) This is a time-consuming problem. Thanks to Kieran Barr for finding two strategies that yield this same outcome!
10. (a) The answer here depends on your assumptions. See below for my take on it. . . .
- (b) The situation is Pareto inefficient.
- (c) The key issue here is that Butch Cassidy is a bank robber, and hence cannot be bound to contracts or other agreements. Sure, Harriman could pay him the money, but what guarantee does he have that this will make Butch stop robbing his train? A more likely outcome is that Butch will take the money and continue to rob the train, and then Harriman will be out even more money. So Harriman hires the superposse instead, even though both he and Butch would be better off with an alternative outcome.
- (d) Contracts can help by forcing players to act in certain ways; then the Coase theorem allows them to negotiate an efficient outcome. The Coase Theorem doesn't work in the case of Butch Cassidy because he's an outlaw: there's no way to bind an outlaw to an enforceable contract.
11. (a) If you make eye contact with the driver, the driver will pretend that she's not going to stop, and then you'll get scared and won't go for it, so then the driver *won't* stop.

- (b) The game tree here has you choosing to look or not look. If you choose not to look, the driver chooses to stop or not, and the payoffs are obvious. If you choose to look, the driver chooses to stop or not, and in each of those situations you must choose whether or not to push the issue.
- 12.
- (a) The game tree is pictured and described in the text.
 - (b) As discussed in the text, backward induction predicts that Player 1 will immediately choose \$2 and end the game, yielding an outcome of $(2, 0)$.
 - (c) No. There are many Pareto improvements, e.g., $(2, 2)$.
 - (d) You can do this with induction; this exercise also suggests why backward induction has the name it does.
 - (e) This is a truly difficult philosophical question. If you're interested, there's an interesting chapter (and a great bibliography) on this topic, in the guise of "the unexpected hanging", in Martin Gardner's 1991 book, *The Unexpected Hanging, and Other Mathematical Diversions*.
13. Well, the exam can't be on Friday, because then on Thursday night you'd think, "Aha! The exam's got to be Friday!" So then you wouldn't be surprised; so the exam can't be on Friday. But then the exam can't be on Thursday, because then on Wednesday night you'd think, "Aha! The exam can't be on Friday, so it's got to be Thursday!" So then you wouldn't be surprised; so the exam can't be on Thursday. But then the exam can't be on Wednesday, or Tuesday, or even Monday. An apparently non-controversial statement by your teacher turns out to be quite treacherous!

Chapter 10: Simultaneous Move Games

1.
 - (a) A good prediction is that everybody would drive to work because driving is a dominant strategy: no matter what everybody else does, you always get there 20 minutes faster by driving.
 - (b) This outcome is not Pareto efficient because the commute takes 2 hours; a Pareto improvement would be for everybody to take the bus, in which case the commute would only take 40 minutes.
 - (c) The central difficulty is *not* that you don't know what others are going to do; you have a dominant strategy, so the other players' strategies are irrelevant for determining your optimal strategy.
 - (d) A reasonable mechanism might be passing a law that everybody has to take the bus or pay a large fine.
2.
 - (a) A good prediction is that everybody will invest in the private good because it's a dominant strategy: no matter what everybody else does, you always get \$1 more by investing privately.

- (b) This outcome is not Pareto efficient because each player only gets a return of \$2; a Pareto improvement would be for everybody to invest in the public good, in which case each player would get a return of \$10.
 - (c) The central difficulty is *not* that you don't know what others are going to do; you have a dominant strategy, so the other players' strategies are irrelevant for determining your optimal strategy.
 - (d) A reasonable mechanism might be passing a law that everybody has to invest in the public good or pay a large fine.
3. (a) Playing D is a dominant strategy for each player, so we can expect an outcome of (D, D).
 - (b) Using backward induction, we start at the end of the game, i.e., in the second round. Playing D is a dominant strategy for each player in this round, so we can expect an outcome of (D, D) in the second round. But the players will anticipate this outcome, so playing D becomes a dominant strategy in the first round too! As a result, the expected outcome is for both players to play D both times.

Chapter 11: Iterated Dominance and Nash Equilibrium

1. This is explained to the best of my abilities in the text. The key idea is to anticipate your opponent's behavior.
2. This is a hard philosophical problem. Show your work to me if you decide to tackle it.
3. U is not the best strategy for Player 1 if Player 2 plays R, M is not the best strategy for Player 1 if Player 2 plays C, and D is not the best strategy for Player 1 if Player 2 plays R. Similarly, there are no strictly dominant strategies for Player 2: L is not the best strategy for Player 2 if Player 1 plays U, C is not the best strategy for Player 2 if Player 1 plays U, and R is not the best strategy for Player 2 if Player 1 plays D.
4. (a) Here U is dominated by M for player 1, then C is dominated by L (or R) for player 2, then D is dominated by M for player 1, then R is dominated by L for player 2. The result: (M, L), with a payoff of (4,8). This is also the unique Nash equilibrium of the game; it is not a Pareto efficient outcome because of (D, R).
- (b) Here D is dominated by M for player 1, then C is dominated by L for player 2, then U is dominated by M for player 1, then L is dominated by R for player 2. The result: (M, R), with a payoff of (6,2). This is also the unique Nash equilibrium of the game; it is not a Pareto efficient outcome because of (U, C).

		Player 2		
		R	P	S
Player 1	R	0,0	-1,1	1,-1
	P	1,-1	0,0	-1,1
	S	-1,1	1,-1	0,0

Table D.1: The payoff matrix for the game “Rock, Paper, Scissors”

- (c) There are no strictly dominated strategies. The NE are (U, C) , (D, L) and (M, R) .
- (d) Here M is dominated by U for player 1, then L and R are dominated by C for player 2, then D is dominated by U for player 1. The result: (U, C) , with a payoff of (5,4). This is also the unique Nash equilibrium of the game; it is a Pareto efficient outcome.
- (e) Here L is dominated by C for player 2, and that is as far as iterated dominance can take us. We do not get a unique prediction for the outcome of this game. (All that we can say is that a rational player 2 would never play L.) With Nash, the NE are (M, C) and (D, R) . Note that these Nash equilibria are a subset of the iterated dominance solutions; see the next problem for details.
5. (a) The payoff matrix is shown in Table D.1.
- (b) Iterated dominance does not help you solve this game because there are no dominated strategies.
- (c) In accordance with intuition, the NE is for both players to choose randomly among the three strategies.
6. Not fair game for the exam.
- 7.

		Entrant	
		Enter	Stay Out
Monopolist	War	10, -10	100, 0
	Peace	35, 5	100, 0

The Nash equilibria are (War, Stay Out) and (Peace, Enter). This suggests that backward induction is in fact a refinement or strengthening of Nash equilibrium (which it is, namely **subgame perfect Nash equilibrium**).

8. Not fair game for the exam.
9. Yes! As long as the interest rate is sufficiently low, the players can also cooperate by taking turns: (C, D), (D, C), (C, D),... Instead of gaining \$1 every stage (which is the result with the trigger strategies), each player now gains \$5 every two stages. As an exercise, you can formally define this strategy (what happens if the other player doesn't cooperate?) and determine the interest rates that allow this strategy as a Nash equilibrium and those that make this strategy a Pareto improvement over the trigger strategy. There are also plenty of other strategies, e.g., tit-for-tat, that you can play around with if you wish.
10. One possibility is to proceed as follows: U is weakly dominated by M for Player 1, and then R is weakly dominated by L for Player 2, and then M is (strictly) dominated by D for Player 1, yielding a prediction of (D, L). But another possibility is to proceed like this: M is weakly dominated by D for Player 1, then L is weakly dominated by R for Player 2, then U is (strictly) dominated by D for Player 1, yielding a prediction of (D, R). Conclusion: the order of elimination matters for iterated weak dominance!

Calculus Problems

C-1. Player 1 chooses p to maximize

$$\begin{aligned} E(\pi_1) &= p[q(1) + (1-q)(0)] + (1-p)[q(0) + (1-q)(3)] \\ &= pq + (1-p)(1-q)(3). \end{aligned}$$

Similarly, player 2 chooses q to maximize

$$\begin{aligned} E(\pi_2) &= q[p(3) + (1-p)(0)] + (1-q)[p(0) + (1-p)(1)] \\ &= 3pq + (1-q)(1-p). \end{aligned}$$

Now, we want to find p and q that form a Nash equilibrium, i.e., that are mutual best responses. To do this, we take derivatives and set them equal to zero.

So: player 1 wants to choose p to maximize $E(\pi_1) = pq + 3(1-p)(1-q)$. Any value of p that maximizes this is either a corner solution (i.e., one of the pure strategies $p = 1$ or $p = 0$) or an interior solution with $0 < p < 1$, in which case the partial derivative of $E(\pi_1)$ with respect to p must be zero:

$$\frac{\partial E(\pi_1)}{\partial p} = 0 \implies q - 3(1-q) = 0 \implies 4q = 3 \implies q = \frac{3}{4}.$$

This tells us that *any* interior value of p is a candidate maximum as long as $q = \frac{3}{4}$. Mathematically, this makes sense because if $q = \frac{3}{4}$ then player

1's expected payoff (no matter what his choice of p) is always

$$E(\pi_1) = pq + 3(1-p)(1-q) = \frac{3}{4}p + 3\frac{1}{4}(1-p) = \frac{3}{4}.$$

If player 2 chooses $q \neq \frac{3}{4}$ then player 1's best response is to choose $p = 1$ (if $q > \frac{3}{4}$) or $p = 0$ (if $q < \frac{3}{4}$).

We can now do the math for player 2 and come up with a similar conclusion. Player 2's expected payoff is $\pi_2 = 3pq + (1-q)(1-p)$. Any value of q that maximizes this is either a corner solution (i.e., one of the pure strategies $q = 1$ or $q = 0$) or an interior solution with $0 < q < 1$, in which case

$$\frac{\partial E(\pi_2)}{\partial q} = 0 \implies 3p - (1-p) = 0 \implies 4p = 1 \implies p = \frac{1}{4}.$$

So if player 1 chooses $p = \frac{1}{4}$ then any choice of q is a best response for player 2. But if player 1 chooses $p \neq \frac{1}{4}$ then player 2's best response is a pure strategy: if player 1 chooses $p > \frac{1}{4}$ then player 2's best response is to choose $q = 1$; if player 1 chooses $p < \frac{1}{4}$ then player 2's best response is to choose $q = 0$.

Now we can put our results together to find the Nash equilibria in this game. If player 1's choice of p is a best response to player 2's choice of q then either $p = 1$ or $p = 0$ or $q = \frac{3}{4}$ (in which case any p is a best response). And if player 2's choice of q is a best response to player 1's choice of p then either $q = 1$ or $q = 0$ or $p = \frac{1}{4}$ (in which case any q is a best response).

Three choices for player 1 and three choices for player 2 combine to give us nine candidate Nash equilibria:

Four pure strategy candidates : $(p = 1, q = 1), (p = 1, q = 0), (p = 0, q = 1), (p = 0, q = 0)$.

One mixed strategy candidate : $(0 < p < 1, 0 < q < 1)$.

Four pure/mixed combinations : $(p = 1, 0 < q < 1), (p = 0, 0 < q < 1), (0 < p < 1, q = 1), (0 < p < 1, q = 0)$.

We can see from the payoff matrix that there are two Nash equilibria among the four pure strategy candidates: $(p = 1, q = 1)$ and $(p = 0, q = 0)$. The other other two are not Nash equilibria. We can also see that the four pure/mixed combinations are not best responses; for example, $(p = 1, 0 < q < 1)$ is not a Nash equilibrium because if player 1 chooses $p = 1$ then player 2's best response is to choose $q = 1$, not $0 < q < 1$.

But the mixed strategy candidate does yield a Nash equilibrium: player 1's choice of $0 < p < 1$ is a best response as long as $q = \frac{3}{4}$. And player 2's choice of $0 < q < 1$ is a best response as long as $p = \frac{1}{4}$. So the player's strategies are mutual best responses if $(p = \frac{1}{4}, q = \frac{3}{4})$.

So this game has three Nash equilibria: two in pure strategies and one in mixed strategies.

C-2. Player 1 chooses p to maximize

$$\begin{aligned} E(\pi_1) &= p[q(0) + (1-q)(-1)] + (1-p)[q(-2) + (1-q)(1)] \\ &= p(q-1) + (1-p)(1-3q). \end{aligned}$$

Similarly, player 2 chooses q to maximize

$$\begin{aligned} E(\pi_2) &= q[p(0) + (1-p)(1)] + (1-q)[p(5) + (1-p)(-2)] \\ &= q(1-p) + (1-q)(7p-2). \end{aligned}$$

Now, we want to find p and q that form a Nash equilibrium, i.e., that are mutual best responses. To do this, we take derivatives and set them equal to zero.

So: player 1 wants to choose p to maximize $E(\pi_1) = p(q-1) + (1-p)(1-3q)$. Any value of p that maximizes this is either a corner solution (i.e., one of the pure strategies $p = 1$ or $p = 0$) or an interior solution with $0 < p < 1$, in which case the partial derivative of $E(\pi_1)$ with respect to p must be zero:

$$\frac{\partial E(\pi_1)}{\partial p} = 0 \implies q - 1 - (1 - 3q) = 0 \implies 4q = 2 \implies q = \frac{1}{2}.$$

This tells us that *any* interior value of p is a candidate maximum as long as $q = \frac{1}{2}$. Mathematically, this makes sense because if $q = \frac{1}{2}$ then player 1's expected payoff (no matter what his choice of p) is always

$$E(\pi_1) = p(q-1) + (1-p)(1-3q) = -\frac{1}{2}p + (1-p)\frac{-1}{2} = -\frac{1}{2}.$$

If player 2 chooses $q \neq \frac{1}{2}$ then player 1's best response is to choose $p = 1$ (if $q > \frac{1}{2}$) or $p = 0$ (if $q < \frac{1}{2}$).

We can now do the math for player 2 and come up with a similar conclusion. Player 2's expected payoff is $q(1-p) + (1-q)(7p-2)$. Any value of q that maximizes this is either a corner solution (i.e., one of the pure strategies $q = 1$ or $q = 0$) or an interior solution with $0 < q < 1$, in which case

$$\frac{\partial E(\pi_2)}{\partial q} = 0 \implies 1 - p - (7p - 2) = 0 \implies 8p = 3 \implies p = \frac{3}{8}.$$

So if player 1 chooses $p = \frac{3}{8}$ then any choice of q is a best response for player 2. But if player 1 chooses $p \neq \frac{3}{8}$ then player 2's best response is a pure strategy: if player 1 chooses $p > \frac{3}{8}$ then player 2's best response is to choose $q = 0$; if player 1 chooses $p < \frac{3}{8}$ then player 2's best response is to choose $q = 1$.

Now we can put our results together to find the Nash equilibria in this game. If player 1's choice of p is a best response to player 2's choice of q then either $p = 1$ or $p = 0$ or $q = \frac{1}{2}$ (in which case any p is a best response). And if player 2's choice of q is a best response to player 1's choice of p then either $q = 1$ or $q = 0$ or $p = \frac{3}{8}$ (in which case any q is a best response).

Three choices for player 1 and three choices for player 2 combine to give us nine candidate Nash equilibria:

Four pure strategy candidates : $(p = 1, q = 1), (p = 1, q = 0), (p = 0, q = 1), (p = 0, q = 0)$.

One mixed strategy candidate : $(0 < p < 1, 0 < q < 1)$.

Four pure/mixed combinations : $(p = 1, 0 < q < 1), (p = 0, 0 < q < 1), (0 < p < 1, q = 1), (0 < p < 1, q = 0)$.

We can see from the payoff matrix that there are no Nash equilibria among the four pure strategy candidates: We can also see that the four pure/mixed combinations are not best responses; for example, $(p = 1, 0 < q < 1)$ is not a Nash equilibrium because if player 1 chooses $p = 1$ then player 2's best response is to choose $q = 0$, not $0 < q < 1$.

But the mixed strategy candidate does yield a Nash equilibrium: player 1's choice of $0 < p < 1$ is a best response as long as $q = \frac{1}{2}$. And player 2's choice of $0 < q < 1$ is a best response as long as $p = \frac{3}{8}$. So the player's strategies are mutual best responses if $(p = \frac{3}{8}, q = \frac{1}{2})$.

So this game has one (mixed strategy) Nash equilibrium.

Chapter 12: Application: Auctions

1. (a) Auctions pit different suppliers against each other, and their individual incentives lead them to drive down the price. This helps ensure that Company X will not be paying much more for springs than it costs the suppliers to produce them.
- (b) The example in the ad above may not be as impressive as it sounds because of the potential for gaming: if Company X knows that a number of firms can produce the springs for about \$350,000, it has essentially nothing to lose by indicating a willingness-to-pay of \$500,000—or even \$1,000,000—because the auction dynamics will drive the price down toward \$350,000. An analogy may help: say I want to purchase a \$20 bill. As long as there are enough competitive bidders, I can more-or-less fearlessly say that I'm willing to pay up to \$1,000 for that \$20 bill; competitive pressures will force the winning bid down to about \$20.
2. (a) The intuition can be seen from an example: say you're willing to pay up to \$100, but you only bid \$90. Let y be the highest bid not

including your bid. If $y < 90$ then you win the auction and pay y ; in this case, bidding \$90 instead of \$100 doesn't help you or hurt you. If $y > 100$ then you lose the auction and would have lost even if you bid \$100; again, bidding \$90 instead of \$100 doesn't help you or hurt you. But if y is between \$90 and \$100 (say, $y = \$95$) then bidding \$90 instead of \$100 actively hurts you: you end up losing the auction when you would have liked to have won it. (You had a chance to get something you value at \$100 for a payment of only \$95, but you didn't take it.)

- (b) Again, the intuition can be seen in the same example in which you're willing to pay up to \$100. Assume that you bid \$110 and that y is the highest bid not including your bid. If $y < \$100$ then you win the auction and pay y ; in this case bidding \$110 instead of \$100 doesn't help you or hurt you. If $y > \$110$ then you lose the auction; again, bidding \$110 instead of \$100 doesn't help you or hurt you. But if y is between \$100 and \$110 (say, $y = \$105$) then bidding \$110 instead of \$100 actively hurts you: you end up winning the auction when you would have liked to have lost it. (You pay \$105 for something you only value at \$100.)
3. You should bid less than your true value. If your true value is, say, \$100, then you are indifferent between having the object and having \$100. If you bid \$100, winning the auction won't make you better off; if you bid more than \$100, winning the auction will actually make you worse off. The only strategy that makes it possible for you to be better off is for you to bid less than \$100.
 4. A reasonable response might start off by noting that bidders will behave differently in the two auctions: bidders will shade their bids in a first-price auction, but not in a second-price auction. So in a first-price auction you get the highest bid from among a set of relatively low bids, and in a second-price auction you get the second-highest bid from among a set of relatively high bids. It's no longer clear which auction has the higher payoff. (In fact, there is a deeper result in game theory, called the Revenue Equivalence Theorem, which predicts that both types of auctions will yield the same expected payoff.)
 5. (a) There are two possible outcomes: either $\$x$ is the highest bid and you win the auction, or $\$x$ isn't the highest bid and you lose the auction.
 - (b) Your expected value from bidding $\$x$ in the auction is

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

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

something like

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

- (c) The expression above simplifies to

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

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

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

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

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

Chapter 13: Application: Marine Affairs

There are no problems in this chapter.

Chapter 14: Transition: Game Theory v. Price Theory

1. Ford, GM, Toyota, and a few other manufacturers dominate the market for new cars, so it is not a competitive market. In contrast, the market for used cars is pretty close to the competitive ideal. There are lots of small sellers—individuals looking to sell their cars—and lots of small buyers—individuals looking to buy a used car.

Appendix E

Answer Key for Part III: Many v. Many

Chapter 15: Supply and Demand Basics

1. This is the price at which the amount that buyers want to buy equals the amount that sellers want to sell. At a higher price, sellers want to sell more than buyers want to buy, creating incentives that push prices down toward the equilibrium. At a lower price, buyers want to buy more than sellers want to sell, creating incentives that push prices up toward the equilibrium.
2. (a) Demand increases. Equilibrium price up, quantity up.
(b) Supply decreases. Equilibrium price up, quantity down.
(c) Demand decreases. Equilibrium price down, quantity down.
(d) Demand decreases. Equilibrium price down, quantity down.
3. (a) Demand increases. Equilibrium price up, quantity up.
(b) Supply decreases. Equilibrium price up, quantity down.
4. (a) Demand increases. Equilibrium price up, quantity up.
(b) Supply decreases. Equilibrium price up, quantity down.
5. (a) Supply increases. Equilibrium price down, quantity up.
(b) Demand decreases. Equilibrium price down, quantity down.
(c) Demand increases. Equilibrium price up, quantity up.
(d) Demand increases. Equilibrium price up, quantity up.
(e) Supply increases. Equilibrium price down, quantity up.
6. (a) Supply decreases. Equilibrium price up, quantity down.

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

Chapter 16: Taxes


1. This is the logic identified in the text, e.g., "At a price of x with a \$.40 tax, buyers should be willing to buy exactly as much as they were willing to buy at a price of $\$(x + .40)$ without the tax." For per-unit taxes, you can also use the ideas of marginal cost and marginal benefit: a tax on the sellers increases marginal costs by the amount of the tax, and a tax on the buyers reduces marginal benefits by the amount of the tax. (Applying this marginal approach is a bit tricky for ad valorem taxes such as sales taxes. You need an additional assumption here about firm profits in equilibrium....)
2. (a) The equilibrium price is \$.80 per pound; the equilibrium quantity is 8 million pounds per day.
- (b) To find the slope of the supply curve, pick any two points on it—say, (8, .80) and (12, 1.00). Then the slope of the supply curve is

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

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


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

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




- The supply curve shifts up by \$.60.
- The new equilibrium features a price of \$1.20 per pound and a quantity of 4 million pounds per day.
- A tax of \$.60 per pound levied on 4 million pounds per day yields revenue of $(.60)(4) = \$2.4$ million per day.
- The buyers pay \$1.20 for each pound of oranges.
- Before paying the \$.60 tax the sellers receive \$1.20 per pound, so after paying the tax the sellers receive \$.60 per pound.
- Without the tax, buyers paid \$.80 per pound; they now pay \$1.20 per pound, so they are worse off by $T_B = .40$ per pound. Similarly, the sellers received \$.80 per pound without the tax, but now they only receive \$.60 per pound, so they are worse off by $T_S = .20$ per pound. (As a check here, note that the \$.40 per pound impact on the buyers plus the \$.20 per pound impact on the sellers equals the \$.60 per pound tax.) The ratio of the tax burdens is $\frac{T_B}{T_S} = \frac{.40}{.20} = 2$.
- The tax burden ratio is the same magnitude as the ratio of the slopes calculated previously! Intuitively, this is because the ratio of the slopes measures the relative responsiveness of buyers and sellers to price changes. The side that is most responsive to price changes (in


this case, the sellers) can push the lion's share of the tax burden onto the other side.



- (a) The demand curve shifts down by \$.60.
 - (b) The new equilibrium features a price of \$.60 per pound and a quantity of 4 million pounds per day.
 - (c) A tax of \$.60 per pound levied on 4 million pounds per day yields revenue of $(\$0.60)(4) = \2.4 million per day.
 - (d) The buyers pay \$.60 to the sellers for each pound of oranges, plus the tax of \$.60 to the government, so they pay a total of \$1.20 per pound.
 - (e) The sellers receive \$.60 per pound of oranges.
 - (f) Without the tax, buyers paid \$.80 per pound; they now pay \$1.20 per pound, so they are worse off by $T_B = \$0.40$ per pound. Similarly, the sellers received \$.80 per pound without the tax, but now they only receive \$.60 per pound, so they are worse off by $T_S = \$0.20$ per pound. (As a check here, note that the \$.40 per pound impact on the buyers plus the \$.20 per pound impact on the sellers equals the \$.60 per pound tax.) The ratio of the tax burdens is $\frac{T_B}{T_S} = \frac{.40}{.20} = 2$.
 - (g) The tax burden ratio is the same magnitude as the ratio of the slopes calculated previously!
5. The answers are essentially identical—regardless of the *legal* incidence of the tax (i.e., whether it's levied on the buyers or the sellers), the *economic* incidence of the tax comes out the same, i.e., the buyers always end up bearing \$.40 of the tax burden and the sellers end up bearing \$.20 of the tax burden. This is an example of the **tax incidence** result.



The supply curve rotates as shown above. The remaining answers are identical to those above; note that at the equilibrium price of \$1.20 per pound, the 50% tax on the sellers amounts to \$.60 per pound.



The demand curve rotates as shown above. The remaining answers are identical to those above; note that at the equilibrium price of \$.60 per pound, the 100% tax on the buyers amounts to \$.60 per pound.

8. The answers are the same! Again, this is an example of tax equivalence, but with sales taxes the comparison is a little bit less obvious than with per-unit taxes: in the case of sales taxes, a 50% tax on the sellers is equivalent to a 100% tax on the buyers. This is because the market price is the price the buyer pays the seller; since the market price is twice as

high when the tax is on the seller (\$1.20 versus \$.60), the sales tax rate on the sellers needs to be only half as high as the sales tax rate on the buyers in order to yield an equivalent result.

Algebra Problems

- A-1. (a) Simultaneously solving the demand and supply equations gives us a market equilibrium of $p = \$.80$ per pound and $q = 8$ million pounds per day. Total revenue is therefore \$6.4 million per day.
- (b) The slope of the demand curve is -10 , so the price elasticity of demand at the market equilibrium is $-10 \frac{.8}{8} = -1$. The slope of the supply curve is 20 , so the price elasticity of supply at the market equilibrium is $20 \frac{.8}{8} = 2$.
- (c) A tax of \$.60 on the buyers will change the demand curve to $q = 16 - 10(p + .6)$, i.e., $q = 10 - 10p$. The supply curve is still $q = -8 + 20p$, so the new market equilibrium is at $p = \$.60$ per pound and $q = 4$ million pounds per day. The buyers end up paying $.60 + .60 = \$1.20$ per pound, and the sellers get \$.60 per pound. The original equilibrium was at a price of \$.80 per pound, so the buyers end up paying \$.40 more and the sellers end up getting \$.20 less. The ratio of these tax burdens is $\frac{.40}{.20} = 2$. This is the negative inverse of the ratio of the elasticities.
- (d) A tax of \$.60 on the sellers will change the supply curve to $q = -8 + 20(p - .6)$, i.e., $q = -20 + 20p$. The demand curve is still $q = 16 - 10p$, so the new market equilibrium is at $p = \$1.20$ per pound and $q = 4$ million pounds per day. The buyers end up paying \$1.20 per pound, and the sellers get $1.20 - .60 = \$.60$ per pound. This is the same result as above, demonstrating the tax equivalence result.
- (e) A tax of 50% on the sellers will change the supply curve to $q = -8 + 20(.5p)$, i.e., $q = -8 + 10p$. The demand curve is still $q = 16 - 10p$, so the new market equilibrium is at $p = \$1.20$ per pound and $q = 4$ million pounds per day. The buyers end up paying \$1.20 per pound, and the sellers get $.5(1.20) = \$.60$ per pound.
- (f) A tax of 100% on the buyers will change the demand curve to $q = 16 - 10(2p)$, i.e., $q = 16 - 20p$. The supply curve is still $q = -8 + 20p$, so the new market equilibrium is at $p = \$.60$ per pound and $q = 4$ million pounds per day. The buyers end up paying $2(.60) = \$1.20$ per pound, and the sellers get \$.60 per pound. This is the sales tax version of the tax equivalence result.
- A-2. (a) Solving the two equations simultaneous we have $500 - 20p = 50 + 25p$,

which simplifies to $45p = 450$ or $p = 10$. Plugging this back in to either of the two original equations yields $q = 300$.

- (b) The supply equation is unchanged and the demand equation becomes $q = 500 - 20(p + 2)$, i.e., $q = 440 - 20p$. Solve this equation and the supply equation simultaneously to get the new equilibrium price and quantity.
 - (c) The demand equation is unchanged and the supply equation becomes $q = 50 + 25(p - 2)$, i.e., $q = 25p$. Solve this equation and the demand equation simultaneously to get the new equilibrium price and quantity.
 - (d) The supply equation is unchanged and the demand equation becomes $q = 500 - 20(1.2p)$, i.e., $q = 500 - 24p$. Solve this equation and the supply equation simultaneously to get the new equilibrium price and quantity.
 - (e) The demand equation is unchanged and the supply equation becomes $q = 50 + 25(.8p)$, i.e., $q = 50 + 20p$. Solve this equation and the demand equation simultaneously to get the new equilibrium price and quantity.
- A-3. (a) Solving simultaneously we get $50 - 6p = 20p - 28$, which yields $p = 3$. Plugging this into either the market demand curve or the market supply curve yields $q = 32$.
- (b) The supply curve is unchanged. The demand curve becomes

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

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

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

Chapter 17: Elasticities

1. N/A
2. Long-run demand and supply curves are more elastic, i.e., more responsive to price changes; in graphical terms this means that they are flatter. The reason is that buyers and sellers are more limited in short run. For example, in the short run buyers of gasoline are unlikely to respond to a price increase by buying a different car or moving where they live and/or where they work; but they can do these sorts of things in the long run. Similarly, in the short run sellers don't have enough time to build extra factories, or they might be stuck with too many factories; in the long run they can build extra factories or close down existing factories when their useful lifetime expires.

3. (a) The point Y corresponds to point A in the elasticity formula, so we have $p_A = \$.80$ and $q_A = 8$. For point B we can take any other point, e.g., the convenient point with $p_B = \$.60$ and $q_B = 10$. Then

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

- (b) Point X is the equilibrium during bad years, when frost reduces supply.
- (c) Total revenue at the three points are pq , i.e., $(1.2)(4) = \$4.8$ million per day at point X, $(.8)(8) = \$6.4$ million per day at point Y, and $(.2)(14) = \$2.8$ million per day at point Z.
- (d) Growers make higher profits during “bad” years: their revenue is higher and their costs are assumed to be identical. This is basically a Prisoner’s Dilemma situation for the growers: they would all be better off if they could restrict supply during “good” years, but the individual incentives lead them to flood the market and get low prices and low profits.
- (e) The point Y corresponds to point A in the elasticity formula, so we have $p_A = \$.80$ and $q_A = 8$. For point B we can take any other point on the supply curve, e.g., the convenient point with $p_B = \$.60$ and $q_B = 4$. Then

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

So the ratio of the elasticities is $\frac{\varepsilon_S}{\varepsilon_D} = \frac{2}{-1} = -2$. This is the same as the ratio of the slopes calculated previously! (This result follows from problem 4.)

4.

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

5. (a) At some market price p , firms making widgets earn the market rate of return; in the long run, then, firms are indifferent between making widgets and making other things, so they are willing to produce any number of widgets at price p . At any price less than p , firms would earn less than the market rate of return; in the long run, then, no firms would be willing to produce widgets, meaning that quantity supplied would be zero at any price less than p . Similarly, at any price greater than p , firms would earn more than the market rate of return;

in the long run, then, everybody would rush into the widget-making business, meaning that the quantity supplied would be infinite at any price greater than p .

- (b) A perfectly elastic supply curve for widgets is a horizontal line at some price p .
- (c) A tax on the seller would shift the supply curve up by the amount of the tax. Since the supply curve is horizontal, the equilibrium price would increase by the full amount of the tax, meaning that buyers would pay the entire tax burden. (Similarly, a tax on the buyer would shift the demand curve down, but the equilibrium price would not change, meaning that the buyers bear the full burden of the tax.) This makes sense because of the analysis above: if sellers bear part of the tax burden then they would be earning less than the market rate of return. So in the long run buyers bear the entire burden of taxes in a competitive market.

Chapter 18: Supply and Demand: Some Details

1. The market demand curve is

$$q = 500(20 - 2p) + 500(10 - 5p) = 15000 - 3500p.$$

The market supply curve is

$$q = 500(6p - 10) = 3000p - 5000.$$

2. Let the relevant variables be $p^H > p^L$ and $q^H > q^L$. A downward-sloping supply curve means q^L is optimal at the higher price (so that $p^H q^L - C(q^L)$ maximizes profits at price p^H) but that q^H is optimal at the lower price (so that $p^L q^H - C(q^H)$ maximizes profits at price p^L). To proceed by contradiction, note that profit maximization at the lower market price yields

$$p^L q^H - C(q^H) \geq p^L q^L - C(q^L).$$

It follows that q^L is not profit-maximizing at the higher price:

$$\begin{aligned} p^H q^H - C(q^H) &\geq (p^H - p^L)q^H + p^L q^L - C(q^L) \\ &= (p^H - p^L)(q^H - q^L) + p^H q^L - C(q^L) \\ &> p^H q^L - C(q^L). \end{aligned}$$

3. (a) The market demand curve is $q = 300(25 - 2p)$, i.e., $q = 7500 - 600p$. The market supply curve is $q = 500(5 + 3p)$, i.e., $q = 2500 + 1500p$.
- (b) Solving simultaneously we get $p = \frac{50}{21} \approx 2.38$ and $q \approx 6071$. Total revenue is therefore $pq \approx (2.38)(6071) \approx 14450$.

- (c) The price elasticity of demand at the market equilibrium is given by

$$\varepsilon_D = \frac{dq}{dp} \frac{p}{q} = -600 \frac{2.38}{6071} \approx -.235.$$

The price elasticity of supply at the market equilibrium is given by

$$\varepsilon_S = \frac{dq}{dp} \frac{p}{q} = 1500 \frac{2.38}{6071} \approx .588.$$

The ratio of the elasticities is $\frac{-600}{1500} = -.4$.

- (d) With a \$1 per-unit tax on the buyers, the market demand curve becomes $q = 7500 - 600(p + 1)$, i.e., $q = 6900 - 600p$. The market supply curve is still $q = 2500 + 1500p$, so the new equilibrium is at $p = \frac{44}{21} \approx 2.10$ and $q \approx 5643$. The seller therefore receives \$2.10 for each unit, and the buyer pays a total of $2.10 + 1.00 = \$3.10$ for each unit. Compared with the original equilibrium price of \$2.38, the seller is worse off by $2.38 - 2.10 = .28$, and the buyer is worse off by $3.10 - 2.38 = .72$. The ratio of these two is $\frac{.72}{.28} \approx 2.5$. Since $(.4)^{-1} = 2.5$, the tax burden ratio is the negative inverse of the ratio of the elasticities.
- (e) With a \$1 per-unit tax on the sellers, the market supply curve becomes $q = 2500 + 1500(p - 1)$, i.e., $q = 1000 + 1500p$. The market demand curve is, as originally, $q = 7500 - 600p$. So the new equilibrium is at $p = \frac{65}{21} \approx \3.10 and $q \approx 5643$. This shows the tax equivalence result: the buyers and sellers end up in the same spot regardless of whether the tax is placed on the buyer or the seller.
- (f) With a 50% sales tax on the sellers, the market supply curve becomes $q = 2500 + 1500(.5p)$, i.e., $q = 2500 + 750p$. The demand curve is, as originally, $q = 7500 - 600p$. So the new equilibrium is at $p = \frac{500}{135} \approx \3.70 and $q \approx 5278$.
- (g) With a 100% sales tax on the buyers, the demand curve becomes $q = 7500 - 600(2p)$, i.e., $q = 7500 - 1200p$. The supply curve is, as originally, $q = 2500 + 1500p$. So the new equilibrium is at $p = \frac{50}{27} \approx \1.85 and $q \approx 5278$. This shows the tax equivalence result for sales taxes: the buyers and sellers end up in the same spot regardless of whether there's a 50% sales tax on the sellers or a 100% sales tax on the buyers.

Chapter 19: Margins

There are no problems in this chapter.

Chapter 20: *Math*: Deriving Supply and Demand Curves

C-1. Microeconomics is about the actions and interactions of optimizing agents (e.g., profit-maximizing firms, utility-maximizing consumers). For differentiable functions with interior maxima or minima, the way to find those interior maxima or minima is to take a derivative and set it equal to zero. This gives you *candidate values* for maxima or minima; the reason is that slopes (i.e., derivatives) are equal to zero at the top of a hill (a maximum) or at the bottom of a valley (a minimum).

C-2. These are explained to the best of my abilities in the text.

- C-3. (a) The equation of the indifference curve is $L^2K = 9000$.
- (b) We can rewrite the indifference curve as $K = 9000L^{-2}$, which has a slope of $\frac{dK}{dL} = -18000L^{-3}$.
- (c) The marginal rate of substitute tells you how much cake the individual is willing to trade for one more latte. If $MRS = -3$, the firm should be willing to trade up to 3 pieces of cake to gain one latte; such a trade would leave the individual on the same indifference curve, i.e., allow the individual to reach the same utility level.
- (d) The individual wants to choose L and K to minimize $C(L, K) = 2L + 3K$ subject to the utility constraint $L^2K = 9000$. The choice variables are L and K ; the objective function is the cost function $C(L, K)$.
- (e) Two equations in two unknowns usually yield a unique solution, so to solve this problem we will find two relevant equations involving L and K and solve them simultaneously. The first equation is the constraint, $U(L, K) = 9000$. The second equation comes from the last dollar rule: $\frac{\frac{\partial U}{\partial L}}{p_L} = \frac{\frac{\partial U}{\partial K}}{p_K}$.
- (f) The marginal utility of lattes is $\frac{\partial U}{\partial L} = 2LK$ and the marginal utility of cake is $\frac{\partial U}{\partial K} = L^2$, so the last dollar rule gives us

$$\frac{\frac{\partial U}{\partial L}}{p_L} = \frac{\frac{\partial U}{\partial K}}{p_K} \implies \frac{2LK}{2} = \frac{L^2}{3},$$

which simplifies as

$$3LK = L^2 \implies 3K = L.$$

Substituting this into our first equation, the constraint $L^2K = 9000$, yields $(3K)^2K = 9000$, which simplifies to $9K^3 = 9000$, and then to $K^3 = 1000$, and finally to $K = 10$. It follows from either of our equations that $L = 30$, so the minimum cost required to reach a utility level of 9000 is $p_LL + p_KK = 2(30) + 3(10) = 90$.

- C-4. (a) The isoquant is $f(L, K) = q$, i.e., $L^{\frac{1}{4}}K^{\frac{1}{2}} = q$.
- (b) Squaring both sides yields $L^{\frac{1}{2}}K = q^2$, i.e., $K = q^2L^{-\frac{1}{2}}$. The slope of this is $\frac{dK}{dL} = -\frac{1}{2}q^2L^{-\frac{3}{2}}$.
- (c) The marginal rate of technical substitute tells you how much capital the firm is willing to trade for one more unit of labor. If $MRTS = -3$, the firm should be willing to trade up to 3 units of capital to gain one unit of labor; such a trade would leave the firm on the same isoquant, i.e., allow the firm to produce the same level of output.
- (d) The firm wants to choose L and K to minimize $C(L, K) = 2L + 2K$ subject to $f(L, K) = q$. The choice variables are L and K ; the objective function is the cost function $C(L, K)$.
- (e) Two equations in two unknowns usually yield a unique solution, so to solve this problem we will find two relevant equations involving L and K and solve them simultaneously. The first equation is the constraint, $f(L, K) = q$. The second equation comes from the last dollar rule: $\frac{\frac{\partial f}{\partial L}}{p_L} = \frac{\frac{\partial f}{\partial K}}{p_K}$.
- (f) The marginal product of labor is $\frac{\partial f}{\partial L} = \frac{1}{4}L^{-\frac{3}{4}}K^{\frac{1}{2}}$ and the marginal product of capital is $\frac{\partial f}{\partial K} = \frac{1}{2}L^{\frac{1}{4}}K^{-\frac{1}{2}}$, so the last dollar rule gives us

$$\frac{\frac{\partial f}{\partial L}}{p_L} = \frac{\frac{\partial f}{\partial K}}{p_K} \implies \frac{\frac{1}{4}L^{-\frac{3}{4}}K^{\frac{1}{2}}}{2} = \frac{\frac{1}{2}L^{\frac{1}{4}}K^{-\frac{1}{2}}}{2},$$

which simplifies as

$$\frac{1}{4}L^{-\frac{3}{4}}K^{\frac{1}{2}} = \frac{1}{2}L^{\frac{1}{4}}K^{-\frac{1}{2}} \implies K = 2L.$$

Substituting this into our first equation, the constraint $L^{\frac{1}{4}}K^{\frac{1}{2}} = q$, yields $L^{\frac{1}{4}}(2L)^{\frac{1}{2}} = q$, which simplifies to $L^{\frac{3}{4}}2^{\frac{1}{2}} = q$, and then to $L^{\frac{3}{4}} = q \cdot 2^{-\frac{1}{2}}$, and finally to $L = q^{\frac{4}{3}}2^{-\frac{2}{3}}$. Substituting this value of L into either of our two equations yields $K = q^{\frac{4}{3}}2^{\frac{1}{3}}$. So the minimum cost to produce 10 units of output is $p_LL + p_KK = 2(q^{\frac{4}{3}}2^{-\frac{2}{3}}) + 2(q^{\frac{4}{3}}2^{\frac{1}{3}}) = q^{\frac{4}{3}}(2^{\frac{1}{3}} + 2^{\frac{4}{3}}) \approx 3.78q^{\frac{4}{3}}$.

- (g) Plugging in $q = 10$ yields $L \approx 13.57$, $K \approx 27.14$, and a minimum cost of $p_LL + p_KK \approx 2(13.57) + 2(27.14) = 81.42$.
- C-5. (a) The firm's job is to choose q to maximize profits $\pi = pq - C(q) = pq - 3.78q^{\frac{4}{3}}$.
- (b) Take a derivative with respect to the choice variable (q) and set it equal to zero. Cost minimization is a necessary condition for profit maximization because a firm that is producing q units of output in a non-cost-minimizing way can always increase profits by producing q units of output in the cost-minimizing way. So a profit-maximizing firm must be producing its optimal level of output at least cost.

- (c) Taking a derivative of the profit function with respect to q and setting it equal to zero yields

$$\frac{d\pi}{dq} = 0 \implies p - \frac{4}{3}(3.78)q^{\frac{1}{3}} = 0 \implies p \approx 5.04q^{\frac{1}{3}}$$

Cubing both sides yields $p^3 \approx 128q$, i.e., $q = \frac{p^3}{128}$. This is the firm's supply curve.

- (d) Plugging $p = 16$ into the supply curve yields $q \approx 32$. So its profits are $\pi = pq - C(q) \approx 16(32) - 3.78(32)^{\frac{4}{3}} \approx 127.98$.
- C-6. (a) The isoquant is $L^{\frac{1}{4}}(4)^{\frac{1}{2}} = q$, i.e., $L^{\frac{1}{4}} = \frac{q}{2}$, i.e., $L = \frac{q^4}{16}$. Note that the isoquant is not a line but *just a single point*. This is because capital is fixed at $K = 4$, so the firm has no ability to trade-off between capital and labor.
- (b) The firm wants to choose L to minimize $C(L, K) = 2L + 2K = 2L + 8$ subject to $f(L, K) = q$.
- (c) Since the amount of capital the firm has is fixed, the firm cannot substitute between labor and capital. So the marginal rate of technical substitution is irrelevant in this problem.
- (d) Capital is a sunk cost, so the price of capital will not affect the firm's behavior.
- (e) In order to produce output of q , the firm has to hire $L = \frac{q^4}{16}$ units of labor. So the cost of producing q units of output is $C(q) = 2(\frac{q^4}{16}) + 2(4) = \frac{1}{8}q^4 + 8$.
- (f) The firm wants to choose q to maximize profits $\pi = pq - C(q) = pq - (\frac{1}{8}q^4 + 8)$. To solve this problem we take a derivative with respect to q and set it equal to zero, yielding

$$\frac{d\pi}{dq} = 0 \implies p - \frac{1}{2}q^3 = 0 \implies 2p = q^3 \implies q = (2p)^{\frac{1}{3}}.$$

- C-7. (a) The profit function is $\pi = pq - C(q) = pq - 2q\sqrt{2}$. Taking a derivative and setting it equal to zero we get $p - 2\sqrt{2} = 0$, i.e., $p = 2\sqrt{2}$.
- (b) At a price of $p = 2$, the firm will supply $q = 0$: its cost of producing each unit of output is $2\sqrt{2} > 2$, so it loses money on each unit it sells!
- (c) At a price of $p = 4$, the firm will supply infinitely many units of output.
- (d) We have $f(2L, 2K) = (2L)^{\frac{1}{2}}(2K)^{\frac{1}{2}} = 2L^{\frac{1}{2}}K^{\frac{1}{2}} = 2f(K, L)$.
- (e) Since doubling inputs doubles output, the firm can double and redouble its profits simply by doubling and redoubling production (i.e.,

its choice of inputs). This (hopefully) helps explain why we get a corner solution (of $q = \infty$) when we attempt to maximize profits with $p = 4$.

- C-8. (a) The individual wants to choose L and K to maximize utility $U(L, K) = L^2K$ subject to the budget constraint $p_L L + p_K K = M$.
- (b) The solution method is to find two equations involving L and K and then solve them simultaneously. One equation is the constraint, $p_L L + p_K K = M$; the other is the last dollar rule, $\frac{\partial U}{\partial L} = \frac{\partial U}{\partial K}$.
- (c) The last dollar rule gives us $\frac{2LK}{p_L} = \frac{L^2}{p_K}$, which simplifies to $3K = L$ when $p_L = 2$ and $p_K = 3$. Substituting this into the budget constraint we have $2(3K) + 3K = 90$, i.e., $K = 10$. Substituting back into either of our equations yields $L = 30$.
- (d) The last dollar rule gives us $\frac{2LK}{p_L} = \frac{L^2}{p_K}$, which simplifies to $6K = p_L L$ when $p_K = 3$. Solving for K and substituting this into the budget constraint yields $p_L L + 3(\frac{1}{6}p_L L) = 90$, i.e., $1.5p_L L = 90$. This simplifies to $L = \frac{60}{p_L}$, which is the Marshallian demand curve for lattes when $p_K = 3$ and $M = 90$. The slope of this demand curve is $\frac{dL}{dp_L} = -60p_L^{-2}$; when $p_L = 2$, this simplifies to $\frac{-60}{4} = -15$.
- (e) The last dollar rule gives us $\frac{2LK}{p_L} = \frac{L^2}{p_K}$, which simplifies to $p_K K = L$ when $p_L = 2$. Using this to substitute for L in the budget constraint yields $2(p_K K) + p_K K = 90$, i.e., $K = \frac{30}{p_K}$, which is the Marshallian demand curve for cake when $p_L = 2$ and $M = 90$. The slope of this demand curve is $\frac{dK}{dp_K} = -30(p_K)^{-2}$; when $p_K = 3$, this simplifies to $-\frac{30}{9} \approx -3.33$.
- (f) The last dollar rule gives us $\frac{2LK}{p_L} = \frac{L^2}{p_K}$, which simplifies to $3K = L$ when $p_L = 2$ and $p_K = 3$. Using this to substitute for L in the budget constraint yields $2(3K) + 3K = M$, i.e., $K = \frac{M}{9}$, which is the Engel curve for cake when $p_L = 2$ and $p_K = 3$. Using the last dollar rule result $3K = L$ to substitute for K in the budget constraint yields $2L + L = M$, i.e., $L = \frac{M}{3}$, which is the Engel curve for cake when $p_L = 2$ and $p_K = 3$.
- C-9. (a) Choose L and K to minimize $C(L, K) = p_L L + p_K K$ subject to the constraint $L^2K = 9000$.
- (b) Combine the constraint with the last dollar rule to get two equations in two unknowns. Solve these simultaneously to get the optimal values of L and K .
- (c) We previously calculated the last dollar rule to yield $3K = L$ when $p_L = 2$ and $p_K = 3$. Substituting this into the constraint yields $(3K)^2K = 9000$, which simplifies to $9K^3 = 9000$, i.e., $K = 10$. It follows from either equation that $L = 30$. The cost-minimizing cost is therefore $p_L L + p_K K = 2(30) + 3(10) = 90$.

- (d) The last dollar rule gives us $\frac{2LK}{p_L} = \frac{L^2}{p_K}$, which simplifies to $6K = p_L L$ when $p_K = 3$. Using this to substitute for K in the constraint $L^2 K = 9000$ yields $L^2 \frac{p_L L}{6} = 9000$, which simplifies to $L^3 = 54000 p_L^{-1}$, and then to $L = (54000 p_L^{-1})^{\frac{1}{3}} = (30)2^{\frac{1}{3}} p_L^{-\frac{1}{3}}$. The slope of this demand curve is $\frac{dL}{dp_L} = -\frac{1}{3}(30)2^{\frac{1}{3}} p_L^{-\frac{4}{3}}$. When $p_L = 2$ this simplifies to $(-10)2^{-1} = -5$.
- (e) The last dollar rule gives us $\frac{2LK}{p_L} = \frac{L^2}{p_K}$, which simplifies to $p_K K = L$ when $p_L = 2$. Using this to substitute for L in the constraint $L^2 K = 9000$ yields $(p_K K)^2 K = 9000$, which simplifies to $p_K^2 K^3 = 9000$, and then to $K = (9000 p_K^{-2})^{\frac{1}{3}} = (10)3^{\frac{2}{3}} p_K^{-\frac{2}{3}}$. The slope of this demand curve is $\frac{dK}{dp_K} = -\frac{2}{3}(10)3^{\frac{2}{3}} p_K^{-\frac{5}{3}}$. When $p_K = 3$ this simplifies to $-\frac{2}{3}(10)3^{-1} = -\frac{20}{9} \approx -2.22$.

C-10. The two problems are two sides of the same coin: if $U = 9000$ is the maximum utility that can be achieved with a budget of $M = 90$, then $M = 90$ is the minimum budget required to reach a utility level of $U = 9000$.

- C-11. (a) What is the maximum amount of money this individual could exchange for 10 more lattes *and still be on the same indifference curve*, i.e., still have a utility level of 9000?
- (b) We must have $L^2 K = 9000$. So (1) if $L = 30$ then we need $K = 10$, which at a price of \$3 per cake requires a budget of \$30; (2) if $L = 40$ then we need $K = 5.625$, which at a price of \$3 per cake requires a budget of \$16.875; (3) subtracting yields \$30-\$16.875 = \$13.125 as this individual's willingness-to-pay for those 10 lattes.
- (c) $p_L = 54000 L^{-3}$
- (d) See figure E.1.
- (e) The area is the same as the willingness-to-pay calculated above!

$$\int_{30}^{40} 54000 L^{-3} dL = -27000 L^{-2} \Big|_{30}^{40} = -16.875 + 30 = \$13.125$$

- (f) Various texts (e.g., Silberberg's *Structure of Economics*) insist that the area under Marshallian demand curves is meaningless.

Chapter 21: Transition: Welfare Economics

There are no problems in this chapter.




Figure E.1: The area under the *Hicksian* (utility-held-constant) demand curve measures willingness to pay.

Appendix F

Answer Key for Appendixes

Appendix A: Government in Practice

There are no problems in this chapter.

Appendix B: *Math*: Monopoly and Oligopoly

1. Firm 1's profits are

$$\pi_1 = pq_1 - C(q_1) = (10 - q_1 - q_2)q_1 - 3q_1^2 = (10 - q_2)q_1 - 4q_1^2.$$

Firm 2's profits are

$$\pi_2 = pq_2 - C(q_2) = (10 - q_1 - q_2)q_2 - 3q_2^2 = (10 - q_1)q_2 - 4q_2^2.$$

With collusion, the firms choose q_1 and q_2 to maximize joint profits

$$\pi_1 + \pi_2 = (10 - q_2)q_1 - 4q_1^2 + (10 - q_1)q_2 - 4q_2^2.$$

To solve this problem, we take partial derivatives with respect to each choice variable and set them equal to zero. This will give us two necessary first-order conditions (NFOCs) in two unknowns (q_1 and q_2); solving these simultaneously gives us our optimum.

So: the NFOCs are

$$\frac{\partial(\pi_1 + \pi_2)}{\partial q_1} = 0 \implies 10 - q_2 - 8q_1 - q_2 = 0$$

and

$$\frac{\partial(\pi_1 + \pi_2)}{\partial q_2} = 0 \implies -q_1 + (10 - q_1) - 8q_2 = 0$$

Solving these jointly yields $q_1 = q_2 = 1$. The price is therefore $p = 10 - 2 = 8$ and industry profits are

$$\pi_1 + \pi_2 = 2(pq_1 - C(q_1)) = 2(8(1) - 3(1^2)) = 2(5) = 10.$$

Next, the Cournot problem. Here Firm 1 chooses q_1 to maximize its profits and Firm 2 chooses q_2 to maximize its profits. (The profit functions are given above.) To solve this problem we take a partial derivative of π_1 with respect to q_1 to get a necessary first-order condition (NFOC) for Firm 1. We then take a partial derivative of π_2 with respect to q_2 to get a necessary first-order condition (NFOC) for Firm 2. Solving these NFOCs simultaneously gives us the Cournot outcome.

So: the NFOCs are

$$\frac{\partial(\pi_1)}{\partial q_1} = 0 \implies 10 - q_2 - 8q_1 = 0$$

and

$$\frac{\partial(\pi_2)}{\partial q_2} = 0 \implies (10 - q_1) - 8q_2 = 0$$

Solving these jointly yields $q_1 = q_2 = \frac{10}{9}$. The price is therefore $p = 10 - \frac{20}{9} = \frac{70}{9}$.

Finally, the Stackleberg problem. The objective functions look the same as in the Cournot case, but here we use backward induction to solve. Firm 2 sees what Firm 1 has chosen, and so chooses q_2 to maximize its profits. It does this by taking a partial derivative of π_2 (given above) with respect to q_2 to get an NFOC that is Firm 2's best response function to Firm 1's choice of q_1 . Next, Firm 1 must anticipate Firm 2's reaction to its choice of q_1 , and substitute this reaction function $q_2(q_1)$ into its profit function. Taking a partial derivative of the resulting profit function π_1 with respect to q_1 yields an NFOC that identifies Firm 1's profit-maximizing choice of q_1 . Plugging this solution into Firm 2's best response function identifies Firm 2's profit-maximizing response of q_2 .

So: The NFOC for Firm 2 is exactly as above: $10 - q_1 - 8q_2 = 0$, i.e., $q_2 = \frac{10 - q_1}{8}$. Substituting this into Firm 1's profit function yields

$$\pi_1 = (10 - q_2)q_1 - 4q_1^2 = \left(10 - \frac{10 - q_1}{8}\right)q_1 - 4q_1^2 = \frac{1}{8}(70q_1 - 31q_1^2).$$

Taking a derivative with respect to q_1 gives us the NFOC

$$\frac{\partial\pi_1}{\partial q_1} = 0 \implies \frac{1}{8}(70 - 62q_1) = 0 \implies q_1 = \frac{35}{31} \approx 1.13.$$

Plugging this into Firm 2's best response function yields

$$q_2 = \frac{10 - q_1}{8} \approx \frac{10 - 1.13}{8} \approx 1.11.$$

2. Firm 1's profits are

$$\pi_1 = pq_1 - C(q_1) = (10 - q_1 - q_2)q_1 - 4q_1^2 = (10 - q_2)q_1 - 5q_1^2.$$

Firm 2's profits are, as above

$$\pi_2 = pq_2 - C(q_2) = (10 - q_1 - q_2)q_2 - 3q_2^2 = (10 - q_1)q_2 - 4q_2^2.$$

With collusion, the firms choose q_1 and q_2 to maximize joint profits

$$\pi_1 + \pi_2 = (10 - q_2)q_1 - 5q_1^2 + (10 - q_1)q_2 - 4q_2^2.$$

To solve this problem, we take partial derivatives with respect to each choice variable and set them equal to zero. This will give us two necessary first-order conditions (NFOCs) in two unknowns (q_1 and q_2); solving these simultaneously gives us our optimum.

So: the NFOCs are

$$\frac{\partial(\pi_1 + \pi_2)}{\partial q_1} = 0 \implies 10 - q_2 - 10q_1 - q_2 = 0$$

and

$$\frac{\partial(\pi_1 + \pi_2)}{\partial q_2} = 0 \implies -q_1 + (10 - q_1) - 8q_2 = 0$$

Solving these jointly yields $q_1 = \frac{15}{19} \approx .79$ and $q_2 = \frac{4}{3}q_1 \approx 1.05$. The price is therefore $p \approx 10 - .79 - 1.05 = 8.16$ and industry profits are

$$\begin{aligned} \pi_1 + \pi_2 &\approx [(8.16)(.79) - 4(.79)^2] + [(8.16)(1.05) - 3(1.05)^2] \\ &\approx 3.95 + 5.26 = 9.21 \end{aligned}$$

Next, the Cournot problem. Here Firm 1 chooses q_1 to maximize its profits and Firm 2 chooses q_2 to maximize its profits. (The profit functions are given above.) To solve this problem we take a partial derivative of π_1 with respect to q_1 to get a necessary first-order condition (NFOC) for Firm 1. We then take a partial derivative of π_2 with respect to q_2 to get a necessary first-order condition (NFOC) for Firm 2. Solving these NFOCs simultaneously gives us the Cournot outcome.

So: the NFOCs are

$$\frac{\partial(\pi_1)}{\partial q_1} = 0 \implies 10 - q_2 - 10q_1 = 0$$

and

$$\frac{\partial(\pi_2)}{\partial q_2} = 0 \implies (10 - q_1) - 8q_2 = 0$$

Solving these jointly yields $q_1 = \frac{70}{79} \approx .89$ and $q_2 = \frac{90}{79} \approx 1.14$. The price is therefore $p \approx 10 - .89 - 1.14 = 7.97$.

Finally, the Stackleberg problem. The objective functions look the same as in the Cournot case, but here we use backward induction to solve. Firm 2 sees what Firm 1 has chosen, and so chooses q_2 to maximize its profits. It does this by taking a partial derivative of π_2 (given above) with respect to q_2 to get an NFOC that is Firm 2's best response function to Firm 1's choice of q_1 . Next, Firm 1 must anticipate Firm 2's reaction to its choice of q_1 , and substitute this reaction function $q_2(q_1)$ into its profit function. Taking a partial derivative of the resulting profit function π_1 with respect to q_1 yields an NFOC that identifies Firm 1's profit-maximizing choice of q_1 . Plugging this solution into Firm 2's best response function identifies Firm 2's profit-maximizing response of q_2 .

So: The NFOC for Firm 2 is exactly as above: $10 - q_1 - 8q_2 = 0$, i.e., $q_2 = \frac{10-q_1}{8}$. Substituting this into Firm 1's profit function yields

$$\pi_1 = (10 - q_2)q_1 - 5q_1^2 = \left(10 - \frac{10 - q_1}{8}\right)q_1 - 5q_1^2 = \frac{1}{8}(70q_1 - 39q_1^2).$$

Taking a derivative with respect to q_1 gives us the NFOC

$$\frac{\partial \pi_1}{\partial q_1} = 0 \implies \frac{1}{8}(70 - 78q_1) = 0 \implies q_1 = \frac{70}{78} \approx .90.$$

Plugging this into Firm 2's best response function yields

$$q_2 = \frac{10 - q_1}{8} \approx \frac{10 - .90}{8} \approx 1.14.$$

3. The only difference between this problem and the previous problem is that Firm 1 and Firm 2 have switched places. So the solutions to the collusion and Cournot problems will be symmetric to the solutions above. Where the switch *is* important is in the Stackleberg game, because here the sequence of moves matters.

So: the Stackleberg problem. The objective functions look the same as in the Cournot case, but here we use backward induction to solve. Firm 2 sees what Firm 1 has chosen, and so chooses q_2 to maximize its profits. It does this by taking a partial derivative of π_2 (given above) with respect to q_2 to get an NFOC that is Firm 2's best response function to Firm 1's choice of q_1 . Next, Firm 1 must anticipate Firm 2's reaction to its choice of q_1 , and substitute this reaction function $q_2(q_1)$ into its profit function. Taking a partial derivative of the resulting profit function π_1 with respect to q_1 yields an NFOC that identifies Firm 1's profit-maximizing choice of q_1 . Plugging this solution into Firm 2's best response function identifies Firm 2's profit-maximizing response of q_2 .

So: The NFOC for Firm 2 is symmetric with Firm 1's best response function from the previous problem: $10 - q_1 - 10q_2 = 0$, i.e., $q_2 = \frac{10-q_1}{10}$.

Substituting this into Firm 1's profit function yields

$$\pi_1 = (10 - q_2)q_1 - 4q_1^2 = \left(10 - \frac{10 - q_1}{10}\right)q_1 - 4q_1^2 = \frac{1}{10}(90q_1 - 39q_1^2).$$

Taking a derivative with respect to q_1 gives us the NFOC

$$\frac{\partial \pi_1}{\partial q_1} = 0 \implies \frac{1}{10}(90 - 78q_1) = 0 \implies q_1 = \frac{90}{78} \approx 1.15.$$

Plugging this into Firm 2's best response function yields

$$q_2 = \frac{10 - q_1}{10} \approx \frac{10 - 1.15}{10} \approx .89.$$

4. Firm 1's profits are

$$\pi_1 = p_1 q_1 - C(q_1) = (10 - 2q_1 - q_2)q_1 - 3q_1^2 = (10 - q_2)q_1 - 5q_1^2.$$

Firm 2's profits are

$$\pi_2 = p_2 q_2 - C(q_2) = (10 - q_1 - 2q_2)q_2 - 3q_2^2 = (10 - q_1)q_2 - 5q_2^2.$$

With collusion, the firms choose q_1 and q_2 to maximize joint profits

$$\pi_1 + \pi_2 = (10 - q_2)q_1 - 5q_1^2 + (10 - q_1)q_2 - 5q_2^2.$$

To solve this problem, we take partial derivatives with respect to each choice variable and set them equal to zero. This will give us two necessary first-order conditions (NFOCs) in two unknowns (q_1 and q_2); solving these simultaneously gives us our optimum.

So: the NFOCs are

$$\frac{\partial(\pi_1 + \pi_2)}{\partial q_1} = 0 \implies 10 - q_2 - 10q_1 - q_2 = 0$$

and

$$\frac{\partial(\pi_1 + \pi_2)}{\partial q_2} = 0 \implies -q_1 + (10 - q_1) - 10q_2 = 0$$

Solving these jointly yields $q_1 = q_2 = \frac{10}{12} \approx .83$. The prices are therefore $p_1 = p_2 \approx 10 - 3(.83) \approx 7.51$ and industry profits are

$$\pi_1 + \pi_2 = 2(p_1 q_1 - C(q_1)) \approx 2(7.51(.83) - 3(.83^2)) \approx 4.17.$$

Next, the Cournot problem. Here Firm 1 chooses q_1 to maximize its profits and Firm 2 chooses q_2 to maximize its profits. (The profit functions are given above.) To solve this problem we take a partial derivative of π_1 with respect to q_1 to get a necessary first-order condition (NFOC) for

Firm 1. We then take a partial derivative of π_2 with respect to q_2 to get a necessary first-order condition (NFOC) for Firm 2. Solving these NFOCs simultaneously gives us the Cournot outcome.

So: the NFOCs are

$$\frac{\partial(\pi_1)}{\partial q_1} = 0 \implies 10 - q_2 - 10q_1 = 0$$

and

$$\frac{\partial(\pi_2)}{\partial q_2} = 0 \implies (10 - q_1) - 10q_2 = 0$$

Solving these jointly yields $q_1 = q_2 = \frac{10}{11} \approx .91$. The prices are therefore $p_1 = p_2 \approx 10 - 3(.91) = 7.27$.

Finally, the Stackleberg problem. The objective functions look the same as in the Cournot case, but here we use backward induction to solve. Firm 2 sees what Firm 1 has chosen, and so chooses q_2 to maximize its profits. It does this by taking a partial derivative of π_2 (given above) with respect to q_2 to get an NFOC that is Firm 2's best response function to Firm 1's choice of q_1 . Next, Firm 1 must anticipate Firm 2's reaction to its choice of q_1 , and substitute this reaction function $q_2(q_1)$ into its profit function. Taking a partial derivative of the resulting profit function π_1 with respect to q_1 yields an NFOC that identifies Firm 1's profit-maximizing choice of q_1 . Plugging this solution into Firm 2's best response function identifies Firm 2's profit-maximizing response of q_2 .

So: The NFOC for Firm 2 is exactly as above: $10 - q_1 - 10q_2 = 0$, i.e., $q_2 = \frac{10-q_1}{10}$. Substituting this into Firm 1's profit function yields

$$\pi_1 = (10 - q_2)q_1 - 5q_1^2 = \left(10 - \frac{10 - q_1}{10}\right)q_1 - 5q_1^2 = \frac{1}{10}(90q_1 - 49q_1^2).$$

Taking a derivative with respect to q_1 gives us the NFOC

$$\frac{\partial\pi_1}{\partial q_1} = 0 \implies \frac{1}{10}(90 - 98q_1) = 0 \implies q_1 = \frac{90}{98} \approx .92.$$

Plugging this into Firm 2's best response function yields

$$q_2 = \frac{10 - q_1}{10} \approx \frac{10 - .92}{10} \approx .91.$$

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 12 for descriptions of different kinds of auctions.

auction A method of selling an item that involves pitting potential buyers against each other. See Chapter 12 for descriptions of the different kinds of auctions.

backward induction A solution concept for **sequential move games** that involves reasoning backward from the end of the **game tree** to the beginning.

barrier to entry A legal or economic barrier that protects a monopoly by preventing other firms from entering the market.

capital theory The branch of microeconomics dealing with investment decisions.

collective action problem A game, such as the **prisoners' dilemma**, in which decisions that are individually optimal lead to results that are collectively sub-optimal (and, in particular, to results that are **Pareto inefficient**).

competitive market A market with many buyers, each small in relation to all the buyers together, and many sellers, each small in relation to all the sellers together. Also called (somewhat redundantly) a **perfectly competitive market**.

complement Intuitively, a good that is used in combination with another good: bread is a complement to butter. (Compare with **substitute**)

consumer surplus The gains from trade accruing to the buyer. (Compare with **producer surplus**)

decision tree A method of visualizing how individuals make decisions. The branches coming out of each node represent the range of possible choices for the individual at that point. (Compare with **game tree**)

demand curve A curve relating the **market price** to desired purchase amounts of some good. An *individual demand curve* relates how much of that good some individual wants to buy at any given price; a *market demand curve* relates how much of that good all of the buyers together want to buy at any given price.

descending price auction See Chapter 12 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 12 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 12 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 12 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 8.3.

price elasticity See **elasticity**.

price-taker An individual, for example a buyer or seller in a **competitive market**, who takes the market price as given because the scope of his activities are too small to affect big-picture variables.

price theory The branch of microeconomics dealing with market interactions between a large number of individuals. (Compare with **game theory**)

prisoners' dilemma In general, a prisoners' dilemma is a **simultaneous move game** with two or more players that features (1) symmetric strategies and payoffs for all the players; (2) **strictly dominant strategies** for all the players; and (3) a predicted outcome that is Pareto inefficient. (Such problems are also called **collective action problems**.) The phrase "prisoners' dilemma" may also refer to a specific game, one featuring two prisoners who must choose whether or not to accept a plea-bargain that involves testifying against the other prisoner.

producer surplus The gains from trade accruing to the seller. (Compare with **consumer surplus**)

profit Loosely defined, money in minus money out. Slightly more accurate is to account for the changing value of money over time by defining profit as the **present value** of money in minus money out. Other definitions consider this to be *accounting profit* rather than *economic profit*, which subtracts out *opportunity costs*.

quantum *Physics.* A minimum amount of a physical quantity which can exist and by multiples of which changes in the quantity occur (Oxford English Dictionary). The "quantum" of economics is the optimizing individual.

real interest rate The interest rate in terms of *purchasing power*, i.e., in terms of your ability to buy stuff. Real interest rates account for **inflation**. (Compare with **nominal interest rate**)

rent-seeking Behavior (usually by a monopolist) intended to increase profits.

repeated game A game in which a **stage game** is played multiple times.

Revenue Equivalence Theorem A theoretical result indicating that a wide variety of auctions all generate the same **expected** revenue. (Note that this only hold under certain conditions.)

risk A **risk-averse** individual prefers to avoid risks, a **risk-loving** individual prefers to take risks, and a **risk-neutral** individual is indifferent to risk.

risk premium An added payment made to avoid risk, or accepted to take on risk.

Robinson Crusoe model A simple economic model based on Daniel Defoe's 1719 novel about life on a desert island for a shipwrecked sailor.

sales tax A tax—usually an **ad valorem tax**—on the sale of some good.

second price auction See Chapter 12 for descriptions of different kinds of auctions.

Second Welfare Theorem See **welfare theorems**.

sequential move game A game, such as chess, in which players take turns moving. These games can be analyzed with **game trees**. (Compare with **simultaneous move game**)

shading one's bid An auction strategy involving lowering one's bid.

simultaneous move game A game, such as Rock, Paper, Scissors, in which players move simultaneously; one important example is the **Prisoners' Dilemma**. These games can be analyzed with **payoff matrices**. (Compare with **sequential move game**)

stage game A game, usually a simple one, that is played multiple times to yield a **repeated game**.

strictly dominant strategy A strategy for one player that yields a payoff that is strictly greater than her payoff from any other strategy, regardless of the other players' strategies. (Compare with **weakly dominant strategy**)

strictly dominated strategy A strategy for one player that yields a payoff that is strictly less than his payoff from some other strategy, regardless of the other players' strategies. (Compare with **weakly dominated strategy**)

subgame Generally used in reference to a **game tree**, a subgame is a subset of a game, e.g., one section of a game tree.

substitute Intuitively, a good that can be used to replace another good: margarine is a substitute for butter. (Compare with **complement**)

substitution effect Together with **income effect**, the components of how individuals respond to price changes. The substitution effect focuses on the effect of price changes on *relative prices*: since a price increase in some good makes it more expensive relative to similar goods, individuals will respond by *substituting out of* the now-more-expensive good and *substituting into* the related goods.

sunk cost A cost that will be borne regardless; for example, when deciding whether or not to sell some stock, the amount you spend on buying the stock is a sunk cost. As in this example, sunk costs often are costs incurred in the past.

supply curve A curve relating the **market price** to desired sales amounts of some good. An **individual supply curve** relates how much of that good some individual wants to sell at any given price; a **market supply curve** relates how much of that good all of the sellers together want to sell at any given price.

tax See **per-unit tax** or **ad valorem tax**.

trigger strategy A strategy in **repeated games** that involves cooperating as long as the other players also cooperate, but ending cooperation as soon as any of the other players stop cooperating.

unfair bet A bet with an expected value less than zero.

uniform pricing The practice of charging all customers the same price for the same good; compare with **price discrimination**.

unit elastic Indicating an elasticity of exactly $+1$ or -1 ; used to describe a degree of responsiveness, e.g., to price changes, that is neither relatively large nor relatively small. (Compare with **elastic** and **inelastic**)

weakly dominant strategy A strategy for one player that yields a payoff that is greater than *or equal to* his payoff from any other strategy, regardless of the other players' strategies. (Compare with **strictly dominant strategy**)

weakly dominated strategy A strategy for one player that yields a payoff that is less than *or equal to* her payoff from some other strategy, regardless of the other players' strategies. (Compare with **strictly dominated strategy**)

welfare economics The branch of microeconomics dealing with social welfare.

welfare theorems Two important results from **welfare economics**. The First Welfare Theorem says that complete and competitive markets yield a Pareto efficient outcome. The Second Welfare Theorem says that *any* **Pareto efficient** outcome can be reached via complete and competitive markets, provided that trade is preceded by an appropriate reallocation of resources.

winner's curse Used in reference to auctions, e.g., of oil-drilling rights, in which the winning bidder overbids. Note that this pertains only to *common-value* auctions.

Index

- Aalsmeer Flower Auction, 121
- absolute advantage, 77
- annuity, 31–33, 39
- arbitrage, 61–65
 - equity premium puzzle, 65
 - explorers and sheep, 62–64
 - financial, 63–65
 - no arbitrage result, 62
 - rush-hour, 62–63
- auction, 26, 70, 121–130
 - all-pay, 126
 - code bidding, 127
 - common values, 127
 - eBay, 125
 - equivalences, 124–126
 - FCC spectrum, 127
 - independent values, 127
 - multi-unit, 127
 - Revenue Equivalence Theorem, 126
 - shading one's bid, 123
 - types of, 122–123
 - winner's curse, 127
- backward induction, 89
- Bank
 - of America, 34, 35, 50–53, 65
 - of Fish, 34, 35, 50–53
 - of Gold, 35
 - of Microsoft, 35
 - of Oil, 35
- barriers to entry, 139
- Bazerman, M.H., 127
- benefit
 - total, 187
- best response function, 230
- Binmore, Ken, 69
- blue
 - and yellow makes green, 146
- Brams, Steven, 74, 95
- budget constraint, 37
 - and present value, 37
- business
 - game theory applied to, 70, 93
- cake cutting algorithms
 - applied to divorce and Middle East, 74
- cake-cutting algorithms, 69–78
 - Adjusted Winner, 74
 - books about, 73, 74
 - Divide and Choose, 71, 72
 - divide or choose?, 72
 - Moving Knife, 71, 72
 - objectives, 74
 - Successive Pairs, 72
 - three player envy-free, 73
 - with many kids, 72
- calculus, 7, 13, 126, 166, 169, 185, 188
- Camp David peace accord, 74
- capital
 - fish are, 34
 - theory, 34–37
- carbon tax, 163
- Cassidy, Butch, 96
- certainty equivalent, 23
- climate change, 41, 163
- Coase Theorem, 75
- Cobb-Douglas
 - production function, 192
 - utility function, 191
- collective action problem, 101
- comparative advantage, 78
- compensated demand curve, 205
- competitive markets, 137

- complement
 - defined using elasticity, 170
- complementary good, 148
- compound interest, 40
 - continuous, 38
- consumer surplus, 187
- continuous compounding, 38
- Conway, John, 73
- corner solution, 50, 113–116, 200, 231
- cost
 - fixed, 185, 186, 188
 - function, 195
 - minimization, 191
 - opportunity, 6
 - sunk, 5, 13, 14, 46
 - total variable, 185
- cost minimization
 - explained graphically, 197
 - explained intuitively, 195
 - explained mathematically, 198
 - for a firm, 195
 - for an individual, 194
 - last dollar rule, 195
- Crampton, Peter, 127
- Crowe, Russell, 110
- decision tree, 5, 8, 19, 29, 89–91, 173, 174
- demand, 143, 173
- demand curve
 - and income effect, 206
 - and substitution effect, 206
 - factor, 199
 - Hicksian (compensated), 204
 - in terms of margins, 186
 - income effect, 178
 - increases and decreases, 146
 - individual, 174
 - inverse, 225
 - market, 144, 174
 - Marshall versus Hicks, 205
 - Marshallian (uncompensated), 202
 - substitution effect, 178
- derivative, 7
 - partial, 9
- differentiation, *see* derivative, calculus
- diversification, 21
- divorce settlement, 74
- Dixit, Avinash K., 94
- dominance
 - iterated, 107
 - strict, 106
 - weak, 108
- duopoly, 99, 228
 - collusion, 228, 229
 - Cournot, 228, 230
 - Stackleberg, 228, 232
- eBay, 121, 125
- economic models
 - Robinson Crusoe, 76
- economics
 - basic assumption of, 3
 - experimental, 102
- elasticity, 163–172
 - X elasticity of Y , 167
 - and calculus, 169
 - and monopoly behavior, 166
 - and taxes, 169
 - and total revenue, 166
 - cross-price, 170
 - elastic, 164
 - formula, 165
 - inelastic demand, 164
 - price elasticity of demand, 164
 - price elasticity of supply, 167
 - short-run versus long-run, 170
 - summed up
 - in English, 165
 - in math, 165
 - in pictures, 165
 - unit, 164
 - versus slope, 163
- Engel curve, 201
- equilibrium
 - market, 145
 - Nash, 108
- equity premium puzzle, 65
- evolution
 - game theory applied to, 70
- expected value, 19–27, 41, 62, 63, 112, 113, 129

- and indifference curve, 23–24
- experimental economics, 102
- factor demand curve, 199
- fair bet, 20, 25
- federal government, *see* government
- Federal Reserve Board, 220
- first mover advantage, 233
- First Welfare Theorem, 213
- fish
 - are capital, 34
 - economic management of, 48–53
 - bang-bang solution, 50
 - Maximum Sustainable Yield, 34, 35, 51
 - open-access fishery, 35
 - slow-growing, 35
- free-rider, 101
- Friedman, David D., 62
- future value, 30, 37, 39
- game
 - subgame, 89
 - theory, 69
- game theory
 - applications of, 70–71
- game tree, 89
- games
 - Battle of the Sexes, 111
 - centipede, 97
 - centipede game, 103
 - draft, 95
 - duopoly, 99
 - finitely repeated, 102
 - grenade, 89–91
 - ice cream pie, 94
 - infinitely repeated, 110
 - investment, 104
 - matching pennies, 112
 - pedestrian in Rome, 97
 - poker, 69
 - Prisoners' Dilemma, 100–103
 - Rock, Paper, Scissors, 99, 117
 - sequential move, 89–98
 - simultaneous move, 99–104
 - sticks, 94
 - surprise exam paradox, 98
- Giffen good, 207
- global warming, 41, 163
- government
 - block grants, 220
 - City of Seattle budget, 223
 - federal budget, 221
 - fiscal activities, 219
 - regulatory activities, 220
 - size of, 219
 - Washington State budget, 222
- green
 - yellow and blue makes, 146
- Guy, Richard, 73
- Hall, Monty, 24
- hand
 - invisible, 4, 213, 214
- Harriman, E.H., 96
- Hicksian (compensated) demand curve, 204
- income effect, 178, 206
- indifference, 23
- indifference curve, 191
 - and expected value, 23–24
- induction
 - backward, 89
- inferior good, 178, 206
- inflation, 55–59
- integral, 185, 188
- integrals, 13
- interest
 - compound, 38, 40
 - nominal rate of, 55–59
 - real rate of, 55–59
 - rule of 72, 39
- interior solution, 7, 114–116, 200
- inverse demand curve, 225
- invisible hand, 4, 213, 214
- isoquant, 192
- jokes
 - \$10 bill on the ground, 61
 - José the bank robber, 121
 - three wishes, 69

- Zen monk, 3
- Lagrangian, 198
- last dollar rule, 195
- latex (or L^AT_EX), viii
- Linux, vii
- liquidity, 64
- local government, *see* government
- lump sum, 29–31, 39, 42, 60, 214, 215
- macroeconomics, vii, 220
- marginal
 - analysis, 6, 51
 - benefit, 6, 183, 186–189, 200, 201
 - cost, 6, 183–189, 200, 201, 226, 227
 - product (of labor/capital), 193, 197, 201
 - value of, 201
 - rate of substitution (MRS), 192
 - rate of technical substitution (MRTS), 193
 - rate of technical transformation (MRTT), 197
 - rate of transformation (MRT), 197
 - revenue, 226, 227
 - utility, 192, 196
- margins, 183–189
 - and demand curve, 186
 - and supply curve, 183
 - in terms of sticks and carrots, 188
- market
 - competitive, 137
 - demand curve, 144, 174
 - equilibrium, 145
 - market-clearing price, 145
 - perfectly competitive, 138
 - price, 143
 - quantity transacted, 143
 - supply curve, 144, 173
 - total expenditure, 143
 - total revenue, 143
- Marshall, Alfred, 183
- Marshallian demand curve, 202
- Maximum Sustainable Yield, 34, 35, 51
- mechanism design, 71
- microeconomics, vii, 72
- monopoly, vii, 6, 83, 99, 138, 139, 166, 225–229, 233, 290
 - and elasticity, 166
 - barriers to entry, 139
 - government-regulated, 7
 - in the long run, 138
 - price discrimination, 228
 - uniform pricing, 85
- Monty Hall Problem, 24
- Nalebuff, Barry J., 94
- Napster, 75
- Nasar, Sylvia, 110
- Nash equilibrium, 108
 - algorithms for finding, 109
 - mixed strategy, 111–116
- Nash, John, 110
- necessary first-order condition, 9, 45, 50, 194, 198, 199, 226, 229–232
- Newman, Paul, 96
- Nobel Prize, 65
 - Coase, Ronald, 75
 - King, B.B., 75
 - Nash, John, 110
- normal good, 178, 206
- oligopoly, 233
- Olson, Mancur, 3
- OPEC, 101
- open source, vii, viii
- opportunity cost, 6
- Pareto
 - efficient, 79–87
 - improvement, 79–87
 - inefficient, 79–87
- Pareto, Vilfrado, 79
- patent race, 126
- payoff matrix, 99
- perpetuity, 33, 35, 39, 51, 110
- poker, 69
- present value, 29–31, 35, 37–42, 45–48, 50–53, 55, 59, 60, 110, 111
 - and budget constraint, 37

- of annuity, 32
 - of Louisiana Purchase, 41, 60
 - of perpetuity, 33
- price discrimination, 228
- price theory, 137
- price-taker, 138
- principal-agent problem, 228
- Prisoners' Dilemma, 100–103
 - finitely repeated, 102
- producer surplus, 185
- production function, 192
 - Cobb-Douglas, 192
- profit, 4
 - accounting, 4
 - economic, 4
- purchasing power, 206
- quantum, vii, 292
- Raiffa
 - Howard, 25
- Ramen noodles, 178, 206
- Raymond, Eric S., viii
- Redford, Robert, 96
- repeated game, 102
- repeated games
 - infinitely, 110
 - trigger strategy, 110
- risk, 4, 19, 25, 64, 65
 - premium, 64, 65
 - equity premium puzzle, 65
 - reducing with diversification, 21
- Robertson, Jack, 73
- Robinson Crusoe model, 76
- Rock, Paper, Scissors, 99, 117
- rule
 - last dollar, 195
- rule of 72, 39
- Samuelson, W.F., 127
- Schwartz, Jesse A., 127
- Second Welfare Theorem, 214
- Selfridge, John, 73
- Shapiro, Carl, 84, 228
- Silberberg, Eugene, 10
- slope, 159
 - and derivative, 7
 - and elasticity, 169, 171
 - and partial derivative, 9
 - and tax burdens, 160, 171
 - of budget constraint, 37, 197, 198, 203
 - of indifference curve, 192, 197, 203
 - of isocost, 197, 198
 - of isoquant, 193, 198
 - of supply/demand curve, 159, 163, 169, 171
 - versus elasticity, 163
- Smith
 - Adam, 213
- sports
 - game theory applied to, 70, 95
- stage game, 102
- state government, *see* government
- subgame, 89
- substitute, 148
 - defined using elasticity, 170
- substitution effect, 178, 206
- Suen, Wing, 10
- sufficient second-order condition, 9
- supply, 143, 173
- supply curve
 - in terms of margins, 183
 - increases and decreases, 146
 - individual, 173
 - market, 144, 173
- surplus
 - consumer, 187
 - producer, 185
- tax
 - carbon, 163
 - equivalence, 156–157
 - per-unit, 153–156
- taxes, 153–161
 - and elasticity, 169
 - in terms of margins, 185, 187
- Taylor, Alan, 74, 95
- total benefit, 187
- tree
 - game, 89
- trees

- economic management of, 45–48
- trigger strategy, 110
- Trump, Donald, 74
- uncertainty, 64
- uncompensated demand curve, 205
- unfair bet, 20
- utility function, 191
 - Cobb-Douglas, 191
- Varian, Hal, 84, 228
- variance, 20, 22
- von Neumann, John, 70
- war
 - game theory applied to, 70
- Webb, William, 73
- welfare economics, 213
 - First Welfare Theorem, 213
 - policy implications, 215
 - Second Welfare Theorem, 214
- winner's curse, 127
- Wolfstetter, Elmar, 121
- yellow
 - and blue makes green, 146