HAL R. VARIAN

INTERMEDIATE MICROECONOMICS WITH CALCULUS

MEDIA UPDATE

Intermediate Microeconomics

With Calculus

First Edition

Intermediate Microeconomics

With Calculus

First Edition

Hal R. Varian

Google and University of California at Berkeley

W. W. Norton & Company • New York • London

W. W. Norton & Company has been independent since its founding in 1923, when William Warder Norton and Mary D. Herter Norton first published lectures delivered at the People's Institute, the adult education division of New York City's Cooper Union. The firm soon expanded its program beyond the Institute, publishing books by celebrated academics from America and abroad. By mid-century, the two major pillars of Norton's publishing program—trade books and college texts—were firmly established. In the 1950s, the Norton family transferred control of the company to its employees, and today—with a staff of four hundred and a comparable number of trade, college, and professional titles published each year—W. W. Norton & Company stands as the largest and oldest publishing house owned wholly by its employees.

Printed in the United States of America
FIRST EDITION MEDIA UPDATE

Editor: Eric Svendsen
Project Editor: David Bradley
Editorial Assistant: Jeannine Hennawi
Managing Editor, College: Marian Johnson
Managing Editor, College Digital Media: Kim Yi
Production Manager: Elizabeth Marotta
Media Editor: Miryam Chandler
Marketing Manager: Janise Turso
Illustrations: Westchester Publishing Services
Manufacturing: LSC Communications, Crawfordsville
TEXnician: Hal Varian

Cover design: M80 design / W. Youssi
Cover photographs: top image (matrix of digits): agsandrew/shutterstock
bottom image (blackboard with a math scribble): ykh/shutterstock

ISBN 978-0-393-68993-8

W. W. Norton & Company, Inc., 500 Fifth Avenue, New York, N.Y. 10110
www.wwnorton.com
W. W. Norton & Company, Ltd., 15 Carlisle Street, London W1D 3BS

1 2 3 4 5 6 7 8 9 0

To Carol

CONTENTS

3 Preferences

4 Utility

5 Choice

6 Demand

7 Revealed Preference

8 Slutsky Equation

9 Buying and Selling

10 Intertemporal Choice

11 Asset Markets

12 Uncertainty

13 Risky Assets

14 Consumer's Surplus

15 Market Demand

16 Equilibrium

17 Measurement

18 Auctions

19 Technology

20 Profit Maximization

21 Cost Minimization

22 Cost Curves

23 Firm Supply

24 Industry Supply

25 Monopoly

26 Monopoly Behavior

27 Factor Markets

28 Oligopoly

29 Game Theory

30 Game Applications

31 Behavioral Economics

32 Exchange

33 Production

34 Welfare

35 Externalities

36 Information Technology

37 Public Goods

38 Asymmetric Information

Mathematical Appendix

PREFACE

This book is my classic *Intermediate Microeconomics* text with the mathematical treatment that was previously in the chapter appendices incorporated into the body of the chapters. This makes the analysis flow somewhat better for those students who are comfortable with elementary calculus.

My aim in writing the original text was to present a treatment of the methods of microeconomics that would allow students to apply these tools on their own and not just passively absorb the predigested cases described in the text. I have found that the best way to do this is to emphasize the fundamental conceptual foundations of microeconomics and to provide concrete examples of their application rather than to attempt to provide an encyclopedia of terminology and anecdote.

The calculus treatment will, I hope, be helpful to students who have appropriate backgrounds. However, it should be remembered that one can go a long way with a few simple facts about linear demand functions and supply functions and some elementary algebra. It is perfectly possible to be analytical without being excessively mathematical.

The distinction is worth emphasizing. An analytical approach to economics is one that uses rigorous, logical reasoning. This does not necessarily require the use of advanced mathematical methods. The language of mathematics certainly helps to ensure a rigorous analysis and using it is undoubtedly the best way to proceed when possible, but it may not be appropriate for all students. This is why there are two versions of the text.

Calculus offers deeper ways to examine the same issues that one can also explore verbally and graphically. Many arguments are much simpler with a little mathematics, and all economics students should learn that. In many cases I've found that with a little motivation, and a few nice economic examples, students become quite enthusiastic about looking at things from an analytic perspective.

There are several other innovations in this text. First, the chapters are generally very short. I've tried to make most of them roughly "lecture size," so that they can be read in one sitting. I have followed the standard order of discussing first consumer theory and then producer theory, but I've spent a bit more time on consumer theory than is normally the case. This is not because I think that consumer theory is necessarily the most important part of microeconomics; rather, I have found that this is the material that students find the most mysterious, so I wanted to provide a more detailed treatment of it.

Second, I've tried to put in a lot of examples of how to use the theories described here. In most books, students look at a lot of diagrams of shifting curves, but they don't see much algebra, or much calculation of any sort for that matter. But it is the algebra that is used to solve problems in practice. Graphs can provide insight, but the real power of economic analysis comes in calculating quantitative answers to economic problems. Every economics student should be able to translate an economic story into an equation or

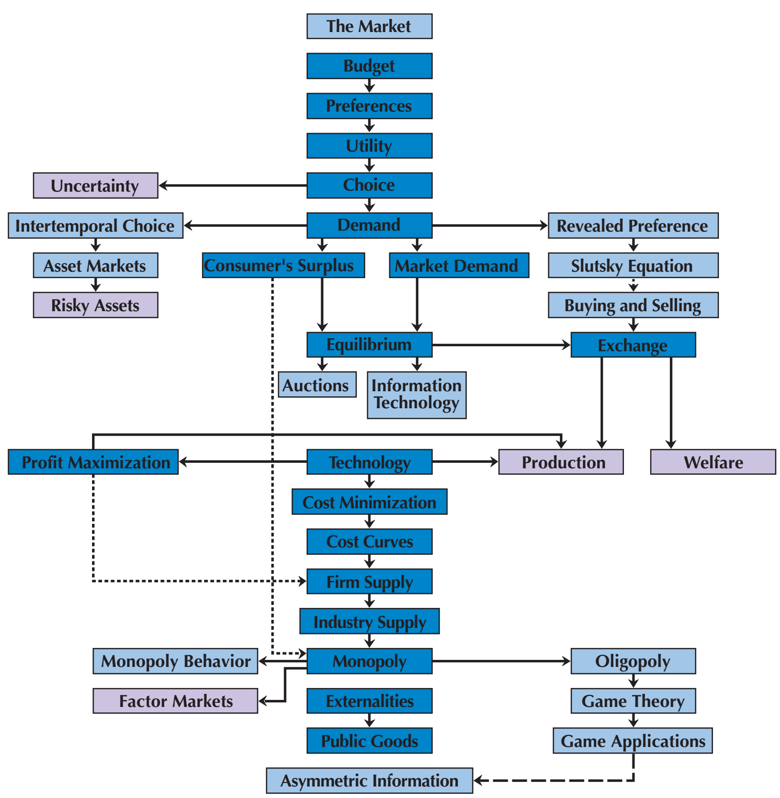

a numerical example, but all too often the development of this skill is neglected. For this reason I have also provided a workbook that I feel is an integral accompaniment to this book. The workbook was written with my colleague Theodore Bergstrom, and we have put a lot of effort into generating interesting and instructive problems. We think that it provides an important aid to the student of microeconomics.

Third, I believe that the treatment of the topics in this book is more accurate than is usually the case in intermediate micro texts. It is true that I've sometimes chosen special cases to analyze when the general case is too difficult, but I've tried to be honest about that when I did it. In general, I've tried to spell out every step of each argument in detail. I believe that the discussion I've provided is not only more complete and more accurate than usual, but this attention to detail also makes the arguments easier to understand than the loose discussion presented in many other books.

There Are Many Paths to Economic Enlightenment

There is more material in this book than can comfortably be taught in one semester, so it is worthwhile picking and choosing carefully the material that you want to study. If you start on page 1 and proceed through the chapters in order, you will run out of time long before you reach the end of the book. The modular structure of the book allows the instructor a great deal of freedom in choosing how to present the material, and I hope that more people will take advantage of this freedom. The chart above illustrates the chapter dependencies.

The darker colored chapters are "core" chapters—they should probably be covered in every intermediate microeconomics course. The lighter-colored chapters are "optional" chapters: I cover some but not all of these every semester. The gray chapters are chapters I usually don't cover in my course, but they could easily be covered in other courses. A solid line going from Chapter A to Chapter B means that Chapter A should be read before chapter B. A broken line means that Chapter B requires knowing some material in Chapter A, but doesn't depend on it in a significant way.

I generally cover consumer theory and markets and then proceed directly to producer theory. Another popular path is to do exchange right after consumer theory; many instructors prefer this route and I have gone to some trouble to make sure that this path is possible.

Some people like to do producer theory before consumer theory. This is possible with this text, but if you choose this path, you will need to supplement the textbook treatment. The material on isoquants, for example, assumes that the students have already seen indifference curves.

Much of the material on public goods, externalities, law, and information can be introduced earlier in the course. I've arranged the material so that it is quite easy to put it pretty much wherever you desire.

Similarly, the material on public goods can be introduced as an illustration of Edgeworth box analysis. Externalities can be introduced right after the discussion of cost curves, and topics from the information chapter can be introduced almost anywhere after students are familiar with the approach of economic analysis.

Changes for this Edition

The text of the book is closely aligned with *Intermediate Microeconomics.* I have added a new chapter on measurement which describes some of the issues involved in estimating economic relationships. The idea is to introduce the student to some basic concepts from econometrics and try to bridge the theoretical treatment in the book with the practical problems encountered in practice.

I have offered some new examples drawn from Silicon Valley firms such as Apple, eBay, Google, Yahoo and others. I discuss topics such as the complementarity between the iPod and iTunes, the positive feedback associated with companies such as Facebook, and the ad auction models used by Google, Microsoft, and Yahoo. I believe that these are fresh and interesting examples of economics in action.

I've also added an extended discussion of mechanism design issues, including two-sided matching markets and the Vickrey-Clarke-Groves mechanisms. This field, which was once primarily theoretical in nature, has now taken on considerable practical importance.

The Test Bank and Workbook

The workbook, *Workouts in Intermediate Microeconomics,* is an integral part of the course. It contains hundreds of fill-in-the-blank exercises that lead the students through the steps of actually applying the tools they have learned in the textbook. In addition to the exercises, *Workouts* contains a collection of short multiple-choice quizzes based on the workbook problems in each chapter. Answers to the quizzes are also included in *Workouts.* These quizzes give a quick way for the student to review the material he or she has learned by working the problems in the workbook.

But there is more . . . instructors who have adopted *Workouts* for their course can make use of the *Test Bank* offered with the textbook. The *Test Bank* contains several alternative versions of each *Workouts* quiz. The questions in these quizzes use different numerical values but the same internal logic. They can be used to provide additional problems for students to practice, or quizzes to be taken in class. Grading is quick and reliable because the quizzes are multiple choice and can be graded electronically.

In our course, we tell the students to work through all the quiz questions for each chapter, either by themselves or with a study group. Then during the term we have a short in-class quiz every other week or so, using the alternative versions from the *Test Bank.* These are essentially the *Workouts* quizzes with different numbers. Hence, students who have done their homework find it easy to do well on the quizzes.

We firmly believe that you can't learn economics without working some problems. The quizzes provided in *Workouts* and in the *Test Bank* make the learning process much easier for both the student and the teacher.

A hard copy of the *Test Bank* is available from the publisher, as is the textbook's *Instructor's Manual*, which includes my teaching suggestions and lecture notes for each chapter of the textbook, and solutions to the exercises in *Workouts.*

What is the Media Update Edition?

Over the years, instructors have requested additional support for this text, especially for a way to assess their students online, and for more in-class lecture support. We also see more students purchasing their texts electronically, expecting more ebook functionality than this text has offered. This Media Update Edition was created to support these needs. While the content in this edition is the same as the current editions of the text, we have made important changes to the text package.

Working with a team of experienced professors and instructors, Norton has developed online homework for this text using their Smartwork5 homework engine. It has a large selection of new, often multi-step problems for each chapter, approximately 50 percent of which are algorithmic. In addition, you will see the figures in this text are now color-coded. Our feedback is that most students have become accustomed to this style of presentation, especially when using electronic formats. These figures are available now electronically for use in your presentations, and a new set of lecture-ready slides has also been created. Finally, also available are a set of graphing interactives on the following topics: the Production Possibilities Frontier, Preferences and the Budget Constraint, Edgeworth boxes, Individual and Market Supply, Individual and Market Demand, Perfect Competition and Profit Maximization, and Long-Run Cost Minimization. They give you and your students the chance to change variables and see the results. For complete information, and to sample these on-line features, please visit `https://digital.wwnorton.com/calcmedia`.

The Production of the Book

The entire book was typeset by the author using TEX, the wonderful typesetting system designed by Donald Knuth. I worked on a Linux system and using GNU `emacs` for editing, `rcs` for version control and the TEX Live system for processing. I used `makeindex` for the index, and Hàn Thế Thành's `pdftex` for the text and figures.

Acknowledgments

Several people contributed to this project. First, I must thank my editorial assistants for the first edition, John Miller and Debra Holt. John provided many comments, suggestions, and exercises based on early drafts of this text and made a significant contribution to the coherence of the final product. Debra did a careful proofreading and consistency check during the final stages and helped in preparing the index.

The following individuals provided me with many useful suggestions and comments during the preparation of the first edition: Ken Binmore (University of Michigan), Mark Bagnoli (Indiana University), Larry Chenault (Miami University), Jonathan Hoag (Bowling Green State University), Allen Jacobs (M.I.T.),

John McMillan (University of California at San Diego), Hal White (University of California at San Diego), and Gary Yohe (Wesleyan University). In particular, I would like to thank Dr. Reiner Buchegger, who prepared the German translation, for his close reading of the first edition and for providing me with a detailed list of corrections. Other individuals to whom I owe thanks for suggestions prior to the first edition are Theodore Bergstrom, Jan Gerson, Oliver Landmann, Alasdair Smith, Barry Smith, and David Winch.

My editorial assistants for the second edition were Sharon Parrott and Angela Bills. They provided much useful assistance with the writing and editing. Robert M. Costrell (University of Massachusetts at Amherst), Ashley Lyman (University of Idaho), Daniel Schwallie (Case-Western Reserve), A. D. Slivinskie (Western Ontario), and Charles Plourde (York University) provided me with detailed comments and suggestions about how to improve the second edition.

In preparing the third edition I received useful comments from the following individuals: Doris Cheng (San Jose), Imre Csekó (Budapest), Gregory Hildebrandt (UCLA), Jamie Brown Kruse (Colorado), Richard Manning (Brigham Young), Janet Mitchell (Cornell), Charles Plourde (York University), Yeung-Nan Shieh (San Jose), and John Winder (Toronto). I especially want to thank Roger F. Miller (University of Wisconsin), and David Wildasin (Indiana) for their detailed comments, suggestions, and corrections.

The fifth edition benefited from comments by Kealoah Widdows (Wabash College), William Sims (Concordia University), Jennifer R. Reinganum (Vanderbilt University), and Paul D. Thistle (Western Michigan University).

I received comments that helped in preparation of the sixth edition from James S. Jordon (Pennsylvania State University), Brad Kamp (University of South Florida), Sten Nyberg (Stockholm University), Matthew R. Roelofs (Western Washington University), Maarten-Pieter Schinkel (University of Maastricht), and Arthur Walker (University of Northumbria).

The seventh edition received reviews by Irina Khindanova (Colorado School of Mines), Istvan Konya (Boston College), Shomu Banerjee (Georgia Tech), Andrew Helms (University of Georgia), Marc Melitz (Harvard University), Andrew Chatterjea (Cornell University), and Cheng-Zhong Qin (UC Santa Barbara).

Finally, I received helpful comments on the eighth edition from Kevin Balsam (Hunter College), Clive Belfield (Queens College, CUNY), Reiner Buchegger (Johannes Kepler University), Lars Metzger (Technische Universitaet Dortmund), Jeffrey Miron (Harvard University), Babu Nahata (University of Louisville), and Scott J. Savage (University of Colorado). I am particularly grateful to Carola Conces who provided research assistance in merging the calculus appendices with the chapter bodies.

I have received helpful comments on the ninth edition from Richard K. Anderson (Texas A&M University), Yong Chao (University of Louisville), Arup Daripa (University of London), Roberta Dessi (Toulouse School of Economics), Chris Doyle (University of Warwick), Thomas A. Gresik (University of Notre Dame), Marco Haan (University of Groningen), Babu Nahata (University of Louisville), Ann-Katrin Roesler (University of Michigan), John G. Riley (UCLA), Kevin Shaver (University of Pittsburgh), Garance Genicot (Georgetown University), and Utku Ünver (Boston College).

Media Update Edition

- Powerpoints: David Lang (Sacramento State University)
- Graphing Interactives: Brian Jenkins (University of California, Irvine)

SmartWork

- Authors: Alan Green (Stetson University), Bob Jones (Rensselaer Polytechnic Institute), Miguel Delgado Helleseter (California State UniversityChannel Islands), Kelvin Wong (Arizona State University), Margherita Negri (University of St. Andrews), Matt Roelofs (Western Washington University), Brian Goegan (Arizona State University), Ty Robbins (Manchester University)
- Reviewers: Benoit Crutzen (Erasmus University Rotterdam), Clive Chapple (University of British Columbia), John S. Thompson (University of Texas at Austin), Ming Li (Concordia University), Gerald Oettinger (University of Texas at Austin)
- Content testers: Billy Foster (Arizona State University), Kacey Douglas (Arizona State University), Jason Query (Western Washington University), Marian Rizov (Lincoln University), Mike Visser (Sonoma State University), Thoedoros Diasokos (University of Stirling)

Berkeley, California
January 2019

CHAPTER 1

THE MARKET

The conventional first chapter of a microeconomics book is a discussion of the "scope and methods" of economics. Although this material can be very interesting, it hardly seems appropriate to *begin* your study of economics with such material. It is hard to appreciate such a discussion until you have seen some examples of economic analysis in action.

So instead, we will begin this book with an *example* of economic analysis. In this chapter we will examine a model of a particular market, the market for apartments. Along the way we will introduce several new ideas and tools of economics. Don't worry if it all goes by rather quickly. This chapter is meant only to provide a quick overview of how these ideas can be used. Later on we will study them in substantially more detail.

1.1 Constructing a Model

Economics proceeds by developing **models** of social phenomena. By a model we mean a simplified representation of reality. The emphasis here is on the word "simple." Think about how useless a map on a one-to-one scale would be. The same is true of an economic model that attempts to describe every aspect of reality. A model's power stems from the elimination of irrelevant detail, which allows the economist to focus on the essential features of the economic reality he or she is attempting to understand.

Here we are interested in what determines the price of apartments, so we want to have a simplified description of the apartment market. There is a certain art to choosing the right simplifications in building a model. In general we want to adopt the simplest model that is capable of describing the economic situation we are examining. We can then add complications one at a time, allowing the model to become more complex and, we hope, more realistic.

The particular example we want to consider is the market for apartments in a medium-size midwestern college town. In this town there are two sorts of apartments. There are some that are adjacent to the university, and others that are farther away. The adjacent apartments are generally considered to be more desirable by students, since they allow easier access to the university. The apartments that are farther away necessitate taking a bus, or a long, cold bicycle ride, so most students would prefer a nearby apartment ... if they can afford one.

We will think of the apartments as being located in two large rings surrounding the university. The adjacent apartments are in the inner ring, while the rest are located in the outer ring. We will focus exclusively on the market for apartments in the inner ring. The outer ring should be interpreted as where people can go who don't find one of the closer apartments. We'll suppose that there are many apartments available in the outer ring, and their price is fixed at some known level. We'll be concerned solely with the determination of the price of the inner-ring apartments and who gets to live there.

An economist would describe the distinction between the prices of the two kinds of apartments in this model by saying that the price of the outer-ring apartments is an **exogenous variable**, while the price of the inner-ring apartments is an **endogenous variable**. This means that the price of the outer-ring apartments is taken as determined by factors not discussed in this particular model, while the price of the inner-ring apartments is determined by forces described in the model.

The first simplification that we'll make in our model is that all apartments are identical in every respect except for location. Thus it will make sense to speak of "the price" of apartments, without worrying about whether the apartments have one bedroom, or two bedrooms, or whatever.

But what determines this price? What determines who will live in the inner-ring apartments and who will live farther out? What can be said about the desirability of different economic mechanisms for allocating apartments? What concepts can we use to judge the merit of different assignments of apartments to individuals? These are all questions that we want our model to address.

1.2 Optimization and Equilibrium

Whenever we try to explain the behavior of human beings we need to have a framework on which our analysis can be based. In much of economics we use a framework built on the following two simple principles.

The optimization principle: People try to choose the best patterns of consumption that they can afford.

The equilibrium principle: Prices adjust until the amount that people demand of something is equal to the amount that is supplied.

Let us consider these two principles. The first is *almost* tautological. If people are free to choose their actions, it is reasonable to assume that they try to choose things they want rather than things they don't want. Of course there are exceptions to this general principle, but they typically lie outside the domain of economic behavior.

The second notion is a bit more problematic. It is at least conceivable that at any given time peoples' demands and supplies are not compatible, and hence something must be changing. These changes may take a long time to work themselves out, and, even worse, they may induce other changes that might "destabilize" the whole system.

This kind of thing can happen . . . but it usually doesn't. In the case of apartments, we typically see a fairly stable rental price from month to month. It is this *equilibrium* price that we are interested in, not in how the market gets to this equilibrium or how it might change over long periods of time.

It is worth observing that the definition used for equilibrium may be different in different models. In the case of the simple market we will examine in this chapter, the demand and supply equilibrium idea will be adequate for our needs. But in more general models we will need more general definitions of equilibrium. Typically, equilibrium will require that the economic agents' actions must be consistent with each other.

How do we use these two principles to determine the answers to the questions we raised above? It is time to introduce some economic concepts.

1.3 The Demand Curve

Suppose that we consider all of the possible renters of the apartments and ask each of them the maximum amount that he or she would be willing to pay to rent one of the apartments.

Let's start at the top. There must be someone who is willing to pay the highest price. Perhaps this person has a lot of money, perhaps he is very lazy and doesn't want to walk far ... or whatever. Suppose that this person is willing to pay \$500 a month for an apartment.

If there is only one person who is willing to pay \$500 a month to rent an apartment, then if the price for apartments were \$500 a month, exactly one apartment would be rented—to the one person who was willing to pay that price.

Suppose that the next highest price that anyone is willing to pay is \$490. Then if the market price were \$499, there would still be only one apartment rented: the person who was *willing* to pay \$500 would rent an apartment, but the person who was willing to pay \$490 wouldn't. And so it goes. Only one apartment would be rented if the price were \$498, \$497, \$496, and so on ... until we reach a price of \$490. At that price, exactly two apartments would be rented: one to the \$500 person and one to the \$490 person.

Similarly, two apartments would be rented until we reach the maximum price that the person with the *third* highest price would be willing to pay, and so on.

Economists call a person's maximum willingness to pay for something that person's **reservation price**. The reservation price is the highest price that a given person will accept and still purchase the good. In other words, a person's reservation price is the price at which he or she is just indifferent between purchasing or not purchasing the good. In our example, if a person has a reservation price p it means that he or she would be just indifferent between living in the inner ring and paying a price p and living in the outer ring.

Thus the number of apartments that will be rented at a given price p^* will just be the number of people who have a reservation price greater than or equal to p^*. For if the market price is p^*, then everyone who is willing to pay at least p^* for an apartment will want an apartment in the inner ring, and everyone who is not willing to pay p^* will choose to live in the outer ring.

We can plot these reservation prices in a diagram as in Figure 1.1. Here the price is depicted on the vertical axis and the number of people who are willing to pay that price or more is depicted on the horizontal axis.

Another way to view Figure 1.1 is to think of it as measuring how many people would want to rent apartments at any particular price. Such a curve is an example of a **demand curve**—a curve that relates the quantity demanded to price. When the market price is above \$500, zero apartments will be rented. When the price is between \$500 and \$490, one apartment will be rented. When it is between \$490 and the third highest reservation price, two apartments will be rented, and so on. The demand curve describes the quantity demanded at each of the possible prices.

The demand curve for apartments slopes down: as the price of apartments decreases more people will be willing to rent apartments. If there are many people and their reservation prices differ only slightly from person to

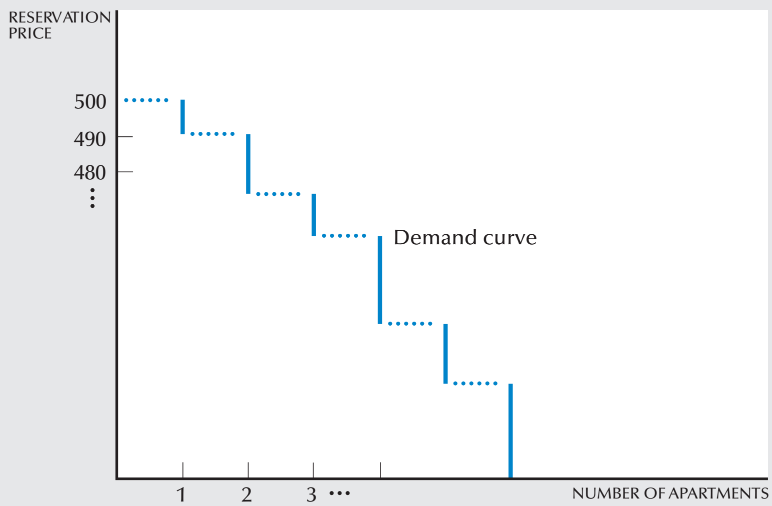

The demand curve for apartments. The vertical axis measures the market price and the horizontal axis measures how many apartments will be rented at each price.

Figure 1.1

person, it is reasonable to think of the demand curve as sloping smoothly downward, as in Figure 1.2. The curve in Figure 1.2 is what the demand curve in Figure 1.1 would look like if there were many people who want to rent the apartments. The “jumps” shown in Figure 1.1 are now so small relative to the size of the market that we can safely ignore them in drawing the market demand curve.

1.4 The Supply Curve

We now have a nice graphical representation of demand behavior, so let us turn to supply behavior. Here we have to think about the nature of the market we are examining. The situation we will consider is where there are many independent landlords who are each out to rent their apartments for the highest price the market will bear. We will refer to this as the case of a **competitive market**. Other sorts of market arrangements are certainly possible, and we will examine a few later.

For now, let's consider the case where there are many landlords who all operate independently. It is clear that if all landlords are trying to do the best they can and if the renters are fully informed about the prices the landlords charge, then the equilibrium price of all apartments in the inner ring must be the same. The argument is not difficult. Suppose instead that there is some high price, p_h, and some low price, p_l, being charged

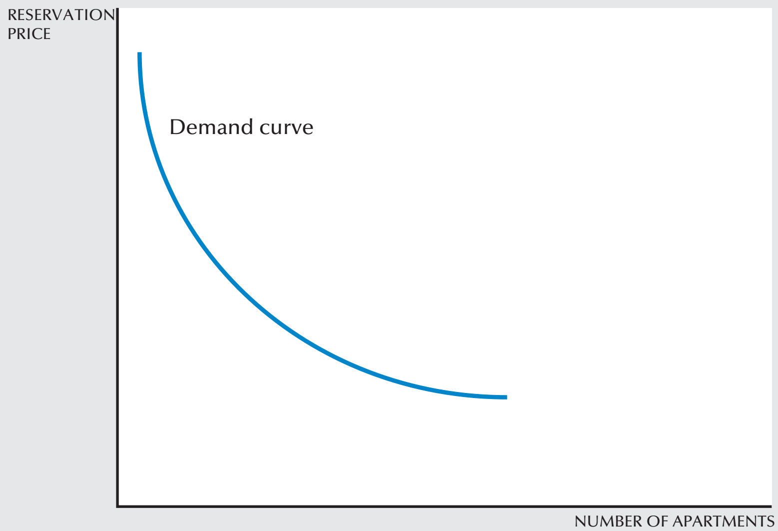

Figure 1.2 **Demand curve for apartments with many demanders.** Because of the large number of demanders, the jumps between prices will be small, and the demand curve will have the conventional smooth shape.

for apartments. The people who are renting their apartments for a high price could go to a landlord renting for a low price and offer to pay a rent somewhere between p_h and p_l. A transaction at such a price would make both the renter and the landlord better off. To the extent that all parties are seeking to further their own interests and are aware of the alternative prices being charged, a situation with different prices being charged for the same good cannot persist in equilibrium.

But what will this single equilibrium price be? Let us try the method that we used in our construction of the demand curve: we will pick a price and ask how many apartments will be supplied at that price.

The answer depends to some degree on the time frame in which we are examining the market. If we are considering a time frame of several years, so that new construction can take place, the number of apartments will certainly respond to the price that is charged. But in the "short run"—within a given year, say—the number of apartments is more or less fixed. If we consider only this short-run case, the supply of apartments will be constant at some predetermined level.

The **supply curve** in this market is depicted in Figure 1.3 as a vertical line. Whatever price is being charged, the same number of apartments will be rented, namely, all the apartments that are available at that time.

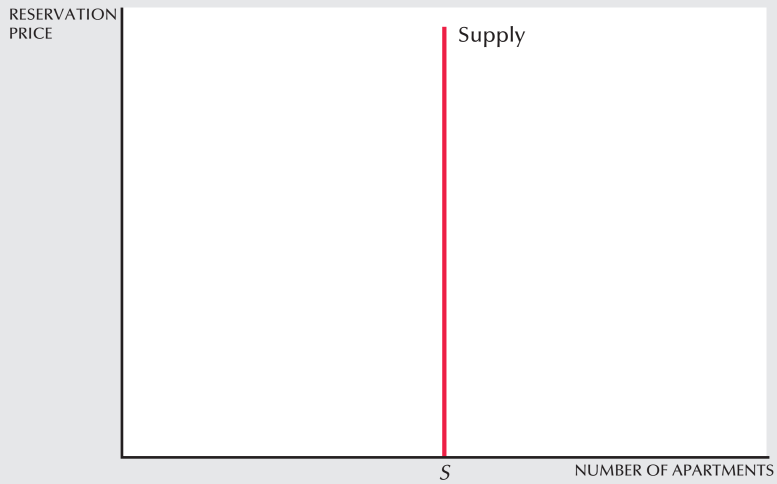

Short-run supply curve. The supply of apartments is fixed in the short run.

Figure 1.3

1.5 Market Equilibrium

We now have a way of representing the demand and the supply side of the apartment market. Let us put them together and ask what the equilibrium behavior of the market is. We do this by drawing both the demand and the supply curve on the same graph in Figure 1.4.

In this graph we have used p^* to denote the price where the quantity of apartments demanded equals the quantity supplied. This is the **equilibrium price** of apartments. At this price, each consumer who is willing to pay at least p^* is able to find an apartment to rent, and each landlord will be able to rent apartments at the going market price. Neither the consumers nor the landlords have any reason to change their behavior. This is why we refer to this as an *equilibrium:* no change in behavior will be observed.

To better understand this point, let us consider what would happen at a price other than p^*. For example, consider some price $p < p^*$ where demand is greater than supply. Can this price persist? At this price at least some of the landlords will have more renters than they can handle. There will be lines of people hoping to get an apartment at that price; there are more people who are willing to pay the price p than there are apartments. Certainly some of the landlords would find it in their interest to raise the price of the apartments they are offering.

Similarly, suppose that the price of apartments is some p greater than p^*.

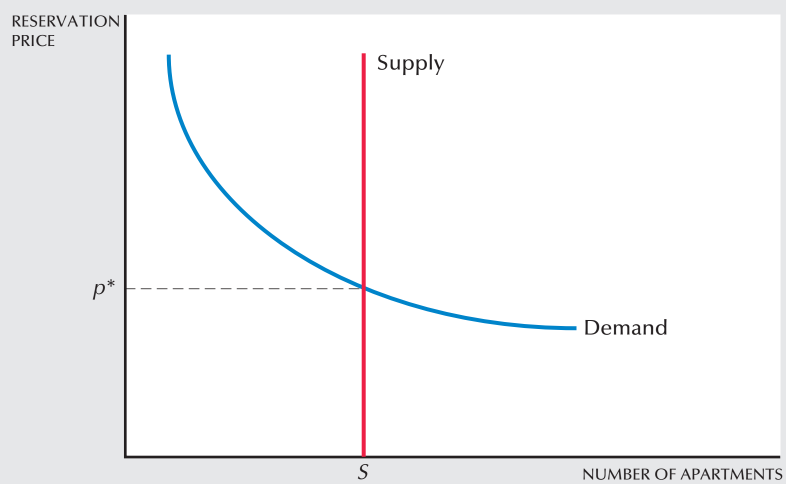

Figure 1.4 **Equilibrium in the apartment market**. The equilibrium price, p^*, is determined by the intersection of the supply and demand curves.

Then some of the apartments will be vacant: there are fewer people who are willing to pay p than there are apartments. Some of the landlords are now in danger of getting no rent at all for their apartments. Thus they will have an incentive to lower their price in order to attract more renters.

If the price is above p^* there are too few renters; if it is below p^* there are too many renters. Only at the price of p^* is the number of people who are willing to rent at that price equal to the number of apartments available for rent. Only at that price does demand equal supply.

At the price p^* the landlords' and the renters' behaviors are compatible in the sense that the number of apartments demanded by the renters at p^* is equal to the number of apartments supplied by the landlords. This is the equilibrium price in the market for apartments.

Once we've determined the market price for the inner-ring apartments, we can ask who ends up getting these apartments and who is exiled to the farther-away apartments. In our model there is a very simple answer to this question: in the market equilibrium everyone who is willing to pay p^* or more gets an apartment in the inner ring, and everyone who is willing to pay less than p^* gets one in the outer ring. The person who has a reservation price of p^* is just indifferent between taking an apartment in the inner ring and taking one in the outer ring. The other people in the inner ring are getting their apartments at less than the maximum they would be willing to pay for them. Thus the assignment of apartments to renters is determined by how much they are willing to pay.

1.6 Comparative Statics

Now that we have an economic model of the apartment market, we can begin to use it to analyze the behavior of the equilibrium price. For example, we can ask how the price of apartments changes when various aspects of the market change. This kind of an exercise is known as **comparative statics**, because it involves comparing two "static" equilibria without worrying about how the market moves from one equilibrium to another.

The movement from one equilibrium to another can take a substantial amount of time, and questions about how such movement takes place can be very interesting and important. But we must walk before we can run, so we will ignore such dynamic questions for now. Comparative statics analysis is only concerned with comparing equilibria, and there will be enough questions to answer in this framework for the present.

Let's start with a simple case. Suppose that the supply of apartments is increased, as in Figure 1.5.

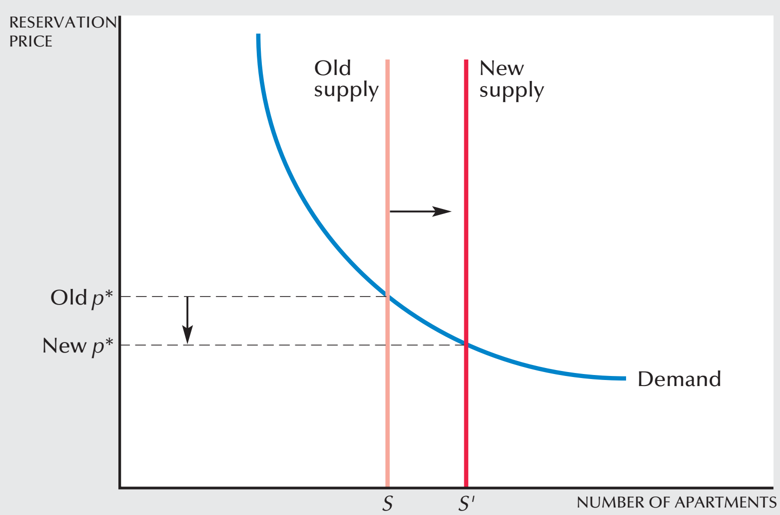

Figure 1.5 **Increasing the supply of apartments**. As the supply of apartments increases, the equilibrium price decreases.

It is easy to see in this diagram that the equilibrium price of apartments will fall. Similarly, if the supply of apartments were reduced the equilibrium price would rise.

Let's try a more complicated—and more interesting—example. Suppose that a developer decides to turn several of the apartments into condominiums. What will happen to the price of the remaining apartments?

Your first guess is probably that the price of apartments will go up, since the supply has been reduced. But this isn't necessarily right. It is true that the supply of apartments to rent has been reduced. But the *demand for apartments* has been reduced as well, since some of the people who were renting apartments may decide to purchase the new condominiums.

It is natural to assume that the condominium purchasers come from those who already live in the inner-ring apartments—those people who are willing to pay more than p^* for an apartment. Suppose, for example, that the demanders with the 10 highest reservation prices decide to buy condos rather than rent apartments. Then the new demand curve is just the old demand curve with 10 fewer demanders at each price. Since there are also 10 fewer apartments to rent, the new equilibrium price is just what it was before, and exactly the same people end up living in the inner-ring apartments. This situation is depicted in Figure 1.6. Both the demand curve and the supply curve shift left by 10 apartments, and the equilibrium price remains unchanged.

Most people find this result surprising. They tend to see just the reduction in the supply of apartments and don't think about the reduction in demand. The case we've considered is an extreme one: *all* of the condo purchasers were former apartment dwellers. But the other case—where none of the condo purchasers were apartment dwellers—is even more extreme.

The model, simple though it is, has led us to an important insight. If we want to determine how conversion to condominiums will affect the apartment market, we have to consider not only the effect on the supply of apartments but also the effect on the demand for apartments.

Let's consider another example of a surprising comparative statics analysis: the effect of an apartment tax. Suppose that the city council decides that there should be a tax on apartments of $50 a year. Thus each landlord will have to pay $50 a year to the city for each apartment that he owns. What will this do to the price of apartments?

Most people would think that at least some of the tax would get passed along to apartment renters. But, rather surprisingly, that is not the case. In fact, the equilibrium price of apartments will remain unchanged!

In order to verify this, we have to ask what happens to the demand curve and the supply curve. The supply curve doesn't change—there are just as many apartments after the tax as before the tax. And the demand curve doesn't change either, since the number of apartments that will be rented at each different price will be the same as well. If neither the demand curve nor the supply curve shifts, the price can't change as a result of the tax.

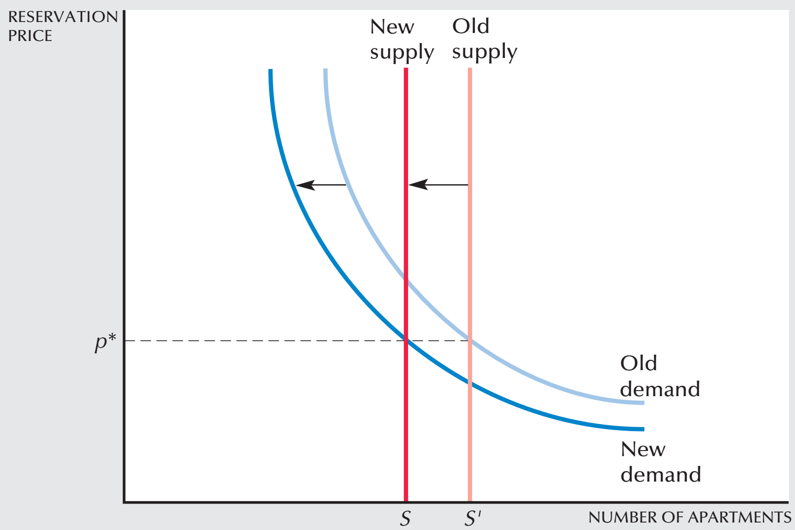

Effect of creating condominiums. If demand and supply both shift left by the same amount the equilibrium price is unchanged.

Figure 1.6

Here is a way to think about the effect of this tax. Before the tax is imposed, each landlord is charging the highest price that he can get that will keep his apartments occupied. The equilibrium price p^* is the highest price that can be charged that is compatible with all of the apartments being rented. After the tax is imposed can the landlords raise their prices to compensate for the tax? The answer is no: if they could raise the price and keep their apartments occupied, they would have already done so. If they were charging the maximum price that the market could bear, the landlords couldn't raise their prices any more: none of the tax can get passed along to the renters. The landlords have to pay the entire amount of the tax.

This analysis depends on the assumption that the supply of apartments remains fixed. If the number of apartments can vary as the tax changes, then the price paid by the renters will typically change. We'll examine this kind of behavior later on, after we've built up some more powerful tools for analyzing such problems.

1.7 Other Ways to Allocate Apartments

In the previous section we described the equilibrium for apartments in a competitive market. But this is only one of many ways to allocate a resource; in this section we describe a few other ways. Some of these may sound rather strange, but each will illustrate an important economic point.

The Discriminating Monopolist

First, let us consider a situation where there is one dominant landlord who owns all of the apartments. Or, alternatively, we could think of a number of individual landlords getting together and coordinating their actions to act as one. A situation where a market is dominated by a single seller of a product is known as a **monopoly**.

In renting the apartments the landlord could decide to auction them off one by one to the highest bidders. Since this means that different people would end up paying different prices for apartments, we will call this the case of the **discriminating monopolist**. Let us suppose for simplicity that the discriminating monopolist knows each person's reservation price for apartments. (This is not terribly realistic, but it will serve to illustrate an important point.)

This means that he would rent the first apartment to the fellow who would pay the most for it, in this case $500. The next apartment would go for $490 and so on as we moved down the demand curve. Each apartment would be rented to the person who was willing to pay the most for it.

Here is the interesting feature of the discriminating monopolist: *exactly the same people will get the apartments as in the case of the market solution,* namely, everyone who valued an apartment at more than p^*. The last person to rent an apartment pays the price p^*—the same as the equilibrium price in a competitive market. The discriminating monopolist's attempt to maximize his own profits leads to the same allocation of apartments as the supply and demand mechanism of the competitive market. The amount the people *pay* is different, but who gets the apartments is the same. It turns out that this is no accident, but we'll have to wait until later to explain the reason.

The Ordinary Monopolist

We assumed that the discriminating monopolist was able to rent each apartment at a different price. But what if he were forced to rent all apartments at the same price? In this case the monopolist faces a tradeoff: if he chooses a low price he will rent more apartments, but he may end up making less money than if he sets a higher price.

Let us use $D(p)$ to represent the demand function—the number of apartments demanded at price p. Then if the monopolist sets a price p, he will rent $D(p)$ apartments and thus receive a revenue of $pD(p)$. The revenue that the monopolist receives can be thought of as the area of a box: the height of the box is the price p and the width of the box is the number of apartments $D(p)$. The product of the height and the width—the area of the box—is the revenue the monopolist receives. This is the box depicted in Figure 1.7.

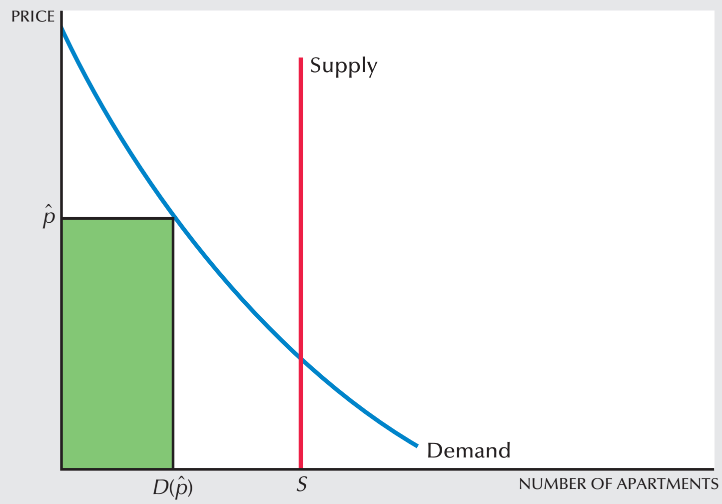

Revenue box. The revenue received by the monopolist is just the price times the quantity, which can be interpreted as the area of the box illustrated.

Figure 1.7

If the monopolist has no costs associated with renting an apartment, he would want to choose a price that has the largest associated revenue box. The largest revenue box in Figure 1.7 occurs at the price $\hat{p}$. In this case the monopolist will find it in his interest *not* to rent all of the apartments. In fact this will generally be the case for a monopolist. The monopolist will want to restrict the output available in order to maximize his profit. This means that the monopolist will generally want to charge a price that is higher than the equilibrium price in a competitive market, p^*. In the case of the ordinary monopolist, fewer apartments will be rented, and each apartment will be rented at a higher price than in the competitive market.

Rent Control

A third and final case that we will discuss will be the case of rent control. Suppose that the city decides to impose a maximum rent that can be charged for apartments, say p_{max}. We suppose that the price p_{max} is less than the equilibrium price in the competitive market, p^*. If this is so we would have a situation of **excess demand**: there are more people who are willing to rent apartments at p_{max} than there are apartments available. Who will end up with the apartments?

The theory that we have described up until now doesn't have an answer to this question. We can describe what will happen when supply equals demand, but we don't have enough detail in the model to describe what will happen if supply doesn't equal demand. The answer to who gets the apartments under rent control depends on who has the most time to spend looking around, who knows the current tenants, and so on. All of these things are outside the scope of the simple model we've developed. It may be that exactly the same people get the apartments under rent control as under the competitive market. But that is an extremely unlikely outcome. It is much more likely that some of the formerly outer-ring people will end up in some of the inner-ring apartments and thus displace the people who would have been living there under the market system. So under rent control the same number of apartments will be rented at the rent-controlled price as were rented under the competitive price: they'll just be rented to different people.

1.8 Which Way Is Best?

We've now described four possible ways of allocating apartments to people:

- The competitive market.
- A discriminating monopolist.
- An ordinary monopolist.
- Rent control.

These are four different economic institutions for allocating apartments. Each method will result in different people getting apartments or in different prices being charged for apartments. We might well ask which economic institution is best. But first we have to define "best." What criteria might we use to compare these ways of allocating apartments?

One thing we can do is to look at the economic positions of the people involved. It is pretty obvious that the owners of the apartments end up with the most money if they can act as discriminating monopolists: this would generate the most revenues for the apartment owner(s). Similarly the rent-control solution is probably the worst situation for the apartment owners.

What about the renters? They are probably worse off on average in the case of a discriminating monopolist—most of them would be paying a higher price than they would under the other ways of allocating apartments.

Are the consumers better off in the case of rent control? Some of them are: the consumers *who end up getting the apartments* are better off than they would be under the market solution. But the ones who didn't get the apartments are *worse off* than they would be under the market solution.

What we need here is a way to look at the economic position of all the parties involved—all the renters *and* all the landlords. How can we examine the desirability of different ways to allocate apartments, taking everybody into account? What can be used as a criterion for a "good" way to allocate apartments taking into account *all* of the parties involved?

1.9 Pareto Efficiency

One useful criterion for comparing the outcomes of different economic institutions is a concept known as Pareto efficiency or economic efficiency.[1] We start with the following definition: if we can find a way to make some people better off without making anybody else worse off, we have a **Pareto improvement**. If an allocation allows for a Pareto improvement, it is called **Pareto inefficient**; if an allocation is such that no Pareto improvements are possible, it is called **Pareto efficient**.

A Pareto inefficient allocation has the undesirable feature that there is some way to make somebody better off without hurting anyone else. There may be other positive things about the allocation, but the fact that it is Pareto inefficient is certainly one strike against it. If there is a way to make someone better off without hurting anyone else, why not do it?

The idea of Pareto efficiency is an important one in economics and we will examine it in some detail later on. It has many subtle implications that we will have to investigate more slowly, but we can get an inkling of what is involved even now.

Here is a useful way to think about the idea of Pareto efficiency. Suppose that we assigned the renters to the inner- and outer-ring apartments randomly, but then allowed them to sublet their apartments to each other. Some people who really wanted to live close in might, through bad luck, end up with an outer-ring apartment. But then they could sublet an inner-ring apartment from someone who was assigned to such an apartment but who didn't value it as highly as the other person. If individuals were assigned randomly to apartments, there would generally be some who would want to trade apartments, if they were sufficiently compensated for doing so.

For example, suppose that person A is assigned an apartment in the inner ring that he feels is worth $200, and that there is some person B in the outer ring who would be willing to pay $300 for A's apartment. Then there is a

[1] Pareto efficiency is named after the nineteenth-century economist and sociologist Vilfredo Pareto (1848–1923) who was one of the first to examine the implications of this idea.

"gain from trade" if these two agents swap apartments and arrange a side payment from B to A of some amount of money between \$200 and \$300. The exact amount of the transaction isn't important. What is important is that the people who are willing to pay the most for the apartments get them—otherwise, there would be an incentive for someone who attached a low value to an inner-ring apartment to make a trade with someone who placed a high value on an inner-ring apartment.

Suppose that we think of all voluntary trades as being carried out so that all gains from trade are exhausted. The resulting allocation must be Pareto efficient. If not, there would be some trade that would make two people better off without hurting anyone else—but this would contradict the assumption that all voluntary trades had been carried out. An allocation in which all voluntary trades have been carried out is a Pareto efficient allocation.

1.10 Comparing Ways to Allocate Apartments

The trading process we've described above is so general that you wouldn't think that anything much could be said about its outcome. But there is one very interesting point that can be made. Let us ask who will end up with apartments in an allocation where all of the gains from trade have been exhausted.

To see the answer, just note that anyone who has an apartment in the inner ring must have a higher reservation price than anyone who has an apartment in the outer ring—otherwise, they could make a trade and make both people better off. Thus if there are S apartments to be rented, then the S people with the highest reservation prices end up getting apartments in the inner ring. This allocation is Pareto efficient—anything else is not, since any other assignment of apartments to people would allow for some trade that would make at least two of the people better off without hurting anyone else.

Let us try to apply this criterion of Pareto efficiency to the outcomes of the various resource allocation devices mentioned above. Let's start with the market mechanism. It is easy to see that the market mechanism assigns the people with the S highest reservation prices to the inner ring—namely, those people who are willing to pay more than the equilibrium price, p^*, for their apartments. Thus there are no further gains from trade to be had once the apartments have been rented in a competitive market. The outcome of the competitive market is Pareto efficient.

What about the discriminating monopolist? Is that arrangement Pareto efficient? To answer this question, simply observe that the discriminating monopolist assigns apartments to exactly the same people who receive apartments in the competitive market. Under each system everyone who is willing to pay more than p^* for an apartment gets an apartment. Thus the discriminating monopolist generates a Pareto efficient outcome as well.

Although both the competitive market and the discriminating monopolist generate Pareto efficient outcomes in the sense that there will be no further trades desired, they can result in quite different distributions of income. Certainly the consumers are much worse off under the discriminating monopolist than under the competitive market, and the landlord(s) are much better off. In general, Pareto efficiency doesn't have much to say about distribution of the gains from trade. It is only concerned with the *efficiency* of the trade: whether all of the possible trades have been made.

What about the ordinary monopolist who is constrained to charge just one price? It turns out that this situation is not Pareto efficient. All we have to do to verify this is to note that, since all the apartments will not in general be rented by the monopolist, he can increase his profits by renting an apartment to someone who doesn't have one at *any* positive price. There is some price at which both the monopolist and the renter must be better off. As long as the monopolist doesn't change the price that anybody else pays, the other renters are just as well off as they were before. Thus we have found a **Pareto improvement**—a way to make two parties better off without making anyone else worse off.

The final case is that of rent control. This also turns out not to be Pareto efficient. The argument here rests on the fact that an arbitrary assignment of renters to apartments will generally involve someone living in the inner ring (say Mr. In) who is willing to pay less for an apartment than someone living in the outer ring (say Ms. Out). Suppose that Mr. In's reservation price is $300 and Ms. Out's reservation price is $500.

We need to find a Pareto improvement—a way to make Mr. In and Ms. Out better off without hurting anyone else. But there is an easy way to do this: just let Mr. In sublet his apartment to Ms. Out. It is worth $500 to Ms. Out to live close to the university, but it is only worth $300 to Mr. In. If Ms. Out pays Mr. In $400, say, and trades apartments, they will both be better off: Ms. Out will get an apartment that she values at more than $400, and Mr. In will get $400 that he values more than an inner-ring apartment.

This example shows that the rent-controlled market will generally not result in a Pareto efficient allocation, since there will still be some trades that could be carried out after the market has operated. As long as some people get inner-ring apartments who value them less highly than people who don't get them, there will be gains to be had from trade.

1.11 Equilibrium in the Long Run

We have analyzed the equilibrium pricing of apartments in the **short run**—when there is a fixed supply of apartments. But in the **long run** the supply of apartments can change. Just as the demand curve measures the number of apartments that will be demanded at different prices, the supply curve measures the number of apartments that will be supplied at different prices.

The final determination of the market price for apartments will depend on the interaction of supply and demand.

And what is it that determines the supply behavior? In general, the number of new apartments that will be supplied by the private market will depend on how profitable it is to provide apartments, which depends, in part, on the price that landlords can charge for apartments. In order to analyze the behavior of the apartment market in the long run, we have to examine the behavior of suppliers as well as demanders, a task we will eventually undertake.

When supply is variable, we can ask questions not only about who gets the apartments, but about how many will be provided by various types of market institutions. Will a monopolist supply more or fewer apartments than a competitive market? Will rent control increase or decrease the equilibrium number of apartments? Which institutions will provide a Pareto efficient number of apartments? In order to answer these and similar questions we must develop more systematic and powerful tools for economic analysis.

Summary

1. Economics proceeds by making models of social phenomena, which are simplified representations of reality.

2. In this task, economists are guided by the optimization principle, which states that people typically try to choose what's best for them, and by the equilibrium principle, which says that prices will adjust until demand and supply are equal.

3. The demand curve measures how much people wish to demand at each price, and the supply curve measures how much people wish to supply at each price. An equilibrium price is one where the amount demanded equals the amount supplied.

4. The study of how the equilibrium price and quantity change when the underlying conditions change is known as comparative statics.

5. An economic situation is Pareto efficient if there is no way to make some group of people better off without making some other group of people worse off. The concept of Pareto efficiency can be used to evaluate different ways of allocating resources.

REVIEW QUESTIONS

1. Suppose that there were 25 people who had a reservation price of \$500, and the 26th person had a reservation price of \$200. What would the demand curve look like?

2. In the above example, what would the equilibrium price be if there were 24 apartments to rent? What if there were 26 apartments to rent? What if there were 25 apartments to rent?

3. If people have different reservation prices, why does the market demand curve slope down?

4. In the text we assumed that the condominium purchasers came from the inner-ring people—people who were already renting apartments. What would happen to the price of inner-ring apartments if all of the condominium purchasers were outer-ring people—the people who were not currently renting apartments in the inner ring?

5. Suppose now that the condominium purchasers were all inner-ring people, but that each condominium was constructed from two apartments. What would happen to the price of apartments?

6. What do you suppose the effect of a tax would be on the number of apartments that would be built in the long run?

7. Suppose the demand curve is $D(p) = 100 - 2p$. What price would the monopolist set if he had 60 apartments? How many would he rent? What price would he set if he had 40 apartments? How many would he rent?

8. If our model of rent control allowed for unrestricted subletting, who would end up getting apartments in the inner circle? Would the outcome be Pareto efficient?

CHAPTER 2

BUDGET CONSTRAINT

The economic theory of the consumer is very simple: economists assume that consumers choose the best bundle of goods they can afford. To give content to this theory, we have to describe more precisely what we mean by "best" and what we mean by "can afford." In this chapter we will examine how to describe what a consumer can afford; the next chapter will focus on the concept of how the consumer determines what is best. We will then be able to undertake a detailed study of the implications of this simple model of consumer behavior.

2.1 The Budget Constraint

We begin by examining the concept of the **budget constraint**. Suppose that there is some set of goods from which the consumer can choose. In real life there are many goods to consume, but for our purposes it is convenient to consider only the case of two goods, since we can then depict the consumer's choice behavior graphically.

We will indicate the consumer's **consumption bundle** by (x_1, x_2). This is simply a list of two numbers that tells us how much the consumer is choosing to consume of good 1, x_1, and how much the consumer is choosing to consume of good 2, x_2. Sometimes it is convenient to denote the consumer's bundle by a single symbol like X, where X is simply an abbreviation for the list of two numbers (x_1, x_2).

We suppose that we can observe the prices of the two goods, (p_1, p_2), and the amount of money the consumer has to spend, m. Then the budget constraint of the consumer can be written as

$$p_1 x_1 + p_2 x_2 \leq m. \tag{2.1}$$

Here $p_1 x_1$ is the amount of money the consumer is spending on good 1, and $p_2 x_2$ is the amount of money the consumer is spending on good 2. The budget constraint of the consumer requires that the amount of money spent on the two goods be no more than the total amount the consumer has to spend. The consumer's *affordable* consumption bundles are those that don't cost any more than m. We call this set of affordable consumption bundles at prices (p_1, p_2) and income m the **budget set** of the consumer.

2.2 Two Goods Are Often Enough

The two-good assumption is more general than you might think at first, since we can often interpret one of the goods as representing everything else the consumer might want to consume.

For example, if we are interested in studying a consumer's demand for milk, we might let x_1 measure his or her consumption of milk in quarts per month. We can then let x_2 stand for everything else the consumer might want to consume.

When we adopt this interpretation, it is convenient to think of good 2 as being the dollars that the consumer can use to spend on other goods. Under this interpretation the price of good 2 will automatically be 1, since the price of one dollar is one dollar. Thus the budget constraint will take the form

$$p_1 x_1 + x_2 \leq m. \tag{2.2}$$

This expression simply says that the amount of money spent on good 1, $p_1 x_1$, plus the amount of money spent on all other goods, x_2, must be no more than the total amount of money the consumer has to spend, m.

We say that good 2 represents a **composite good** that stands for everything else that the consumer might want to consume other than good 1. Such a composite good is invariably measured in dollars to be spent on goods other than good 1. As far as the algebraic form of the budget constraint is concerned, equation (2.2) is just a special case of the formula given in equation (2.1), with $p_2 = 1$, so everything that we have to say about the budget constraint in general will hold under the composite-good interpretation.

2.3 Properties of the Budget Set

The **budget line** is the set of bundles that cost exactly m:

$$p_1 x_1 + p_2 x_2 = m. \tag{2.3}$$

These are the bundles of goods that just exhaust the consumer's income.

The budget set is depicted in Figure 2.1. The heavy line is the budget line—the bundles that cost exactly m—and the bundles below this line are those that cost strictly less than m.

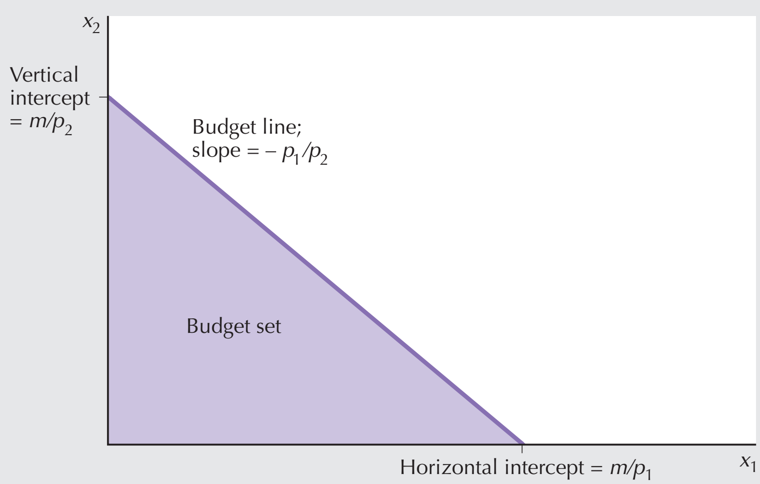

Figure 2.1 **The budget set.** The budget set consists of all bundles that are affordable at the given prices and income.

We can rearrange the budget line in equation (2.3) to give us the formula

$$x_2 = \frac{m}{p_2} - \frac{p_1}{p_2} x_1. \tag{2.4}$$

This is the formula for a straight line with a vertical intercept of m/p_2 and a slope of $-p_1/p_2$. The formula tells us how many units of good 2 the consumer needs to consume in order to just satisfy the budget constraint if she is consuming x_1 units of good 1.

Here is an easy way to draw a budget line given prices (p_1, p_2) and income m. Just ask yourself how much of good 2 the consumer could buy if she spent all of her money on good 2. The answer is, of course, m/p_2. Then ask how much of good 1 the consumer could buy if she spent all of her money on good 1. The answer is m/p_1. Thus the horizontal and vertical intercepts measure how much the consumer could get if she spent all of her money on goods 1 and 2, respectively. In order to depict the budget line just plot these two points on the appropriate axes of the graph and connect them with a straight line.

The slope of the budget line has a nice economic interpretation. It measures the rate at which the market is willing to "substitute" good 1 for good 2. Suppose for example that the consumer is going to increase her consumption of good 1 by dx_1. How much will her consumption of good 2 have to change in order to satisfy her budget constraint? Let us use dx_2 to indicate her change in the consumption of good 2.

Now note that if she satisfies her budget constraint before and after making the change she must satisfy

$$p_1 x_1 + p_2 x_2 = m$$

and

$$p_1(x_1 + dx_1) + p_2(x_2 + dx_2) = m.$$

Subtracting the first equation from the second gives

$$p_1 dx_1 + p_2 dx_2 = 0.$$

This says that the total value of the change in her consumption must be zero. Solving for dx_2/dx_1, the rate at which good 2 can be substituted for good 1 while still satisfying the budget constraint, gives

$$\frac{dx_2}{dx_1} = -\frac{p_1}{p_2}.$$

This is just the slope of the budget line. The negative sign is there since dx_1 and dx_2 must always have opposite signs. If you consume more of good 1, you have to consume less of good 2 and vice versa if you continue to satisfy the budget constraint. Alternatively, we could have taken the implicit derivative of both sides of the budget constraint with respect to x_1 and obtained the same result.

Economists sometimes say that the slope of the budget line measures the **opportunity cost** of consuming good 1. In order to consume more of good 1 you have to give up some consumption of good 2. Giving up the opportunity to consume good 2 is the true economic cost of more good 1 consumption; and that cost is measured by the slope of the budget line.

2.4 How the Budget Line Changes

When prices and incomes change, the set of goods that a consumer can afford changes as well. How do these changes affect the budget set?

Let us first consider changes in income. It is easy to see from equation (2.4) that an increase in income will increase the vertical intercept and not affect the slope of the line. Thus an increase in income will result in a *parallel shift outward* of the budget line as in Figure 2.2. Similarly, a decrease in income will cause a parallel shift inward.

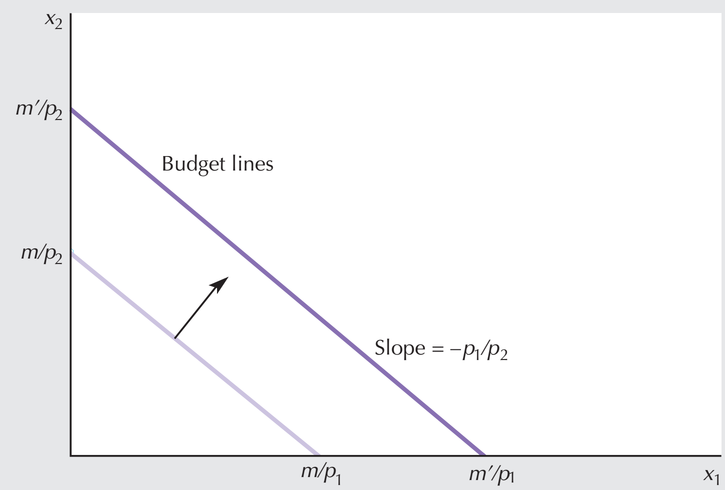

Figure 2.2 **Increasing income.** An increase in income causes a parallel shift outward of the budget line.

What about changes in prices? Let us first consider increasing price 1 while holding price 2 and income fixed. According to equation (2.4), increasing p_1 will not change the vertical intercept, but it will make the budget line steeper since p_1/p_2 will become larger.

Another way to see how the budget line changes is to use the trick described earlier for drawing the budget line. If you are spending all of your money on good 2, then increasing the price of good 1 doesn't change the maximum amount of good 2 you could buy—thus the vertical intercept of the budget line doesn't change. But if you are spending all of your money on good 1, and good 1 becomes more expensive, then your consumption of good 1 must decrease. Thus the horizontal intercept of the budget line must shift inward, resulting in the tilt depicted in Figure 2.3.

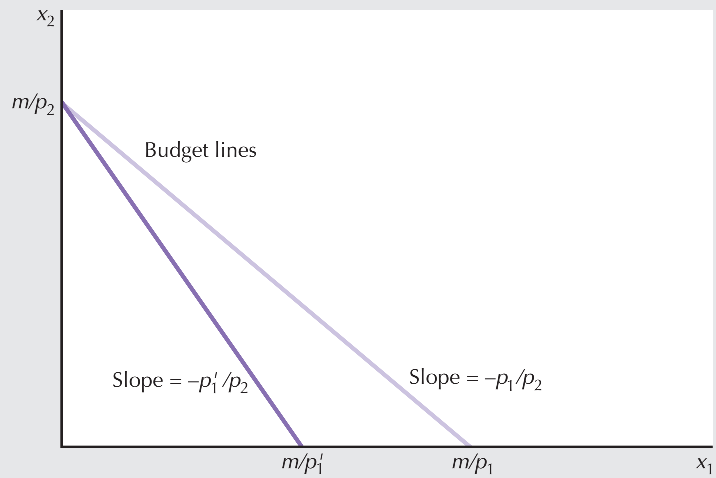

Increasing price. If good 1 becomes more expensive, the budget line becomes steeper.

Figure 2.3

What happens to the budget line when we change the prices of good 1 and good 2 at the same time? Suppose for example that we double the prices of both goods 1 and 2. In this case both the horizontal and vertical intercepts shift inward by a factor of one-half, and therefore the budget line shifts inward by one-half as well. Multiplying both prices by two is just like dividing income by 2.

We can also see this algebraically. Suppose our original budget line is

$$p_1x_1 + p_2x_2 = m.$$

Now suppose that both prices become t times as large. Multiplying both prices by t yields

$$tp_1x_1 + tp_2x_2 = m.$$

But this equation is the same as

$$p_1x_1 + p_2x_2 = \frac{m}{t}.$$

Thus multiplying both prices by a constant amount t is just like dividing income by the same constant t. It follows that if we multiply both prices by t *and* we multiply income by t, then the budget line won't change at all.

We can also consider price and income changes together. What happens if both prices go up and income goes down? Think about what happens to the horizontal and vertical intercepts. If m decreases and p_1 and p_2 both increase, then the intercepts m/p_1 and m/p_2 must both decrease. This means that the budget line will shift inward. What about the slope of the budget line? If price 2 increases more than price 1, so that $-p_1/p_2$ decreases (in absolute value), then the budget line will be flatter; if price 2 increases less than price 1, the budget line will be steeper.

2.5 The Numeraire

The budget line is defined by two prices and one income, but one of these variables is redundant. We could peg one of the prices, or the income, to some fixed value, and adjust the other variables so as to describe exactly the same budget set. Thus the budget line

$$p_1x_1 + p_2x_2 = m$$

is exactly the same budget line as

$$\frac{p_1}{p_2}x_1 + x_2 = \frac{m}{p_2}$$

or

$$\frac{p_1}{m}x_1 + \frac{p_2}{m}x_2 = 1,$$

since the first budget line results from dividing everything by p_2, and the second budget line results from dividing everything by m. In the first case, we have pegged $p_2 = 1$, and in the second case, we have pegged $m = 1$. Pegging the price of one of the goods or income to 1 and adjusting the other price and income appropriately doesn't change the budget set at all.

When we set one of the prices to 1, as we did above, we often refer to that price as the **numeraire** price. The numeraire price is the price relative to which we are measuring the other price and income. It will occasionally be convenient to think of one of the goods as being a numeraire good, since there will then be one less price to worry about.

2.6 Taxes, Subsidies, and Rationing

Economic policy often uses tools that affect a consumer's budget constraint, such as taxes. For example, if the government imposes a **quantity tax**, this means that the consumer has to pay a certain amount to the government for each unit of the good he purchases. In the U.S., for example, we pay about 15 cents a gallon as a federal gasoline tax.

How does a quantity tax affect the budget line of a consumer? From the viewpoint of the consumer the tax is just like a higher price. Thus a quantity tax of t dollars per unit of good 1 simply changes the price of good 1 from p_1 to $p_1 + t$. As we've seen above, this implies that the budget line must get steeper.

Another kind of tax is a **value** tax. As the name implies this is a tax on the value—the price—of a good, rather than the quantity purchased of a good. A value tax is usually expressed in percentage terms. Most states in the U.S. have sales taxes. If the sales tax is 6 percent, then a good that is priced at \$1 will actually sell for \$1.06. (Value taxes are also known as **ad valorem** taxes.)

If good 1 has a price of p_1 but is subject to a sales tax at rate τ, then the actual price facing the consumer is $(1 + \tau)p_1$.[1] The consumer has to pay p_1 to the supplier and τp_1 to the government for each unit of the good so the total cost of the good to the consumer is $(1 + \tau)p_1$.

A **subsidy** is the opposite of a tax. In the case of a **quantity subsidy**, the government *gives* an amount to the consumer that depends on the amount of the good purchased. If, for example, the consumption of milk were subsidized, the government would pay some amount of money to each consumer of milk depending on the amount that consumer purchased. If the subsidy is s dollars per unit of consumption of good 1, then from the viewpoint of the consumer, the price of good 1 would be $p_1 - s$. This would therefore make the budget line flatter.

Similarly an ad valorem subsidy is a subsidy based on the price of the good being subsidized. If the government gives you back \$1 for every \$2 you donate to charity, then your donations to charity are being subsidized at a rate of 50 percent. In general, if the price of good 1 is p_1 and good 1 is subject to an ad valorem subsidy at rate σ, then the actual price of good 1 facing the consumer is $(1 - \sigma)p_1$.[2]

You can see that taxes and subsidies affect prices in exactly the same way except for the algebraic sign: a tax increases the price to the consumer, and a subsidy decreases it.

Another kind of tax or subsidy that the government might use is a **lump-sum** tax or subsidy. In the case of a tax, this means that the government takes away some fixed amount of money, regardless of the individual's behavior. Thus a lump-sum tax means that the budget line of a consumer will shift inward because his money income has been reduced. Similarly, a lump-sum subsidy means that the budget line will shift outward. Quantity

[1] The Greek letter τ, tau, rhymes with "wow" in mathematical discourse, though modern Greeks pronounce it "taf."

[2] The Greek letter σ is pronounced "sig-ma."

taxes and value taxes tilt the budget line one way or the other depending on which good is being taxed, but a lump-sum tax shifts the budget line inward.

Governments also sometimes impose *rationing* constraints. This means that the level of consumption of some good is fixed to be no larger than some amount. For example, during World War II the U.S. government rationed certain foods like butter and meat.

Suppose, for example, that good 1 were rationed so that no more than $\overline{x}_1$ could be consumed by a given consumer. Then the budget set of the consumer would look like that depicted in Figure 2.4: it would be the old budget set with a piece lopped off. The lopped-off piece consists of all the consumption bundles that are affordable but have $x_1 > \overline{x}_1$.

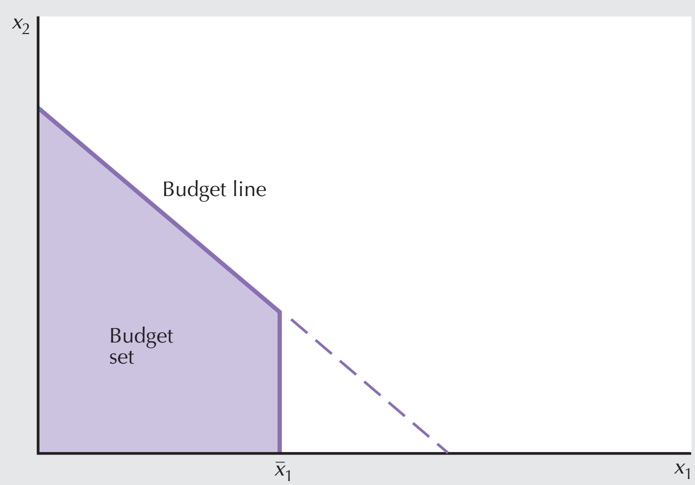

Figure 2.4 **Budget set with rationing**. If good 1 is rationed, the section of the budget set beyond the rationed quantity will be lopped off.

Sometimes taxes, subsidies, and rationing are combined. For example, we could consider a situation where a consumer could consume good 1 at a price of p_1 up to some level $\overline{x}_1$, and then had to pay a tax t on all consumption in excess of $\overline{x}_1$. The budget set for this consumer is depicted in Figure 2.5. Here the budget line has a slope of $-p_1/p_2$ to the left of $\overline{x}_1$, and a slope of $-(p_1 + t)/p_2$ to the right of $\overline{x}_1$.

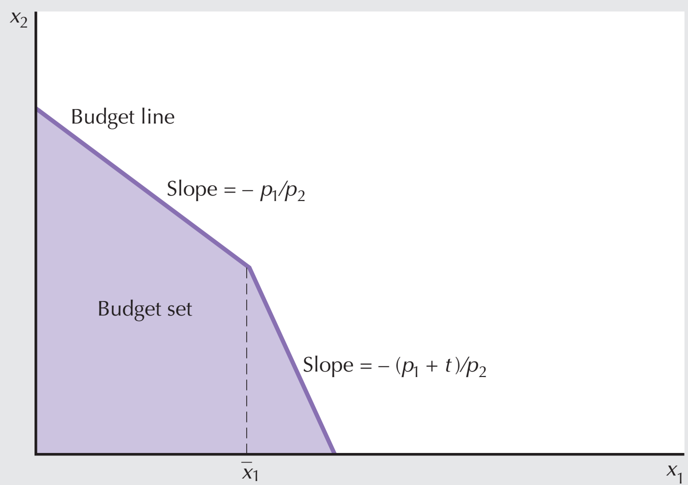

Figure 2.5 **Taxing consumption greater than $\overline{x}_1$.** In this budget set the consumer must pay a tax only on the consumption of good 1 that is in excess of $\overline{x}_1$, so the budget line becomes steeper to the right of $\overline{x}_1$.

EXAMPLE: The Food Stamp Program

Since the Food Stamp Act of 1964 the U.S. federal government has provided a subsidy on food for poor people. The details of this program have been adjusted several times. Here we will describe the economic effects of one of these adjustments.

Before 1979, households who met certain eligibility requirements were allowed to purchase food stamps, which could then be used to purchase food at retail outlets. In January 1975, for example, a family of four could receive a maximum monthly allotment of \$153 in food coupons by participating in the program.

The price of these coupons to the household depended on the household income. A family of four with an adjusted monthly income of \$300 paid \$83 for the full monthly allotment of food stamps. If a family of four had a monthly income of \$100, the cost for the full monthly allotment would have been \$25.[3]

The pre-1979 Food Stamp program was an ad valorem subsidy on food. The rate at which food was subsidized depended on the household income.

[3] These figures are taken from Kenneth Clarkson, *Food Stamps and Nutrition*, American Enterprise Institute, 1975.

The family of four that was charged \$83 for their allotment paid \$1 to receive \$1.84 worth of food (1.84 equals 153 divided by 83). Similarly, the household that paid \$25 was paying \$1 to receive \$6.12 worth of food (6.12 equals 153 divided by 25).

The way that the Food Stamp program affected the budget set of a household is depicted in Figure 2.6A. Here we have measured the amount of money spent on food on the horizontal axis and expenditures on all other goods on the vertical axis. Since we are measuring each good in terms of the money spent on it, the "price" of each good is automatically 1, and the budget line will therefore have a slope of -1.

If the household is allowed to buy \$153 of food stamps for \$25, then this represents roughly an 84 percent ($= 1 - 25/153$) subsidy of food purchases, so the budget line will have a slope of roughly $-.16$ ($= 25/153$) until the household has spent \$153 on food. Each dollar that the household spends on food up to \$153 would reduce its consumption of other goods by about 16 cents. After the household spends \$153 on food, the budget line facing it would again have a slope of -1.

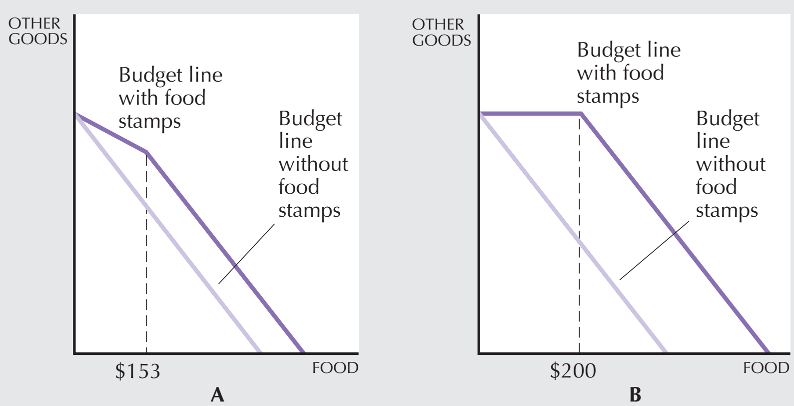

Figure 2.6 **Food stamps.** How the budget line is affected by the Food Stamp program. Part A shows the pre-1979 program and part B the post-1979 program.

These effects lead to the kind of "kink" depicted in Figure 2.6. Households with higher incomes had to pay more for their allotment of food stamps. Thus the slope of the budget line would become steeper as household income increased.

In 1979 the Food Stamp program was modified. Instead of requiring that households purchase food stamps, they are now simply given to qualified households. Figure 2.6B shows how this affects the budget set.

Suppose that a household now receives a grant of $200 of food stamps a month. Then this means that the household can consume $200 more food per month, regardless of how much it is spending on other goods, which implies that the budget line will shift to the right by $200. The slope will not change: $1 less spent on food would mean $1 more to spend on other things. But since the household cannot legally sell food stamps, the maximum amount that it can spend on other goods does not change. The Food Stamp program is effectively a lump-sum subsidy, except for the fact that the food stamps can't be sold.

2.7 Budget Line Changes

In the next chapter we will analyze how the consumer chooses an optimal consumption bundle from his or her budget set. But we can already state some observations here that follow from what we have learned about the movements of the budget line.

First, we can observe that since the budget set doesn't change when we multiply all prices and income by a positive number, the optimal choice of the consumer from the budget set can't change either. Without even analyzing the choice process itself, we have derived an important conclusion: a perfectly balanced inflation—one in which all prices and all incomes rise at the same rate—doesn't change anybody's budget set, and thus cannot change anybody's optimal choice.

Second, we can make some statements about how well-off the consumer can be at different prices and incomes. Suppose that the consumer's income increases and all prices remain the same. We know that this represents a parallel shift outward of the budget line. Thus every bundle the consumer was consuming at the lower income is also a possible choice at the higher income. But then the consumer must be at least as well-off at the higher income as at the lower income—since he or she has the same choices available as before plus some more. Similarly, if one price declines and all others stay the same, the consumer must be at least as well-off. This simple observation will be of considerable use later on.

Summary

1. The budget set consists of all bundles of goods that the consumer can afford at given prices and income. We will typically assume that there are only two goods, but this assumption is more general than it seems.

2. The budget line is written as $p_1x_1+p_2x_2=m$. It has a slope of $-p_1/p_2$, a vertical intercept of m/p_2, and a horizontal intercept of m/p_1.

3. Increasing income shifts the budget line outward. Increasing the price of good 1 makes the budget line steeper. Increasing the price of good 2 makes the budget line flatter.

4. Taxes, subsidies, and rationing change the slope and position of the budget line by changing the prices paid by the consumer.

REVIEW QUESTIONS

1. Originally the consumer faces the budget line $p_1x_1 + p_2x_2 = m$. Then the price of good 1 doubles, the price of good 2 becomes 8 times larger, and income becomes 4 times larger. Write down an equation for the new budget line in terms of the original prices and income.

2. What happens to the budget line if the price of good 2 increases, but the price of good 1 and income remain constant?

3. If the price of good 1 doubles and the price of good 2 triples, does the budget line become flatter or steeper?

4. What is the definition of a numeraire good?

5. Suppose that the government puts a tax of 15 cents a gallon on gasoline and then later decides to put a subsidy on gasoline at a rate of 7 cents a gallon. What net tax is this combination equivalent to?

6. Suppose that a budget equation is given by $p_1x_1 + p_2x_2 = m$. The government decides to impose a lump-sum tax of u, a quantity tax on good 1 of t, and a quantity subsidy on good 2 of s. What is the formula for the new budget line?

7. If the income of the consumer increases and one of the prices decreases at the same time, will the consumer necessarily be at least as well-off?

CHAPTER 3

PREFERENCES

We saw in Chapter 2 that the economic model of consumer behavior is very simple: people choose the best things they can afford. The last chapter was devoted to clarifying the meaning of "can afford," and this chapter will be devoted to clarifying the economic concept of "best things."

We call the objects of consumer choice **consumption bundles**. This is a complete list of the goods and services that are involved in the choice problem that we are investigating. The word "complete" deserves emphasis: when you analyze a consumer's choice problem, make sure that you include all of the appropriate goods in the definition of the consumption bundle.

If we are analyzing consumer choice at the broadest level, we would want not only a complete list of the goods that a consumer might consume, but also a description of when, where, and under what circumstances they would become available. After all, people care about how much food they will have tomorrow as well as how much food they have today. A raft in the middle of the Atlantic Ocean is very different from a raft in the middle of the Sahara Desert. And an umbrella when it is raining is quite a different good from an umbrella on a sunny day. It is often useful to think of the "same" good available in different locations or circumstances as a different good, since the consumer may value the good differently in those situations.

However, when we limit our attention to a simple choice problem, the relevant goods are usually pretty obvious. We'll often adopt the idea described earlier of using just two goods and calling one of them "all other goods" so that we can focus on the tradeoff between one good and everything else. In this way we can consider consumption choices involving many goods and still use two-dimensional diagrams.

So let us take our consumption bundle to consist of two goods, and let x_1 denote the amount of one good and x_2 the amount of the other. The complete consumption bundle is therefore denoted by (x_1, x_2). As noted before, we will occasionally abbreviate this consumption bundle by X.

3.1 Consumer Preferences

We will suppose that given any two consumption bundles, (x_1, x_2) and (y_1, y_2), the consumer can rank them as to their desirability. That is, the consumer can determine that one of the consumption bundles is strictly better than the other, or decide that she is indifferent between the two bundles.

We will use the symbol $\succ$ to mean that one bundle is **strictly preferred** to another, so that $(x_1, x_2) \succ (y_1, y_2)$ should be interpreted as saying that the consumer **strictly prefers** (x_1, x_2) to (y_1, y_2), in the sense that she definitely wants the x-bundle rather than the y-bundle. This preference relation is meant to be an operational notion. If the consumer prefers one bundle to another, it means that he or she would choose one over the other, given the opportunity. Thus the idea of preference is based on the consumer's *behavior.* In order to tell whether one bundle is preferred to another, we see how the consumer behaves in choice situations involving the two bundles. If she always chooses (x_1, x_2) when (y_1, y_2) is available, then it is natural to say that this consumer prefers (x_1, x_2) to (y_1, y_2).

If the consumer is **indifferent** between two bundles of goods, we use the symbol $\sim$ and write $(x_1, x_2) \sim (y_1, y_2)$. Indifference means that the consumer would be just as satisfied, according to her own preferences, consuming the bundle (x_1, x_2) as she would be consuming the other bundle, (y_1, y_2).

If the consumer prefers or is indifferent between the two bundles we say that she **weakly prefers** (x_1, x_2) to (y_1, y_2) and write $(x_1, x_2) \succeq (y_1, y_2)$.

These relations of strict preference, weak preference, and indifference are not independent concepts; the relations are themselves related! For example, if $(x_1, x_2) \succeq (y_1, y_2)$ and $(y_1, y_2) \succeq (x_1, x_2)$ we can conclude that $(x_1, x_2) \sim (y_1, y_2)$. That is, if the consumer thinks that (x_1, x_2) is at least as good as (y_1, y_2) *and* that (y_1, y_2) is at least as good as (x_1, x_2), then the consumer must be indifferent between the two bundles of goods.

Similarly, if $(x_1, x_2) \succeq (y_1, y_2)$ but we know that it is *not* the case that $(x_1, x_2) \sim (y_1, y_2)$, we can conclude that we must have $(x_1, x_2) \succ (y_1, y_2)$. This just says that if the consumer thinks that (x_1, x_2) is at least as good as (y_1, y_2), and she is not indifferent between the two bundles, then it must be that she thinks that (x_1, x_2) is strictly better than (y_1, y_2).

3.2 Assumptions about Preferences

Economists usually make some assumptions about the "consistency" of consumers' preferences. For example, it seems unreasonable—not to say contradictory—to have a situation where $(x_1, x_2) \succ (y_1, y_2)$ and, at the same time, $(y_1, y_2) \succ (x_1, x_2)$. For this would mean that the consumer strictly prefers the x-bundle to the y-bundle ... and vice versa.

So we usually make some assumptions about how the preference relations work. Some of the assumptions about preferences are so fundamental that we can refer to them as "axioms" of consumer theory. Here are three such axioms about consumer preference.

Complete. We assume that any two bundles can be compared. That is, given any x-bundle and any y-bundle, we assume that $(x_1, x_2) \succeq (y_1, y_2)$, or $(y_1, y_2) \succeq (x_1, x_2)$, or both, in which case the consumer is indifferent between the two bundles.

Reflexive. We assume that any bundle is at least as good as itself: $(x_1, x_2) \succeq (x_1, x_2)$.

Transitive. If $(x_1, x_2) \succeq (y_1, y_2)$ and $(y_1, y_2) \succeq (z_1, z_2)$, then we assume that $(x_1, x_2) \succeq (z_1, z_2)$. In other words, if the consumer thinks that X is at least as good as Y and that Y is at least as good as Z, then the consumer thinks that X is at least as good as Z.

The first axiom, completeness, is hardly objectionable, at least for the kinds of choices economists generally examine. To say that any two bundles can be compared is simply to say that the consumer is able to make a choice between any two given bundles. One might imagine extreme situations involving life or death choices where ranking the alternatives might be difficult, or even impossible, but these choices are, for the most part, outside the domain of economic analysis.

The second axiom, reflexivity, is trivial. Any bundle is certainly at least as good as an identical bundle. Parents of small children may occasionally observe behavior that violates this assumption, but it seems plausible for most adult behavior.

The third axiom, transitivity, is more problematic. It isn't clear that transitivity of preferences is *necessarily* a property that preferences would have to have. The assumption that preferences are transitive doesn't seem compelling on grounds of pure logic alone. In fact it's not. Transitivity is a hypothesis about people's choice behavior, not a statement of pure logic. Whether it is a basic fact of logic or not isn't the point: it is whether or not it is a reasonably accurate description of how people behave that matters.

What would you think about a person who said that he preferred a bundle X to Y, and preferred Y to Z, but then also said that he preferred Z to X? This would certainly be taken as evidence of peculiar behavior.

More importantly, how would this consumer behave if faced with choices among the three bundles X, Y, and Z? If we asked him to choose his most preferred bundle, he would have quite a problem, for whatever bundle he chose, there would always be one that was preferred to it. If we are to have a theory where people are making "best" choices, preferences must satisfy the transitivity axiom or something very much like it. If preferences were not transitive there could well be a set of bundles for which there is no best choice.

3.3 Indifference Curves

It turns out that the whole theory of consumer choice can be formulated in terms of preferences that satisfy the three axioms described above, plus a few more technical assumptions. However, we will find it convenient to describe preferences graphically by using a construction known as **indifference curves**.

Consider Figure 3.1 where we have illustrated two axes representing a consumer's consumption of goods 1 and 2. Let us pick a certain consumption bundle (x_1, x_2) and shade in all of the consumption bundles that are weakly preferred to (x_1, x_2). This is called the **weakly preferred set**. The bundles on the boundary of this set—the bundles for which the consumer is just indifferent to (x_1, x_2)—form the **indifference curve**.

We can draw an indifference curve through any consumption bundle we want. The indifference curve through a consumption bundle consists of all bundles of goods that leave the consumer indifferent to the given bundle.

One problem with using indifference curves to describe preferences is that they only show you the bundles that the consumer perceives as being indifferent to each other—they don't show you which bundles are better and which bundles are worse. It is sometimes useful to draw small arrows on the indifference curves to indicate the direction of the preferred bundles. We won't do this in every case, but we will do it in a few of the examples where confusion might arise.

If we make no further assumptions about preferences, indifference curves can take very peculiar shapes indeed. But even at this level of generality, we can state an important principle about indifference curves: *indifference curves representing distinct levels of preference cannot cross.* That is, the situation depicted in Figure 3.2 cannot occur.

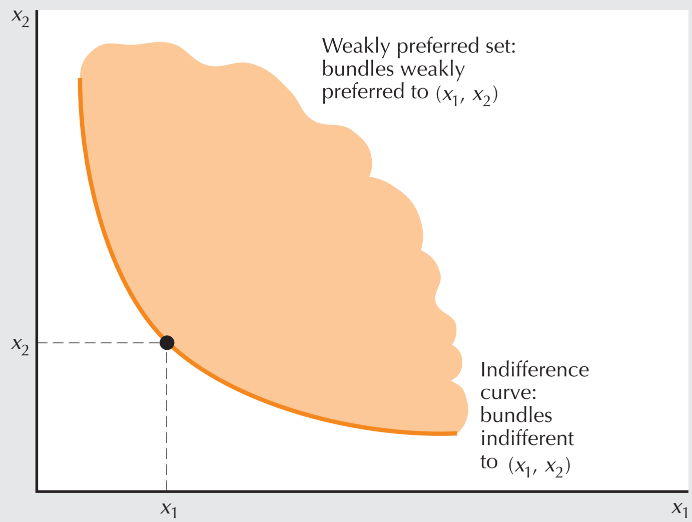

Weakly preferred set. The shaded area consists of all bundles that are at least as good as the bundle (x_1, x_2).

Figure 3.1

In order to prove this, let us choose three bundles of goods, X, Y, and Z, such that X lies only on one indifference curve, Y lies only on the other indifference curve, and Z lies at the intersection of the indifference curves. By assumption the indifference curves represent distinct levels of preference, so one of the bundles, say X, is strictly preferred to the other bundle, Y. We know that $X \sim Z$ and $Z \sim Y$, and the axiom of transitivity therefore implies that $X \sim Y$. But this contradicts the assumption that $X \succ Y$. This contradiction establishes the result—indifference curves representing distinct levels of preference cannot cross.

What other properties do indifference curves have? In the abstract, the answer is: not many. Indifference curves are a way to describe preferences. Nearly any "reasonable" preferences that you can think of can be depicted by indifference curves. The trick is to learn what kinds of preferences give rise to what shapes of indifference curves.

3.4 Examples of Preferences

Let us try to relate preferences to indifference curves through some examples. We'll describe some preferences and then see what the indifference curves that represent them look like.

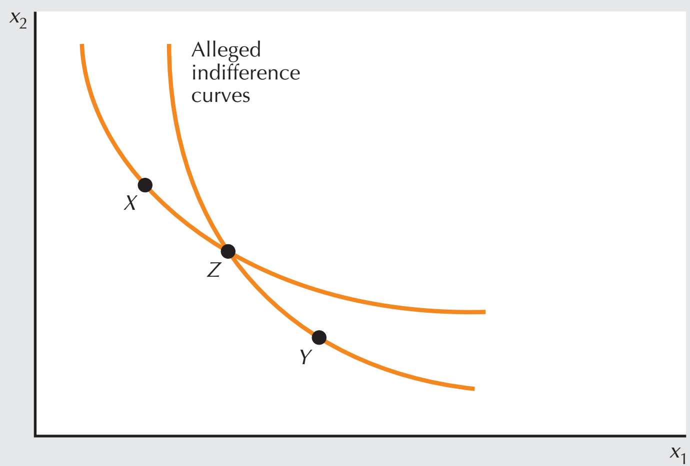

Figure 3.2

Indifference curves cannot cross. If they did, X, Y, and Z would all have to be indifferent to each other and thus could not lie on distinct indifference curves.

There is a general procedure for constructing indifference curves given a "verbal" description of the preferences. First plop your pencil down on the graph at some consumption bundle (x_1, x_2). Now think about giving a little more of good 1, Δx_1, to the consumer, moving him to $(x_1 + \Delta x_1, x_2)$. Now ask yourself how would you have to *change* the consumption of x_2 to make the consumer indifferent to the original consumption point? Call this change Δx_2. Ask yourself the question "For a given change in good 1, how does good 2 have to change to make the consumer just indifferent between $(x_1 + \Delta x_1, x_2 + \Delta x_2)$ and (x_1, x_2)?" Once you have determined this movement at one consumption bundle you have drawn a piece of the indifference curve. Now try it at another bundle, and so on, until you develop a clear picture of the overall shape of the indifference curves.

Perfect Substitutes

Two goods are **perfect substitutes** if the consumer is willing to substitute one good for the other at a *constant* rate. The simplest case of perfect substitutes occurs when the consumer is willing to substitute the goods on a one-to-one basis.

Suppose, for example, that we are considering a choice between red pencils and blue pencils, and the consumer involved likes pencils, but doesn't care about color at all. Pick a consumption bundle, say $(10, 10)$. Then for this consumer, any other consumption bundle that has 20 pencils in it is just as good as $(10, 10)$. Mathematically speaking, any consumption bundle (x_1, x_2) such that $x_1 + x_2 = 20$ will be on this consumer's indifference curve through $(10, 10)$. Thus the indifference curves for this consumer are all parallel straight lines with a slope of -1, as depicted in Figure 3.3. Bundles with more total pencils are preferred to bundles with fewer total pencils, so the direction of increasing preference is up and to the right, as illustrated in Figure 3.3.

How does this work in terms of general procedure for drawing indifference curves? If we are at $(10, 10)$, and we increase the amount of the first good by one unit to 11, how much do we have to change the second good to get back to the original indifference curve? The answer is clearly that we have to decrease the second good by 1 unit. Thus the indifference curve through $(10, 10)$ has a slope of -1. The same procedure can be carried out at any bundle of goods with the same results—in this case all the indifference curves have a constant slope of -1.

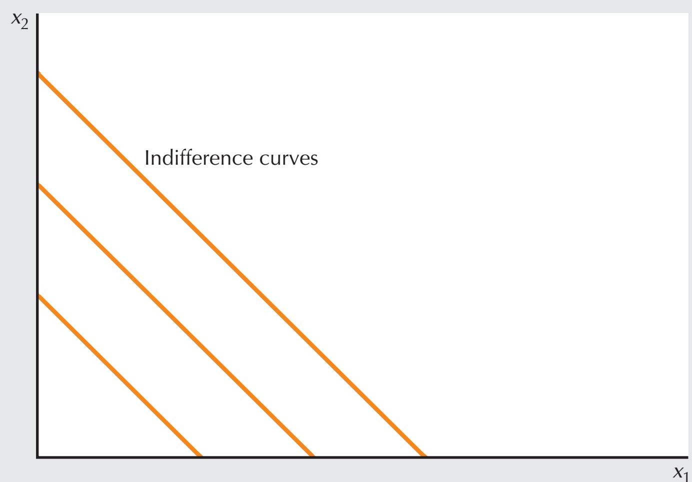

Perfect substitutes. The consumer only cares about the total number of pencils, not about their colors. Thus the indifference curves are straight lines with a slope of -1.

Figure 3.3

The important fact about perfect substitutes is that the indifference curves have a *constant* slope. Suppose, for example, that we graphed blue pencils on the vertical axis and *pairs* of red pencils on the horizontal axis. The indifference curves for these two goods would have a slope of -2, since the consumer would be willing to give up two blue pencils to get one more *pair* of red pencils.

In the textbook we'll primarily consider the case where goods are perfect substitutes on a one-for-one basis, and leave the treatment of the general case for the workbook.

Perfect Complements

Perfect complements are goods that are always consumed together in fixed proportions. In some sense the goods "complement" each other. A nice example is that of right shoes and left shoes. The consumer likes shoes, but always wears right and left shoes together. Having only one out of a pair of shoes doesn't do the consumer a bit of good.

Let us draw the indifference curves for perfect complements. Suppose we pick the consumption bundle $(10, 10)$. Now add 1 more right shoe, so we have $(11, 10)$. By assumption this leaves the consumer indifferent to the original position: the extra shoe doesn't do him any good. The same thing happens if we add one more left shoe: the consumer is also indifferent between $(10, 11)$ and $(10, 10)$.

Thus the indifference curves are L-shaped, with the vertex of the L occurring where the number of left shoes equals the number of right shoes as in Figure 3.4.

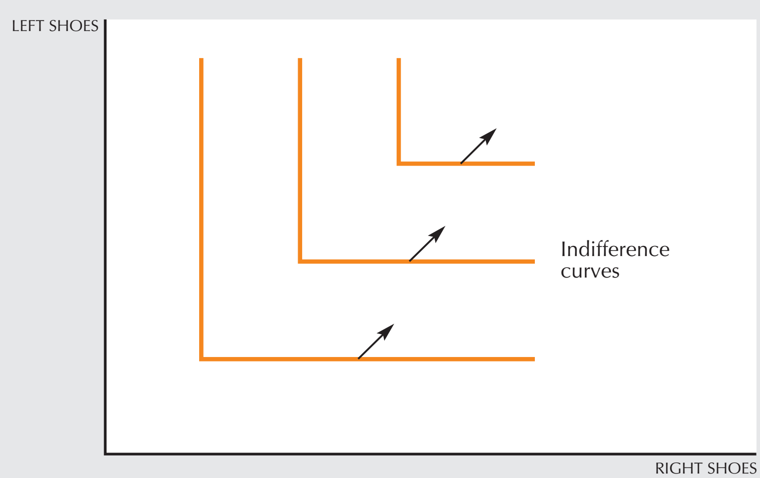

Figure 3.4 **Perfect complements**. The consumer always wants to consume the goods in fixed proportions to each other. Thus the indifference curves are L-shaped.

Increasing both the number of left shoes and the number of right shoes at the same time will move the consumer to a more preferred position, so the direction of increasing preference is again up and to the right, as illustrated in the diagram.

The important thing about perfect complements is that the consumer prefers to consume the goods in fixed proportions, not necessarily that the proportion is one-to-one. If a consumer always uses two teaspoons of sugar in her cup of tea, and doesn't use sugar for anything else, then the indifference curves will still be L-shaped. In this case the corners of the L will occur at (2 teaspoons sugar, 1 cup tea), (4 teaspoons sugar, 2 cups tea) and so on, rather than at (1 right shoe, 1 left shoe), (2 right shoes, 2 left shoes), and so on.

In the textbook we'll primarily consider the case where the goods are consumed in proportions of one-for-one and leave the treatment of the general case for the workbook.

Bads

A **bad** is a commodity that the consumer doesn't like. For example, suppose that the commodities in question are now pepperoni and anchovies—and the consumer loves pepperoni but dislikes anchovies. But let us suppose there is some possible tradeoff between pepperoni and anchovies. That is, there would be some amount of pepperoni on a pizza that would compensate the consumer for having to consume a given amount of anchovies. How could we represent these preferences using indifference curves?

Pick a bundle (x_1, x_2) consisting of some pepperoni and some anchovies. If we give the consumer more anchovies, what do we have to do with the pepperoni to keep him on the same indifference curve? Clearly, we have to give him some extra pepperoni to compensate him for having to put up with the anchovies. Thus this consumer must have indifference curves that slope up and to the right as depicted in Figure 3.5.

The direction of increasing preference is down and to the right—that is, toward the direction of decreased anchovy consumption and increased pepperoni consumption, just as the arrows in the diagram illustrate.

Neutrals

A good is a **neutral good** if the consumer doesn't care about it one way or the other. What if a consumer is just neutral about anchovies?[1] In this case his indifference curves will be vertical lines as depicted in Figure 3.6.

[1] Is anybody neutral about anchovies?

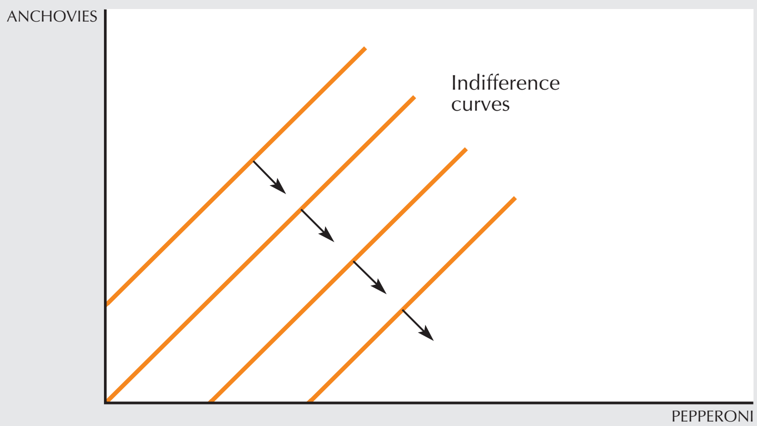

Figure 3.5

Bads. Here anchovies are a "bad," and pepperoni is a "good" for this consumer. Thus the indifference curves have a positive slope.

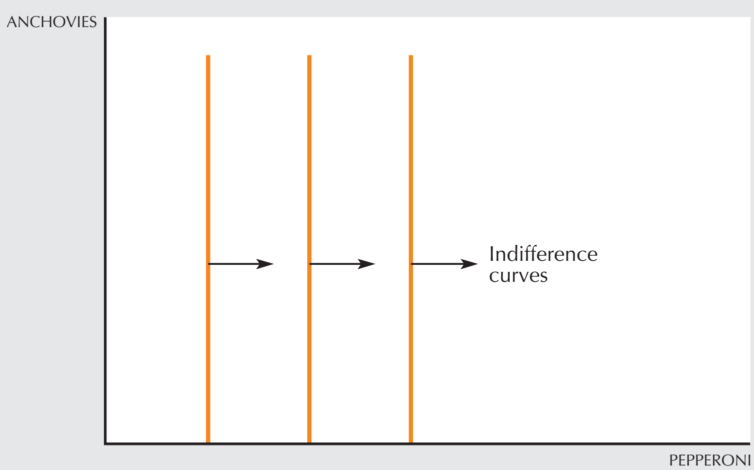

Figure 3.6

A neutral good. The consumer likes pepperoni but is neutral about anchovies, so the indifference curves are vertical lines.

He only cares about the amount of pepperoni he has and doesn't care at all about how many anchovies he has. The more pepperoni the better, but adding more anchovies doesn't affect him one way or the other.

Satiation

We sometimes want to consider a situation involving **satiation**, where there is some overall best bundle for the consumer, and the "closer" he is to that best bundle, the better off he is in terms of his own preferences. For example, suppose that the consumer has some most preferred bundle of goods $(\overline{x}_1, \overline{x}_2)$, and the farther away he is from that bundle, the worse off he is. In this case we say that $(\overline{x}_1, \overline{x}_2)$ is a **satiation** point, or a **bliss** point. The indifference curves for the consumer look like those depicted in Figure 3.7. The best point is $(\overline{x}_1, \overline{x}_2)$ and points farther away from this bliss point lie on "lower" indifference curves.

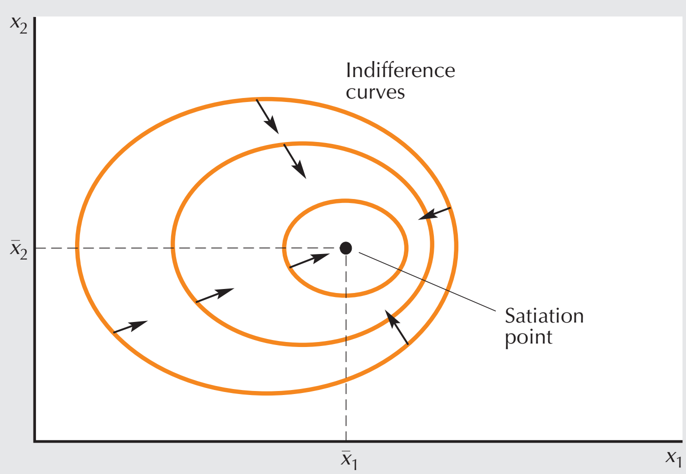

Figure 3.7

Satiated preferences. The bundle $(\overline{x}_1, \overline{x}_2)$ is the satiation point or bliss point, and the indifference curves surround this point.

In this case the indifference curves have a negative slope when the consumer has "too little" or "too much" of both goods, and a positive slope when he has "too much" of one of the goods. When he has too much of one of the goods, it becomes a bad—reducing the consumption of the bad good moves him closer to his "bliss point." If he has too much of both goods, they both are bads, so reducing the consumption of each moves him closer to the bliss point.

Suppose, for example, that the two goods are chocolate cake and ice cream. There might well be some optimal amount of chocolate cake and ice cream that you would want to eat per week. Any less than that amount would make you worse off, but any more than that amount would also make you worse off.

If you think about it, most goods are like chocolate cake and ice cream in this respect—you can have too much of nearly anything. But people would generally not voluntarily *choose* to have too much of the goods they consume. Why would you choose to have more than you want of something? Thus the interesting region from the viewpoint of economic choice is where you have *less* than you want of most goods. The choices that people actually care about are choices of this sort, and these are the choices with which we will be concerned.

Discrete Goods

Usually we think of measuring goods in units where fractional amounts make sense—you might on average consume 12.43 gallons of milk a month even though you buy it a quart at a time. But sometimes we want to examine preferences over goods that naturally come in discrete units.

For example, consider a consumer's demand for automobiles. We could define the demand for automobiles in terms of the time spent using an automobile, so that we would have a continuous variable, but for many purposes it is the actual number of cars demanded that is of interest.

There is no difficulty in using preferences to describe choice behavior for this kind of discrete good. Suppose that x_2 is money to be spent on other goods and x_1 is a **discrete good** that is only available in integer amounts. We have illustrated the appearance of indifference "curves" and a weakly preferred set for this kind of good in Figure 3.8. In this case the bundles indifferent to a given bundle will be a set of discrete points. The set of bundles at least as good as a particular bundle will be a set of line segments.

The choice of whether to emphasize the discrete nature of a good or not will depend on our application. If the consumer chooses only one or two units of the good during the time period of our analysis, recognizing the discrete nature of the choice may be important. But if the consumer is choosing 30 or 40 units of the good, then it will probably be convenient to think of this as a continuous good.

3.5 Well-Behaved Preferences

We've now seen some examples of indifference curves. As we've seen, many kinds of preferences, reasonable or unreasonable, can be described by these

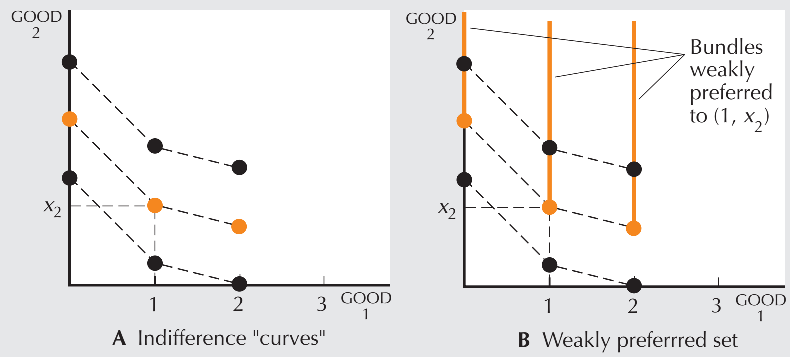

Figure 3.8 **A discrete good**. Here good 1 is only available in integer amounts. In panel A the dashed lines connect together the bundles that are indifferent, and in panel B the vertical lines represent bundles that are at least as good as the indicated bundle.

simple diagrams. But if we want to describe preferences in general, it will be convenient to focus on a few general shapes of indifference curves. In this section we will describe some more general assumptions that we will typically make about preferences and the implications of these assumptions for the shapes of the associated indifference curves. These assumptions are not the only possible ones; in some situations you might want to use different assumptions. But we will take them as the defining features for **well-behaved indifference curves**.

First we will typically assume that more is better, that is, that we are talking about *goods*, not bads. More precisely, if (x_1, x_2) is a bundle of goods and (y_1, y_2) is a bundle of goods with at least as much of both goods and more of one, then $(y_1, y_2) \succ (x_1, x_2)$. This assumption is sometimes called **monotonicity** of preferences. As we suggested in our discussion of satiation, more is better would probably only hold up to a point. Thus the assumption of monotonicity is saying only that we are going to examine situations *before* that point is reached—before any satiation sets in—while more *still is* better. Economics would not be a very interesting subject in a world where everyone was satiated in their consumption of every good.

What does monotonicity imply about the shape of indifference curves? It implies that they have a *negative* slope. Consider Figure 3.9. If we start at a bundle (x_1, x_2) and move anywhere up and to the right, we must be moving to a preferred position. If we move down and to the left we must be moving to a worse position. So if we are moving to an *indifferent* position, we must be moving either left and up or right and down: the indifference curve must have a negative slope.

Second, we are going to assume that averages are preferred to extremes. That is, if we take two bundles of goods (x_1, x_2) and (y_1, y_2) on the same indifference curve and take a weighted average of the two bundles such as

$$\left(\frac{1}{2}x_1 + \frac{1}{2}y_1, \frac{1}{2}x_2 + \frac{1}{2}y_2\right),$$

then the average bundle will be at least as good as or strictly preferred to each of the two extreme bundles. This weighted-average bundle has the average amount of good 1 and the average amount of good 2 that is present in the two bundles. It therefore lies halfway along the straight line connecting the x–bundle and the y–bundle.

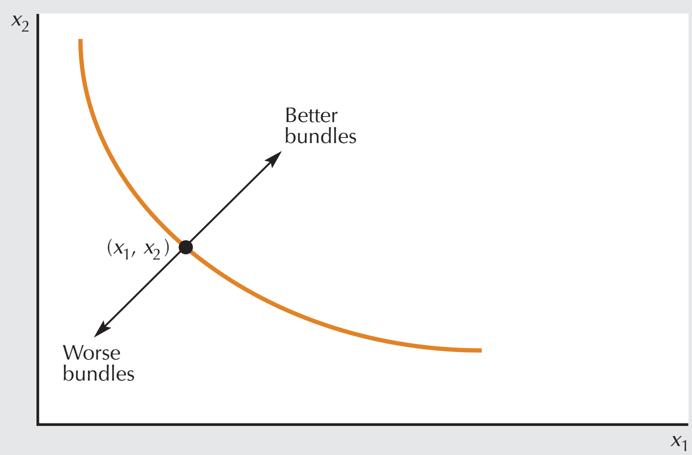

Figure 3.9 **Monotonic preferences**. More of both goods is a better bundle for this consumer; less of both goods represents a worse bundle.

Actually, we're going to assume this for any weight t between 0 and 1, not just 1/2. Thus we are assuming that if $(x_1, x_2) \sim (y_1, y_2)$, then

$$(tx_1 + (1-t)y_1, tx_2 + (1-t)y_2) \succeq (x_1, x_2)$$

for any t such that $0 \leq t \leq 1$. This weighted average of the two bundles gives a weight of t to the x-bundle and a weight of $1 - t$ to the y-bundle.

Therefore, the distance from the x-bundle to the average bundle is just a fraction t of the distance from the x-bundle to the y-bundle, along the straight line connecting the two bundles.

What does this assumption about preferences mean geometrically? It means that the set of bundles weakly preferred to (x_1, x_2) is a **convex set**. For suppose that (y_1, y_2) and (x_1, x_2) are indifferent bundles. Then, if averages are preferred to extremes, all of the weighted averages of (x_1, x_2) and (y_1, y_2) are weakly preferred to (x_1, x_2) and (y_1, y_2). A convex set has the property that if you take *any* two points in the set and draw the line segment connecting those two points, that line segment lies entirely in the set.

Figure 3.10A depicts an example of convex preferences, while Figures 3.10B and 3.10C show two examples of nonconvex preferences. Figure 3.10C presents preferences that are so nonconvex that we might want to call them "concave preferences."

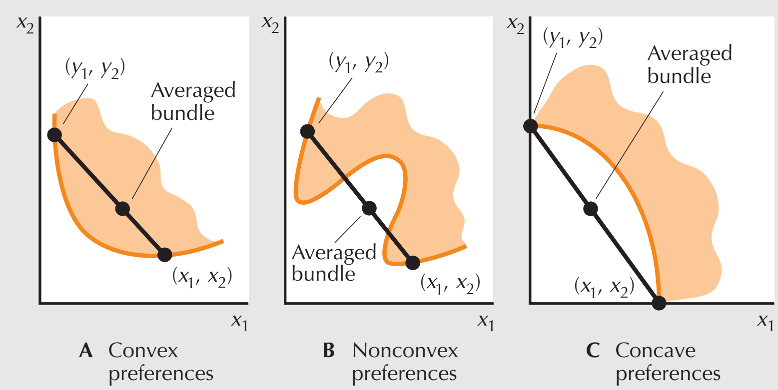

Figure 3.10

Various kinds of preferences. Panel A depicts convex preferences, panel B depicts nonconvex preferences, and panel C depicts "concave" preferences.

Can you think of preferences that are not convex? One possibility might be something like my preferences for ice cream and olives. I like ice cream and I like olives . . . but I don't like to have them together! In considering my consumption in the next hour, I might be indifferent between consuming 8 ounces of ice cream and 2 ounces of olives, or 8 ounces of olives and 2 ounces of ice cream. But either one of these bundles would be better than consuming 5 ounces of each! These are the kind of preferences depicted in Figure 3.10C.

Why do we want to assume that well-behaved preferences are convex? Because, for the most part, goods are consumed together. The kinds of preferences depicted in Figures 3.10B and 3.10C imply that the consumer would prefer to specialize, at least to some degree, and to consume only one of the goods. However, the normal case is where the consumer would want to trade some of one good for the other and end up consuming some of each, rather than specializing in consuming only one of the two goods.

In fact, if we look at my preferences for *monthly* consumption of ice cream and olives, rather than at my immediate consumption, they would tend to look much more like Figure 3.10A than Figure 3.10C. Each month I would prefer having some ice cream and some olives—albeit at different times—to specializing in consuming either one for the entire month.

Finally, one extension of the assumption of convexity is the assumption of **strict convexity**. This means that the weighted average of two indifferent bundles is *strictly* preferred to the two extreme bundles. Convex preferences may have flat spots, while *strictly* convex preferences must have indifferences curves that are "rounded." The preferences for two goods that are perfect substitutes are convex, but not strictly convex.

3.6 The Marginal Rate of Substitution

We will often find it useful to refer to the slope of an indifference curve at a particular point. This idea is so useful that it even has a name: the slope of an indifference curve is known as the **marginal rate of substitution (MRS)**. The name comes from the fact that the MRS measures the rate at which the consumer is just willing to substitute one good for the other.

Suppose that we take a little of good 1, Δx_1, away from the consumer. Then we give him Δx_2, an amount that is just sufficient to put him back on his indifference curve, so that he is just as well off after this substitution of x_2 for x_1 as he was before. We think of the ratio $\Delta x_2/\Delta x_1$ as being the *rate* at which the consumer is willing to substitute good 2 for good 1.

Now think of Δx_1 as being a very small change—a marginal change. When we imagine Δx_1 becoming infintesimally small, we use the notation dx_1. Then the rate dx_2/dx_1 measures the *marginal* rate of substitution of good 2 for good 1, or the slope of the indifference curve, as can be seen in Figure 3.11.

In economics, we commonly write about *marginal* changes, which you will learn to associate with slopes. The ratio defining the MRS will always describe the slope of the indifference curve: the rate at which the consumer is just willing to substitute a little more consumption of good 2 for a little less consumption of good 1.

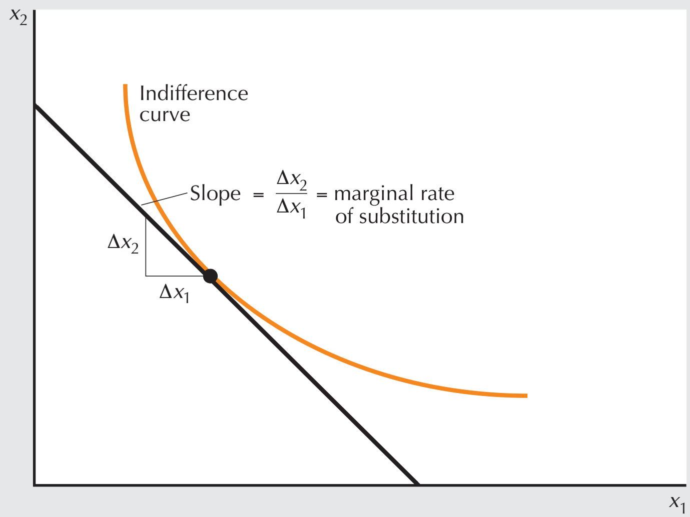

The marginal rate of substitution (MRS). The marginal rate of substitution measures the slope of the indifference curve.

Figure 3.11

One slightly confusing thing about the MRS is that it is typically a *negative* number. We've already seen that monotonic preferences imply that indifference curves must have a negative slope. Since the MRS is the numerical measure of the slope of an indifference curve, it will naturally be a negative number.

The marginal rate of substitution measures an interesting aspect of the consumer's behavior. Suppose that the consumer has well-behaved preferences, that is, preferences that are monotonic and convex, and that he is currently consuming some bundle (x_1, x_2). We now will offer him a trade: he can exchange good 1 for 2, or good 2 for 1, in any amount at a "rate of exchange" of E.

That is, if the consumer gives up dx_1 units of good 1, he can get Edx_1 units of good 2 in exchange. Or, conversely, if he gives up dx_2 units of good 2, he can get dx_2/E units of good 1. Geometrically, we are offering the consumer an opportunity to move to any point along a line with slope $-E$ that passes through (x_1, x_2), as depicted in Figure 3.12. Moving up and to the left from (x_1, x_2) involves exchanging good 1 for good 2, and moving down and to the right involves exchanging good 2 for good 1. In either movement, the exchange rate is E. Since exchange always involves giving up one good in exchange for another, the exchange *rate* E corresponds to a *slope* of $-E$.

We can now ask what would the rate of exchange have to be in order for the consumer to want to stay put at (x_1, x_2)? To answer this question, we simply note that any time the exchange line *crosses* the indifference curve, there will be some points on that line that are preferred to (x_1, x_2)—that lie above the indifference curve. Thus, if there is to be no movement from (x_1, x_2), the exchange line must be tangent to the indifference curve. That is, the slope of the exchange line, $-E$, must be the slope of the indifference curve at (x_1, x_2). At any other rate of exchange, the exchange line would cut the indifference curve and thus allow the consumer to move to a more preferred point.

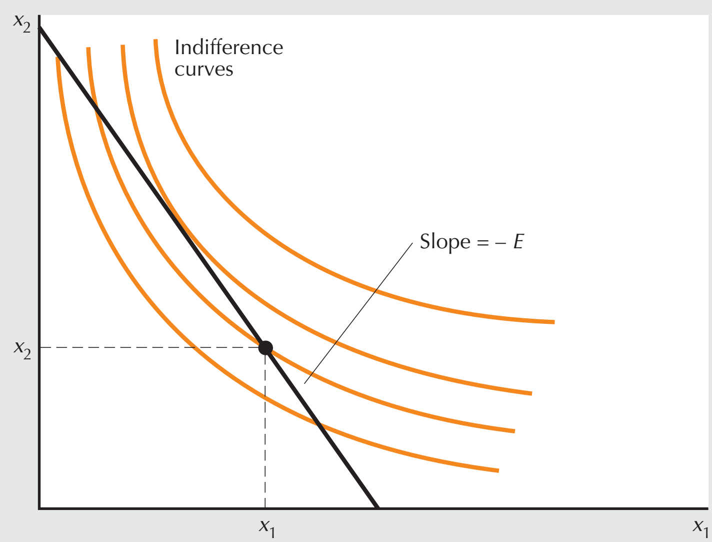

Figure 3.12 **Trading at an exchange rate**. Here we are allowing the consumer to trade the goods at an exchange rate E, which implies that the consumer can move along a line with slope $-E$.

Thus the slope of the indifference curve, the marginal rate of substitution, measures the rate at which the consumer is just on the margin of trading or not trading. At any rate of exchange other than the MRS, the consumer would want to trade one good for the other. But if the rate of exchange equals the MRS, the consumer wants to stay put.

3.7 Other Interpretations of the MRS

We have said that the MRS measures the rate at which the consumer is just on the margin of being willing to substitute good 1 for good 2. We could also say that the consumer is just on the margin of being willing to "pay" some of good 1 in order to buy some more of good 2. So sometimes you hear people say that the slope of the indifference curve measures the **marginal willingness to pay**.

If good 2 represents the consumption of "all other goods," and it is measured in dollars that you can spend on other goods, then the marginal-willingness-to-pay interpretation is very natural. The marginal rate of substitution of good 2 for good 1 is how many dollars you would just be willing to give up spending on other goods in order to consume a little bit more of good 1. Thus the MRS measures the marginal willingness to give up dollars in order to consume a small amount more of good 1. But giving up those dollars is just like paying dollars in order to consume a little more of good 1.

If you use the marginal-willingness-to-pay interpretation of the MRS, you should be careful to emphasize both the "marginal" and the "willingness" aspects. The MRS measures the amount of good 2 that one is *willing* to pay for a *marginal* amount of extra consumption of good 1. How much you actually *have* to pay for some given amount of extra consumption may be different than the amount you are willing to pay. How much you have to pay will depend on the price of the good in question. How much you are willing to pay doesn't depend on the price—it is determined by your preferences.

Similarly, how much you may be willing to pay for a large change in consumption may be different from how much you are willing to pay for a marginal change. How much you actually end up buying of a good will depend on your preferences for that good and the prices that you face. How much you would be willing to pay for a small amount extra of the good is a feature only of your preferences.

3.8 Behavior of the MRS

It is sometimes useful to describe the shapes of indifference curves by describing the behavior of the marginal rate of substitution. For example, the "perfect substitutes" indifference curves are characterized by the fact that the MRS is constant at -1. The "neutrals" case is characterized by the fact that the MRS is everywhere infinite. The preferences for "perfect complements" are characterized by the fact that the MRS is either zero or infinity, and nothing in between.

We've already pointed out that the assumption of monotonicity implies that indifference curves must have a negative slope, so the MRS always involves reducing the consumption of one good in order to get more of another for monotonic preferences.

The case of convex indifference curves exhibits yet another kind of behavior for the MRS. For strictly convex indifference curves, the MRS—the slope of the indifference curve—decreases (in absolute value) as we increase x_1. Thus the indifference curves exhibit a **diminishing marginal rate of substitution**. This means that the amount of good 1 that the person is willing to give up for an additional amount of good 2 increases the amount of good 1 increases. Stated in this way, convexity of indifference curves seems very natural: it says that the more you have of one good, the more willing you are to give some of it up in exchange for the other good. (But remember the ice cream and olives example—for some pairs of goods this assumption might not hold!)

Summary

1. Economists assume that a consumer can rank various consumption possibilities. The way in which the consumer ranks the consumption bundles describes the consumer's preferences.

2. Indifference curves can be used to depict different kinds of preferences.

3. Well-behaved preferences are monotonic (meaning more is better) and convex (meaning averages are preferred to extremes).

4. The marginal rate of substitution (MRS) measures the slope of the indifference curve. This can be interpreted as how much the consumer is willing to give up of good 2 to acquire more of good 1.

REVIEW QUESTIONS

1. If we observe a consumer choosing (x_1, x_2) when (y_1, y_2) is available one time, are we justified in concluding that $(x_1, x_2) \succ (y_1, y_2)$?

2. Consider a group of people A, B, C and the relation "at least as tall as," as in "A is at least as tall as B." Is this relation transitive? Is it complete?

3. Take the same group of people and consider the relation "strictly taller than." Is this relation transitive? Is it reflexive? Is it complete?

4. A college football coach says that given any two linemen A and B, he always prefers the one who is bigger and faster. Is this preference relation transitive? Is it complete?

5. Can an indifference curve cross itself? For example, could Figure 3.2 depict a single indifference curve?

6. Could Figure 3.2 be a single indifference curve if preferences are monotonic?

7. If both pepperoni and anchovies are bads, will the indifference curve have a positive or a negative slope?

8. Explain why convex preferences means that "averages are preferred to extremes."

9. What is your marginal rate of substitution of \$1 bills for \$5 bills?

10. If good 1 is a "neutral," what is its marginal rate of substitution for good 2?

11. Think of some other goods for which your preferences might be concave.

CHAPTER 4

UTILITY

In Victorian days, philosophers and economists talked blithely of "utility" as an indicator of a person's overall well-being. Utility was thought of as a numeric measure of a person's happiness. Given this idea, it was natural to think of consumers making choices so as to maximize their utility, that is, to make themselves as happy as possible.

The trouble is that these classical economists never really described how we were to measure utility. How are we supposed to quantify the "amount" of utility associated with different choices? Is one person's utility the same as another's? What would it mean to say that an extra candy bar would give me twice as much utility as an extra carrot? Does the concept of utility have any independent meaning other than its being what people maximize?

Because of these conceptual problems, economists have abandoned the old-fashioned view of utility as being a measure of happiness. Instead, the theory of consumer behavior has been reformulated entirely in terms of **consumer preferences**, and utility is seen only as a *way to describe preferences.*

Economists gradually came to recognize that all that mattered about utility as far as choice behavior was concerned was whether one bundle had a higher utility than another—how much higher didn't really matter.

Originally, preferences were defined in terms of utility: to say a bundle (x_1, x_2) was preferred to a bundle (y_1, y_2) meant that the x-bundle had a higher utility than the y-bundle. But now we tend to think of things the other way around. The *preferences* of the consumer are the fundamental description useful for analyzing choice, and utility is simply a way of describing preferences.

A **utility function** is a way of assigning a number to every possible consumption bundle such that more-preferred bundles get assigned larger numbers than less-preferred bundles. That is, a bundle (x_1, x_2) is preferred to a bundle (y_1, y_2) if and only if the utility of (x_1, x_2) is larger than the utility of (y_1, y_2): in symbols, $(x_1, x_2) \succ (y_1, y_2)$ if and only if $u(x_1, x_2) > u(y_1, y_2)$.

The only property of a utility assignment that is important is how it *orders* the bundles of goods. The magnitude of the utility function is only important insofar as it *ranks* the different consumption bundles; the size of the utility difference between any two consumption bundles doesn't matter. Because of this emphasis on ordering bundles of goods, this kind of utility is referred to as **ordinal utility**.

Consider for example Table 4.1, where we have illustrated several different ways of assigning utilities to three bundles of goods, all of which order the bundles in the same way. In this example, the consumer prefers A to B and B to C. All of the ways indicated are valid utility functions that describe the same preferences because they all have the property that A is assigned a higher number than B, which in turn is assigned a higher number than C.

Different ways to assign utilities. **Table 4.1**

Bundle	U_1	U_2	U_3
A	3	17	-1
B	2	10	-2
C	1	.002	-3

Since only the ranking of the bundles matters, there can be no unique way to assign utilities to bundles of goods. If we can find one way to assign utility numbers to bundles of goods, we can find an infinite number of ways to do it. If $u(x_1, x_2)$ represents a way to assign utility numbers to the bundles (x_1, x_2), then multiplying $u(x_1, x_2)$ by 2 (or any other positive number) is just as good a way to assign utilities.

Multiplication by 2 is an example of a **monotonic transformation**. A monotonic transformation is a way of transforming one set of numbers into another set of numbers in a way that preserves the order of the numbers.

We typically represent a monotonic transformation by a function $f(u)$ that transforms each number u into some other number $f(u)$, in a way that preserves the order of the numbers in the sense that $u_1 > u_2$ implies $f(u_1) > f(u_2)$. A monotonic transformation and a monotonic function are essentially the same thing.

Examples of monotonic transformations are multiplication by a positive number (e.g., $f(u) = 3u$), adding any number (e.g., $f(u) = u + 17$), raising u to an odd power (e.g., $f(u) = u^3$), and so on.[1]

Provided that a monotonic function $f(u)$ is differentiable, its derivative $f'(u)$ is given by

$$f'(u) = \lim_{\hat{u} \to u} \frac{f(\hat{u}) - f(u)}{\hat{u} - u} > 0$$

For a monotonic transformation, $f(\hat{u}) - f(u)$ always has the same sign as $\hat{u} - u$. Thus a monotonic function always has a positive first derivative. This means that the graph of a monotonic function will always have a positive slope, as depicted in Figure 4.1A.

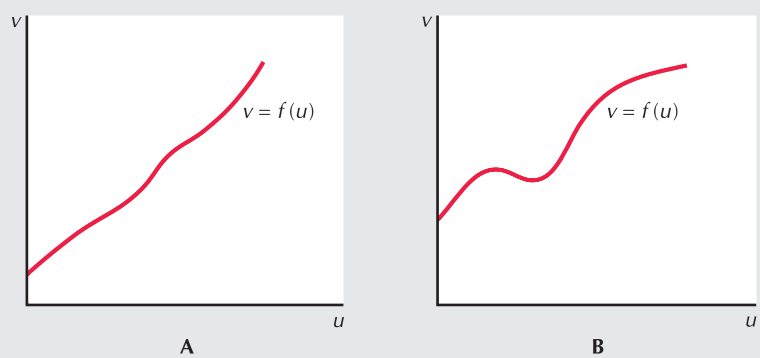

Figure 4.1 **A positive monotonic transformation**. Panel A illustrates a monotonic function—one that is always increasing. Panel B illustrates a function that is *not* monotonic, since it sometimes increases and sometimes decreases.

[1] What we are calling a "monotonic transformation" is, strictly speaking, called a "positive monotonic transformation," in order to distinguish it from a "negative monotonic transformation," which is one that *reverses* the order of the numbers. Monotonic transformations are sometimes called "monotonous transformations," which seems unfair, since they can actually be quite interesting.

If $f(u)$ is *any* monotonic transformation of a utility function that represents some particular preferences, then $f(u(x_1, x_2))$ is also a utility function that represents those same preferences.

Why? The argument is given in the following three statements:

1. To say that $u(x_1, x_2)$ represents some particular preferences means that $u(x_1, x_2) > u(y_1, y_2)$ if and only if $(x_1, x_2) \succ (y_1, y_2)$.
2. But if $f(u)$ is a monotonic transformation, then $u(x_1, x_2) > u(y_1, y_2)$ if and only if $f(u(x_1, x_2)) > f(u(y_1, y_2))$.
3. Therefore, $f(u(x_1, x_2)) > f(u(y_1, y_2))$ if and only if $(x_1, x_2) \succ (y_1, y_2)$, so the function $f(u)$ represents the preferences in the same way as the original utility function $u(x_1, x_2)$.

We summarize this discussion by stating the following principle: *a monotonic transformation of a utility function is a utility function that represents the same preferences as the original utility function.*

Geometrically, a utility function is a way to label indifference curves. Since every bundle on an indifference curve must have the same utility, a utility function is a way of assigning numbers to the different indifference curves in a way that higher indifference curves get assigned larger numbers. Seen from this point of view a monotonic transformation is just a relabeling of indifference curves. As long as indifference curves containing more-preferred bundles get a larger label than indifference curves containing less-preferred bundles, the labeling will represent the same preferences.

4.1 Cardinal Utility

There are some theories of utility that attach a significance to the magnitude of utility. These are known as **cardinal utility** theories. In a theory of cardinal utility, the size of the utility difference between two bundles of goods is supposed to have some sort of significance.

We know how to tell whether a given person prefers one bundle of goods to another: we simply offer him or her a choice between the two bundles and see which one is chosen. Thus we know how to assign an ordinal utility to the two bundles of goods: we just assign a higher utility to the chosen bundle than to the rejected bundle. Any assignment that does this will be a utility function. Thus we have an operational criterion for determining whether one bundle has a higher utility than another bundle for some individual.

But how do we tell if a person likes one bundle twice as much as another? How could you even tell if *you* like one bundle twice as much as another?

One could propose various definitions for this kind of assignment: I like one bundle twice as much as another if I am willing to pay twice as much for it. Or, I like one bundle twice as much as another if I am willing to run twice as far to get it, or to wait twice as long, or to gamble for it at twice the odds.

There is nothing wrong with any of these definitions; each one would give rise to a way of assigning utility levels in which the magnitude of the numbers assigned had some operational significance. But there isn't much right about them either. Although each of them is a possible interpretation of what it means to want one thing twice as much as another, none of them appears to be an especially compelling interpretation of that statement.

Even if we did find a way of assigning utility magnitudes that seemed to be especially compelling, what good would it do us in describing choice behavior? To tell whether one bundle or another will be chosen, we only have to know which is preferred—which has the larger utility. Knowing how much larger doesn't add anything to our description of choice. Since cardinal utility isn't needed to describe choice behavior and there is no compelling way to assign cardinal utilities anyway, we will stick with a purely ordinal utility framework.

4.2 Constructing a Utility Function

But are we assured that there is any way to assign ordinal utilities? Given a preference ordering can we always find a utility function that will order bundles of goods in the same way as those preferences? Is there a utility function that describes any reasonable preference ordering?

Not all kinds of preferences can be represented by a utility function. For example, suppose that someone had intransitive preferences so that $A \succ B \succ C \succ A$. Then a utility function for these preferences would have to consist of numbers $u(A)$, $u(B)$, and $u(C)$ such that $u(A) > u(B) > u(C) > u(A)$. But this is impossible.

However, if we rule out perverse cases like intransitive preferences, it turns out that we will typically be able to find a utility function to represent preferences. We will illustrate one construction here, and another one in Chapter 14.

Suppose that we are given an indifference map as in Figure 4.2. We know that a utility function is a way to label the indifference curves such that higher indifference curves get larger numbers. How can we do this?

One easy way is to draw the diagonal line illustrated and label each indifference curve with its distance from the origin measured along the line.

How do we know that this is a utility function? It is not hard to see that if preferences are monotonic then the line through the origin must intersect every indifference curve exactly once. Thus every bundle is getting a label, and those bundles on higher indifference curves are getting larger labels—and that's all it takes to be a utility function.

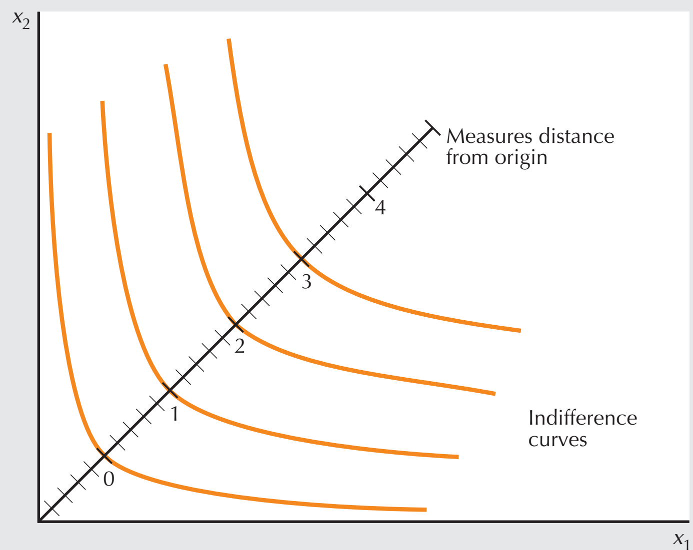

Constructing a utility function from indifference curves. Draw a diagonal line and label each indifference curve with how far it is from the origin measured along the line.

Figure 4.2

This gives us one way to find a labeling of indifference curves, at least as long as preferences are monotonic. This won't always be the most natural way in any given case, but at least it shows that the idea of an ordinal utility function is pretty general: nearly any kind of "reasonable" preferences can be represented by a utility function.

4.3 Some Examples of Utility Functions

In Chapter 3 we described some examples of preferences and the indifference curves that represented them. We can also represent these preferences by utility functions. If you are given a utility function, $u(x_1, x_2)$, it is relatively easy to draw the indifference curves: you just plot all the points (x_1, x_2) such that $u(x_1, x_2)$ equals a constant. In mathematics, the set of all (x_1, x_2) such that $u(x_1, x_2)$ equals a constant is called a **level set**. For each different value of the constant, you get a different indifference curve.

EXAMPLE: Indifference Curves from Utility

Suppose that the utility function is given by: $u(x_1, x_2) = x_1 x_2$. What do the indifference curves look like?

We know that a typical indifference curve is just the set of all x_1 and x_2 such that $k = x_1x_2$ for some constant k. Solving for x_2 as a function of x_1, we see that a typical indifference curve has the formula:

$$x_2 = \frac{k}{x_1}.$$

This curve is depicted in Figure 4.3 for $k = 1, 2, 3 \cdots$.

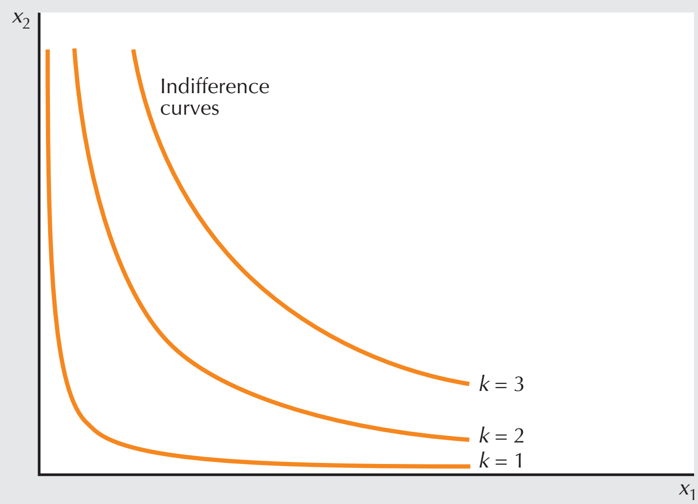

Figure 4.3 **Indifference curves.** The indifference curves $k = x_1x_2$ for different values of k.

Let's consider another example. Suppose that we were given a utility function $v(x_1, x_2) = x_1^2x_2^2$. What do its indifference curves look like? By the standard rules of algebra we know that:

$$v(x_1, x_2) = x_1^2x_2^2 = (x_1x_2)^2 = u(x_1, x_2)^2.$$

Thus the utility function $v(x_1, x_2)$ is just the square of the utility function $u(x_1, x_2)$. Since $u(x_1, x_2)$ cannot be negative, it follows that $v(x_1, x_2)$ is a monotonic transformation of the previous utility function, $u(x_1, x_2)$. This means that the utility function $v(x_1, x_2) = x_1^2x_2^2$ has to have exactly the same shaped indifference curves as those depicted in Figure 4.3. The labeling of the indifference curves will be different—the labels that were

$1, 2, 3, \cdots$ will now be $1, 4, 9, \cdots$—but the set of bundles that has $v(x_1, x_2) = 9$ is exactly the same as the set of bundles that has $u(x_1, x_2) = 3$. Thus $v(x_1, x_2)$ describes exactly the same preferences as $u(x_1, x_2)$ since it *orders* all of the bundles in the same way.

Going the other direction—finding a utility function that represents some indifference curves—is somewhat more difficult. There are two ways to proceed. The first way is mathematical. Given the indifference curves, we want to find a function that is constant along each indifference curve and that assigns higher values to higher indifference curves.

The second way is a bit more intuitive. Given a description of the preferences, we try to think about what the consumer is trying to maximize—what combination of the goods describes the choice behavior of the consumer. This may seem a little vague at the moment, but it will be more meaningful after we discuss a few examples.

Perfect Substitutes

Remember the red pencil and blue pencil example? All that mattered to the consumer was the total number of pencils. Thus it is natural to measure utility by the total number of pencils. Therefore we provisionally pick the utility function $u(x_1, x_2) = x_1 + x_2$. Does this work? Just ask two things: is this utility function constant along the indifference curves? Does it assign a higher label to more-preferred bundles? The answer to both questions is yes, so we have a utility function.

Of course, this isn't the only utility function that we could use. We could also use the *square* of the number of pencils. Thus the utility function $v(x_1, x_2) = (x_1 + x_2)^2 = x_1^2 + 2x_1x_2 + x_2^2$ will also represent the perfect-substitutes preferences, as would any other monotonic transformation of $u(x_1, x_2)$.

What if the consumer is willing to substitute good 1 for good 2 at a rate that is different from one-to-one? Suppose, for example, that the consumer would require *two* units of good 2 to compensate him for giving up one unit of good 1. This means that good 1 is *twice* as valuable to the consumer as good 2. The utility function therefore takes the form $u(x_1, x_2) = 2x_1 + x_2$. Note that this utility yields indifference curves with a slope of -2.

In general, preferences for perfect substitutes can be represented by a utility function of the form

$$u(x_1, x_2) = ax_1 + bx_2.$$

Here a and b are some positive numbers that measure the "value" of goods 1 and 2 to the consumer. Note that the slope of a typical indifference curve is given by $-a/b$.