

EIGHTH EDITION

LABOR ECONOMICS

GEORGE J. BORJAS

**Mc
Graw
Hill**
Education

Labor Economics

Eighth Edition

George J. Borjas

Harvard University





LABOR ECONOMICS, EIGHTH EDITION

Published by McGraw-Hill Education, 2 Penn Plaza, New York, NY 10121. Copyright © 2020 by McGraw-Hill Education. All rights reserved. Printed in the United States of America. Previous editions © 2016, 2013, and 2010. No part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written consent of McGraw-Hill Education, including, but not limited to, in any network or other electronic storage or transmission, or broadcast for distance learning.

Some ancillaries, including electronic and print components, may not be available to customers outside the United States.

This book is printed on acid-free paper.

1 2 3 4 5 6 7 8 9 LCR 21 20 19

ISBN 978-1-260-00472-4 (bound edition)

MHID 1-260-00472-4 (bound edition)

ISBN 978-1-260-48439-7 (loose-leaf edition)

MHID 1-260-48439-4 (loose-leaf edition)

Associate Portfolio Manager: *Kevin White*

Marketing Manager: *Bobby Pearson*

Content Project Managers: *Lisa Bruflodt, Danielle Clement*

Buyer: *Laura Fuller*

Designer: *Beth Blech*

Content Licensing Specialist: *Lorraine Buczek*

Cover Image: *Simfo/Getty Images*

Compositor: *SPi Global*

All credits appearing on page or at the end of the book are considered to be an extension of the copyright page.

Library of Congress Cataloging-in-Publication Data

Names: Borjas, George J., author.

Title: Labor economics / George J. Borjas, Harvard University.

Description: Eighth edition. | New York, NY: McGraw-Hill Education, [2020]

Identifiers: LCCN 2018037650 | ISBN 9781260004724 (alk. paper)

Subjects: LCSH: Labor economics. | Labor market—United States.

Classification: LCC HD4901.B674 2020 | DDC 331.0973—dc23 LC record available at

<https://lccn.loc.gov/2018037650>

The Internet addresses listed in the text were accurate at the time of publication. The inclusion of a website does not indicate an endorsement by the authors or McGraw-Hill Education, and McGraw-Hill Education does not guarantee the accuracy of the information presented at these sites.

About the Author

George J. Borjas

George J. Borjas is the Robert W. Scrivner Professor of Economics and Social Policy at the John F. Kennedy School of Government, Harvard University. He is also a research associate at the National Bureau of Economic Research and a Research Fellow at IZA. Professor Borjas received his Ph.D. in economics from Columbia University.

Professor Borjas has written extensively on labor market issues. He is the author of several books, including *Wage Policy in the Federal Bureaucracy* (American Enterprise Institute, 1980), *Friends or Strangers: The Impact of Immigrants on the U.S. Economy* (Basic Books, 1990), *Heaven's Door: Immigration Policy and the American Economy* (Princeton University Press, 1999), *Immigration Economics* (Harvard University Press, 2014), and *We Wanted Workers: Unraveling the Immigration Narrative* (Norton, 2016). He has published more than 150 articles in books and scholarly journals, including the *American Economic Review*, the *Journal of Political Economy*, and the *Quarterly Journal of Economics*.

Professor Borjas was elected a Fellow of the Econometric Society in 1998, and a Fellow of the Society of Labor Economics in 2004. In 2011, Professor Borjas was awarded the IZA Prize in Labor Economics. He was an editor of the *Review of Economics and Statistics* from 1998 to 2006. He also has served as a member of the Advisory Panel in Economics at the National Science Foundation and has testified frequently before congressional committees and government commissions.

The McGraw-Hill Series Economics

ESSENTIALS OF ECONOMICS

Brue, McConnell, and Flynn

Essentials of Economics

Fourth Edition

Mandel

Economics: The Basics

Third Edition

Schiller

Essentials of Economics

Tenth Edition

PRINCIPLES OF ECONOMICS

Asarta and Butters

Principles of Economics

Colander

Economics, Microeconomics, and Macroeconomics

Eleventh Edition

Frank, Bernanke, Antonovics, and
Heffetz

Principles of Economics, Principles of Microeconomics, Principles of Macroeconomics

Seventh Edition

Frank, Bernanke, Antonovics, and
Heffetz

Streamlined Editions:

Principles of

Economics, Principles of Microeconomics, Principles of Macroeconomics

Third Edition

Karlan and Morduch

Economics, Microeconomics, and Macroeconomics

Second Edition

McConnell, Brue, and Flynn Economics, Microeconomics, Macroeconomics

Twenty-First Edition

McConnell, Brue, and Flynn

Brief Editions: Microeconomics and Macroeconomics

Third Edition

Samuelson and Nordhaus

Economics, Microeconomics, and Macroeconomics

Nineteenth Edition

Schiller

The Economy Today, The Micro Economy Today, and The Macro Economy Today

Fifteenth Edition

Slavin

Economics, Microeconomics, and Macroeconomics

Twelfth Edition

ECONOMICS OF SOCIAL

ISSUES

Guell

Issues in Economics Today

Eighth Edition

Register and Grimes

Economics of Social Issues

Twenty-First Edition

ECONOMETRICS AND DATA

ANALYTICS

Hilmer and Hilmer

Practical Econometrics

First Edition

Prince

Predictive Analytics for Business Strategy

First Edition

MANAGERIAL ECONOMICS

Baye and Prince

Managerial Economics and Business Strategy

Ninth Edition

Brickley, Smith, and Zimmerman

Managerial Economics and Organizational Architecture

Sixth Edition

Thomas and Maurice

Managerial Economics

Twelfth Edition

INTERMEDIATE ECONOMICS

Bernheim and Whinston

Microeconomics

Second Edition

Dornbusch, Fischer, and Startz

Macroeconomics

Thirteenth Edition

Frank

Microeconomics and Behavior

Ninth Edition

ADVANCED ECONOMICS

Romer

Advanced Macroeconomics

Fifth Edition

MONEY AND BANKING

Cecchetti and Schoenholtz

Money, Banking, and Financial Markets

Fifth Edition

URBAN ECONOMICS

O'Sullivan

Urban Economics

Ninth Edition

LABOR ECONOMICS

Borjas

Labor Economics

Eighth Edition

McConnell, Brue, and Macpherson

Contemporary Labor Economics

Eleventh Edition

PUBLIC FINANCE

Rosen and Gayer

Public Finance

Tenth Edition

ENVIRONMENTAL ECONOMICS

Field and Field

Environmental Economics: An Introduction

Seventh Edition

INTERNATIONAL ECONOMICS

Appleyard and Field

International Economics

Ninth Edition

Pugel

International Economics

Seventeenth Edition

To Sarah, Timothy, and Rebecca

Preface to the Eighth Edition

The original motivation for writing *Labor Economics* grew out of my years of teaching labor economics to undergraduates. After trying out many of the textbooks in the market, it seemed to me that students were not being exposed to what the essence of labor economics was about: To try to *understand* how labor markets work. As a result, I felt that students did not really grasp *why* some persons choose to work, while other persons withdraw from the labor market; *why* some firms expand their employment at the same time that other firms are laying off workers; or *why* earnings are distributed unequally.

The key difference between *Labor Economics* and competing textbooks lies in its philosophy. I believe that knowing the *story* of how labor markets work is, in the end, more important than showing off our skills at constructing elegant models of the labor market or remembering hundreds of statistics and institutional details summarizing labor market conditions at a particular point in time.

I doubt that many students will (or should!) remember the mechanics of deriving a labor supply curve or what the unemployment rate was at the peak of the Great Recession 10 or 20 years after they leave college. However, if students could remember the *story* of how the labor market works—and, in particular, that workers and firms respond to changing incentives by altering the amount of labor they supply or demand—the students would be much better prepared to make informed opinions about the many proposed government policies that can have a dramatic impact on labor market opportunities, such as a “workfare” program requiring that welfare recipients work or a payroll tax assessed on employers to fund a national health-care program or a guest worker program that grants tens of thousands of entry visas to high-skill workers. The exposition in this book, therefore, stresses the *ideas* that labor economists use to understand how the labor market works.

The book also makes extensive use of labor market statistics and reports evidence obtained from hundreds of research studies. These data summarize the stylized facts that a good theory of the labor market should be able to explain, as well as help shape our thinking about the way the labor market works. The main objective of the book, therefore, is to survey the field of labor economics with an emphasis on *both* theory and facts. The book relies much more heavily on “the economic way of thinking” than competing textbooks. I believe this approach gives a much better understanding of labor economics than an approach that minimizes or ignores the story-telling aspects of economic theory.

Requirements

The book uses economic analysis throughout. *All* of the theoretical tools are introduced and explained in the text. As a result, the only prerequisite is that the student has some familiarity with the basics of microeconomics, particularly supply and demand curves. The exposure acquired in the typical introductory economics class more than satisfies this prerequisite. All other concepts (such as indifference curves, budget lines, production functions, and isoquants) are motivated, defined, and explained as they appear in our story. The book does not make use of any mathematical skills beyond those taught in high school algebra (particularly the notion of a slope).

Labor economists also make extensive use of econometric analysis in their research. Although the discussion in this book does not require any prior exposure to econometrics, the student will get a much better “feel” for the research findings if they know a little about how labor economists manipulate data to reach their conclusions. The appendix to Chapter 1 provides a simple (and very brief) introduction to econometrics and allows the student to visualize how labor economists conclude, for instance, that winning the lottery reduces labor supply, or that schooling increases earnings. Additional econometric concepts widely used in labor economics—such as the difference-in-differences estimator or instrumental variables—are introduced in the context of policy-relevant examples throughout the text.

Changes in the Eighth Edition

The Eighth Edition offers a thorough rewriting of the entire textbook, making it the most significant revision in quite a few years. As one edition rolls into the next and material gets added to or deleted from the textbook, I think many authors discover that the book keeps moving further away from what the author originally intended. There comes a time when one needs to take a step back, get reacquainted with the entire manuscript free from the pressures of having to get the next edition out the door, take stock of how all the pieces fit together in the context of an ever-evolving field, and do a thorough rethinking of how to best present the material once more as part of a cohesive whole. I experienced that feeling about 3 years ago, shortly after the last edition was published, and decided at the time to tackle the Eighth Edition as if I were writing the textbook for the first time. And that is precisely what I have done.

Readers will find that although much will seem familiar, big chunks of the book have been completely rewritten and streamlined. The book still offers many detailed policy discussions and still uses the evidence reported in state-of-the-art research articles to illustrate the many applications of modern labor economics. The text continues to make frequent use of such econometric tools as fixed effects, the difference-in-differences estimator, and instrumental variables—tools that play a central role in the toolkit of labor economists. And the Eighth Edition even adds to the toolkit by introducing the synthetic control method.

But the text is now much leaner, making it a shorter and easier-to-read book. And it emphasizes, from the very beginning, how these empirical tools are a central part of the methodological revolution that changed labor economics in the past two decades. Empirical analysis must be much more than calculating a correlation describing the relation between two variables. It must instead reflect a well-thought-out strategy that attempts to *identify* the direct consequences of the many shocks that continually hit the labor market.

Among the specific changes in the Eighth Edition are:

1. There are several new extensions of theoretical concepts throughout the book, including a new section on household production (Chapter 2) and on the education production function (Chapter 6). Similarly, there are more detailed discussions of some empirical applications, including the signaling value of the General Equivalency Diploma (GED) and the male–female wage gap in the “gig economy.”
2. The important distinction that empirical labor economics now makes between estimating correlations and identifying consequences from specific labor market shocks is introduced early in the book. Specifically, Chapter 2 has a new section discussing the

- age-old distinction between correlation and causation in the context of evidence from the labor supply literature, which measures the labor supply consequences of winning a lottery or of how taxi drivers are compensated.
3. The section on the employment effects of the minimum wage provides a detailed discussion of the studies that measure the impact of the minimum wage in Seattle, with an illustration of how empirical work in labor economics, particularly when it addresses politically contentious issues, can often lead to wildly different conclusions.
 4. A reorganization of the human capital material in Chapters 6 and 7. Because of the voluminous research on the economics of education, a detailed discussion of the education decision and of how to measure the returns to education now fills up Chapter 6. Chapter 7 continues the study of the human capital model by focusing on postschool investments, on the link between human capital and the wage distribution, and on the determinants of increasing wage inequality. The discussion also introduces the canonical model used in the wage structure literature that uses the Constant Elasticity of Substitution (CES) production function to derive a relative demand curve between high- and low-skill labor. The Mathematical Appendix now includes a detailed derivation of how the model is used to estimate the elasticity of substitution between two labor inputs.
 5. The material on immigration, again one of those topics where the number of studies is growing rapidly, has also been reorganized and tightened. Some users of the earlier edition suggested that because of the intimate link between the wage impact of immigration and the efficiency gains from immigration, the introduction of the immigration surplus should follow immediately after the discussion of the wage impact, and I concur. The immigration material in the geographic mobility chapter now focuses on two issues that are more directly related to the migration decision: The self-selection of immigrants and the assimilation of immigrants in the receiving labor market.

Organization of the Book

The instructor will find that this book is much shorter than competing labor economics textbooks—particularly after the thorough rewriting in the Eighth Edition. The book contains an introductory chapter, plus 11 substantive chapters. If the instructor wished to cover all of the material, each chapter could serve as the basis for about a week’s worth of lectures in a typical undergraduate semester course. Despite the book’s brevity, the instructor will find that all of the key topics in labor economics are covered systematically. The discussion, however, is kept to essentials as I have tried very hard not to deviate into tangential material, or into 10-page-long ruminations on my pet topics.

Chapter 1 presents a brief introduction that exposes the student to the concepts of labor supply, labor demand, and equilibrium. The chapter uses the “real-world” example of the Alaskan labor market during the construction of the oil pipeline to introduce these concepts. In addition, the chapter shows how labor economists contrast the theory with the evidence, as well as discusses the limits of the insights provided by both the theory and the data. The example used to introduce the student to regression analysis is drawn from “real-world” data—and looks at the link between differences in mean wages across occupations and differences in educational attainment as well as the “female-ness” of occupations.

The book begins the detailed analysis of the labor market with a detailed study of labor supply and labor demand. Chapter 2 examines the factors that determine whether a person chooses to work and, if so, how much, while Chapter 3 examines the factors that determine how many workers a firm wants to hire. Chapter 4 puts together the supply decisions of workers with the demand decisions of employers and shows how the labor market “balances out” the conflicting interests of the two parties. These three chapters jointly form the core of the neoclassical approach to labor economics.

The remainder of the book extends and generalizes the basic supply–demand framework. Chapter 5 stresses that jobs differ in their characteristics, so that jobs with unpleasant working conditions may have to offer higher wages in order to attract workers. Chapter 6 stresses that workers are different because they differ in their educational attainment, while Chapter 7 notes that workers also differ in how much on-the-job training they acquire. These investments in human capital help determine the shape of the wage distribution. Chapter 8 describes a key mechanism that allows the labor market to balance out the interests of workers and firms, namely labor turnover and migration.

The final section of the book discusses distortions and imperfections in labor markets. Chapter 9 analyzes how labor market discrimination affects the earnings and employment opportunities of minority workers and women. Chapter 10 discusses how labor unions affect the relationship between the firm and the worker. Chapter 11 notes that employers often find it difficult to monitor the activities of their workers, so that the workers will often want to “shirk” on the job. The chapter discusses how different types of incentive pay systems arise to discourage workers from misbehaving. Finally, Chapter 12 discusses why unemployment can exist and persist in labor markets.

The text uses a number of pedagogical devices designed to deepen the student’s understanding of labor economics. A chapter typically begins by presenting a number of stylized facts about the labor market, such as wage differentials between blacks and whites or between men and women. The chapter then presents the story that labor economists have developed to understand why these facts are observed in the labor market. Finally, the chapter extends and applies the theory to related labor market phenomena. Each chapter typically contains at least one lengthy application of the material to a major policy issue, as well as boxed examples showing the “Theory at Work.”

The end-of-chapter material also contains a number of student-friendly devices. There is a chapter summary describing briefly the main lessons of the chapter; a “Key Concepts” section listing the major concepts introduced in the chapter (when a key concept makes its first appearance, it appears in **boldface**). Each chapter includes “Review Questions” that the student can use to review the major theoretical and empirical issues, a set of 15 problems (many of them brand new) that test the students’ understanding of the material, as well as a list of “Selected Readings” to guide interested students to many of the standard references in a particular area of study.

Supplements for the Book

There are several learning and teaching aids that accompany the eighth edition of *Labor Economics*. These resources are available to instructors for quick download and convenient access via the Instructor Resource material available through McGraw-Hill Connect®.

A *Solutions Manual* and *Test Bank* have been prepared by Robert Lemke of Lake Forest College. The Solutions Manual provides detailed answers to all of the end-of-chapter problems. The comprehensive Test Bank offers over 350 multiple-choice questions in Word and electronic format. Test questions have now been categorized by AACSB learning categories, Bloom's Taxonomy, level of difficulty, and the topic to which they relate. The computerized Test Bank is available through *McGraw-Hill's EZ Test Online*, a flexible and easy-to-use electronic testing program. It accommodates a wide range of question types and you can add your own questions. Multiple versions of the test can be created and any test can be exported for use with course management systems such as Blackboard. The program is available for Windows and Macintosh environments. *PowerPoint Presentations* prepared by Michael Welker of Franciscan University of Steubenville, contain a detailed review of the important concepts presented in each chapter. The slides can be adapted and edited to fit the needs of your course. A *Digital Image Library* is also included, which houses all of the tables and figures featured in this book.

Acknowledgments

I have benefited from countless e-mail messages sent by users of the textbook—both students and instructors. These messages often contained very valuable suggestions, most of which found their way into the Eighth Edition. I strongly encourage users to contact me (gborjas@harvard.edu) with any comments or changes that they would like to see included in the next revision. I am grateful to Robert Lemke of Lake Forest College, who updated the quiz questions for this edition, helped me expand the menu of end-of-chapter problems, and collaborated in and revised the *Solutions Manual* and *Test Bank*; and Michael Welker, Franciscan University of Steubenville, who created the PowerPoint presentation for the Eighth Edition. I am particularly grateful to many friends and colleagues who have generously shared some of their research data so that I could summarize and present it in a relatively simple way throughout the textbook, including Daniel Aaronson, David Autor, William Carrington, Chad Cotti, John Friedman, Barry Hirsch, Lawrence Katz, Alan Krueger, David Lee, Bhashkar Mazumder, and Solomon Polachek. Finally, I have benefited from the countless comments—far too numerous to mention individually—made by many colleagues on the earlier editions.

Contents in Brief

1	Introduction	1	9	Labor Market Discrimination	299
2	Labor Supply	19	10	Labor Unions	341
3	Labor Demand	76	11	Incentive Pay	376
4	Labor Market Equilibrium	122	12	Unemployment	403
5	Compensating Wage Differentials	171		MATHEMATICAL APPENDIX: SOME STANDARD MODELS IN LABOR ECONOMICS	441
6	Education	201		NAME INDEX	453
7	The Wage Distribution	238		SUBJECT INDEX	460
8	Labor Mobility	271			

Contents

Chapter 1

Introduction 1

- 1-1 An Economic Story of the Labor Market 2
- 1-2 The Actors in the Labor Market 3
- 1-3 Why Do We Need a Theory? 7
 - Summary 10
 - Review Questions 10
 - Key Concepts 10

Appendix:

An Introduction to Regression Analysis 11

- Key Concepts 18

Chapter 2

Labor Supply 19

- 2-1 Measuring the Labor Force 20
- 2-2 Basic Facts about Labor Supply 21
- 2-3 The Worker's Preferences 23
- 2-4 The Budget Constraint 28
- 2-5 The Hours of Work Decision 30
- 2-6 To Work or Not to Work? 36
- 2-7 The Labor Supply Curve 39
- 2-8 Estimates of the Labor Supply Elasticity 41
- 2-9 Household Production 44
- 2-10 Correlation versus Causation: Searching for Random Shocks 50
- 2-11 Policy Application: Welfare Programs and Work Incentives 52
- 2-12 Policy Application: The Earned Income Tax Credit 57
- 2-13 Labor Supply over the Life Cycle 61
- 2-14 Policy Application: Disability Benefits and Labor Force Participation 69
 - Theory at Work: Dollars and Dreams* 37
 - Theory at Work: Gaming the EITC* 61
 - Summary 71
 - Key Concepts 71
 - Review Questions 71
 - Problems 72
 - Selected Readings 75

Chapter 3

Labor Demand 76

- 3-1 The Production Function 77
- 3-2 The Short Run 79
- 3-3 The Long Run 85
- 3-4 The Long-Run Demand Curve for Labor 89
- 3-5 The Elasticity of Substitution 95
- 3-6 What Makes Labor Demand Elastic? 96
- 3-7 Factor Demand with Many Inputs 98
- 3-8 Overview of Labor Market Equilibrium 100
- 3-9 Rosie the Riveter as an Instrumental Variable 101
- 3-10 Policy Application: The Minimum Wage 106
 - Theory at Work: California's Overtime Regulations* 94
 - Theory at Work: The Minimum Wage and Drunk Driving* 111
 - Summary 117
 - Key Concepts 118
 - Review Questions 118
 - Problems 118
 - Selected Readings 121

Chapter 4

Labor Market Equilibrium 122

- 4-1 Equilibrium in a Single Labor Market 122
- 4-2 Equilibrium across Labor Markets 125
- 4-3 Policy Application: Payroll Taxes and Subsidies 129
- 4-4 Policy Application: Mandated Benefits 136
- 4-5 The Labor Market Impact of Immigration 139
- 4-6 The Immigration Surplus 150
- 4-7 Policy Application: High-Skill Immigration 152
- 4-8 The Cobweb Model 157
- 4-9 Monopsony 159

Theory at Work: The Intifadah and Palestinian Wages 124

Theory at Work: Hurricanes and the Labor Market 156

Summary 166

Key Concepts 166

Review Questions 166

Problems 167

Selected Readings 170

Chapter 5

Compensating Wage Differentials 171

5-1 The Market for Risky Jobs 172

5-2 The Hedonic Wage Function 178

5-3 Policy Application: How Much Is a Life Worth? 183

5-4 Policy Application: Safety and Health Regulations 186

5-5 Compensating Differentials and Job Amenities 188

5-6 Policy Application: Health Insurance and the Labor Market 192

Theory at Work: Jumpers in Japan 178

Theory at Work: The Value of Life on the Interstate 185

Summary 195

Key Concepts 196

Review Questions 196

Problems 196

Selected Readings 200

Chapter 6

Education 201

6-1 Education in the Labor Market: Some Stylized Facts 202

6-2 Present Value 204

6-3 The Schooling Model 204

6-4 Education and Earnings 210

6-5 Estimating the Rate of Return to School 215

6-6 Policy Application: School Construction in Indonesia 219

6-7 Policy Application: The Education Production Function 220

6-8 Do Workers Maximize Lifetime Earnings? 224

6-9 Signaling 227

Theory at Work: Destiny at Age 6 214

Theory at Work: Booker T. Washington and Julius Rosenwald 223

Summary 233

Key Concepts 233

Review Questions 233

Problems 234

Selected Readings 237

Chapter 7

The Wage Distribution 238

7-1 Postschool Human Capital Investments 239

7-2 On-the-Job Training 239

7-3 The Age–Earnings Profile 244

7-4 Policy Application: Training Programs 248

7-5 Wage Inequality 250

7-6 Measuring Inequality 253

7-7 The Changing Wage Distribution 255

7-8 Policy Application: Why Did Inequality Increase? 257

7-9 Inequality Across Generations 265

Theory at Work: Earnings and Substance Abuse 247

Theory at Work: Computers, Pencils, and the Wage Distribution 262

Summary 267

Key Concepts 267

Review Questions 267

Problems 268

Selected Readings 270

Chapter 8

Labor Mobility 271

8-1 Migration as a Human Capital Investment 271

8-2 Internal Migration 273

8-3 Family Migration 276

8-4 The Self-Selection of Migrants 279

8-5 Immigrant Assimilation 284

8-6 The Job Match and Job Turnover 288

8-7 Job Turnover and the Age–Earnings

Profile 291

Theory at Work: Power Couples 279

Theory at Work: Health Insurance and Job

Lock 291

Summary 293

Key Concepts 294

Review Questions 294

Problems 295

Selected Readings 298

Chapter 9

Labor Market Discrimination 299

9-1 Race and Gender in the Labor Market 299

9-2 The Discrimination Coefficient 301

9-3 Employer Discrimination 303

9-4 Employee Discrimination 309

9-5 Customer Discrimination 310

9-6 Statistical Discrimination 311

9-7 Experimental Evidence 315

9-8 Measuring Discrimination 317

9-9 Policy Application: The Black–White Wage Gap 320

9-10 The Relative Wage of Hispanic and Asians 326

9-11 Policy Application: The Male–Female Wage Gap 328

Theory at Work: Beauty and the Beast 302

Theory at Work: Orchestrating

Impartiality 317

Theory at Work: Shades of Black 327

Summary 333

Key Concepts 334

Review Questions 334

Problems 335

Selected Readings 339

Chapter 10

Labor Unions 341

10-1 A Brief History of American Unions 342

10-2 Determinants of Union Membership 344

10-3 Monopoly Unions 348

10-4 Policy Application: The Efficiency Cost of Unions 350

10-5 Efficient Bargaining 352

10-6 Strikes 358

10-7 Union Wage Effects 363

10-8 Policy Application: Public-Sector Unions 367

Theory at Work: The Rise and Fall of PATCO 348

Theory at Work: The Cost of Labor

Disputes 362

Theory at Work: Occupational

Licensing 367

Summary 369

Key Concepts 370

Review Questions 370

Problems 371

Selected Readings 374

Chapter 11

Incentive Pay 376

11-1 Piece Rates and Time Rates 376

11-2 Tournaments 381

11-3 Policy Application: The Compensation of Executives 385

11-4 Policy Application: Incentive Pay for Teachers 387

11-5 Work Incentives and Delayed Compensation 389

11-6 Efficiency Wages 391

Theory at Work: Windshields by the Piece 381

Theory at Work: How Much Is A Soul Worth? 386

Theory at Work: Did Henry Ford Pay Efficiency Wages? 394

Summary 398

Key Concepts 398

Review Questions 398

Problems 399

Selected Readings 402

Chapter 12

Unemployment 403

12-1 Unemployment in the United States 404

12-2 Types of Unemployment 411

12-3 The Steady-State Rate of Unemployment 412

12-4 Job Search 414

12-5 Policy Application: Unemployment
Compensation 420

12-6 The Intertemporal Substitution
Hypothesis 425

12-7 The Sectoral Shifts Hypothesis 426

12-8 Efficiency Wages and Unemployment 427

12-9 Policy Application: The Phillips Curve 431

*Theory at Work: Graduating During a
Recession 410*

*Theory at Work: Cash Bonuses and
Unemployment 425*

Summary 435

Key Concepts 436

Review Questions 436

Problems 436

Selected Readings 440

**Mathematical Appendix: Some Standard
Models in Labor Economics 441**

Indexes 453

Name Index 453

Subject Index 460

Chapter 1

Introduction

Observations always involve theory.

—*Edwin Hubble*

Most of us will allocate a substantial fraction of our time to the labor market. How we do in the labor market helps determine our wealth, what we can afford to consume, with whom we associate, where we vacation, which schools our children attend, and even who finds us attractive. Not surprisingly, we are all eager to learn how the labor market works. **Labor economics** studies how labor markets work.

Our interest in labor markets, however, is sparked by more than our personal involvement. Many of the central issues in the debate over social policy revolve around the labor market experiences of particular groups of workers or various aspects of the employment relationship between workers and firms. The policy issues examined by modern labor economics include the following:

1. Do welfare programs create work disincentives?
2. What is the impact of immigration on the wage of native-born workers?
3. Do minimum wages increase the unemployment rate of less-skilled workers?
4. What is the impact of occupational safety and health regulations on employment and earnings?
5. Do government subsidies of human capital investments improve the economic well-being of disadvantaged workers?
6. Why did wage inequality in the United States rise so rapidly after 1980?
7. What is the impact of affirmative action programs on the earnings of women and minorities and on the number of women and minorities that firms hire?
8. What is the economic impact of unions, both on their members and on the rest of the economy?
9. Would merit pay for teachers improve the academic achievement of students?
10. Do generous unemployment insurance benefits lengthen the duration of spells of unemployment?

This diverse list of questions clearly illustrates why the study of labor markets is intrinsically more important and more interesting than the study of the market for butter (unless one happens to be in the butter business!). Labor economics helps us understand and address many of the social and economic problems facing modern societies.

1-1 An Economic Story of the Labor Market

This book tells the “story” of how labor markets work. Telling this story involves much more than simply recounting the history and details of labor law or presenting reams of statistics summarizing labor market conditions. Good stories have themes, characters that come alive with vivid personalities, conflicts that have to be resolved, ground rules that limit the set of permissible actions, and events that result inevitably from the interaction among characters.

The story we will tell about the labor market has all these features. Labor economists typically assign motives to the various “actors” in the labor market. Workers, for instance, are trying to find the best possible job and firms are trying to make money. Workers and firms, therefore, enter the labor market with clashing objectives—workers are trying to sell their labor at the highest price and firms are trying to buy labor at the lowest price.

The exchanges between workers and firms are constrained by the ground rules that the government imposes to regulate transactions in the labor market. Changes in these rules and regulations obviously lead to different outcomes. For instance, a minimum wage law prohibits exchanges that pay less than a particular amount per hour worked; occupational safety regulations forbid firms from offering working conditions that are deemed too risky to the worker’s health.

The deals that are struck between workers and firms determine the types of jobs that are offered, the skills that workers acquire, the extent of labor turnover, the structure of unemployment, and the observed earnings distribution. The story thus provides a theory, a framework for understanding, analyzing, and predicting a wide array of labor market outcomes.

The underlying philosophy of the book is that modern economics provides a useful story of how the labor market works. The typical assumptions we make about the behavior of workers and firms, and about the ground rules under which the labor market participants make their transactions, suggest outcomes often corroborated by what we see in real-world labor markets.

The discussion is guided by the belief that learning the story of how labor markets work is as important as knowing basic facts about the labor market. The study of facts without theory is just as empty as the study of theory without facts. Without understanding how labor markets work—that is, without having a theory of why workers and firms pursue some employment relationships and avoid others—we would be hard-pressed to predict the labor market impact of changes in government policies or of changes in the demographic composition of the workforce.

A question often asked is which are more important—ideas or facts? This book stresses that “ideas *about* facts” are most important. We do not study labor economics so that we can construct elegant mathematical theories or to remember that the unemployment rate was 6.9 percent in 1993. Rather, we want to identify which economic and social factors generate a certain level of unemployment, and why.

The main objective of this book is to survey the field of labor economics with an emphasis on *both* theory and facts: Where the theory helps us understand how the facts are generated and where the facts can help shape our thinking about the way labor markets work.

1-2 The Actors in the Labor Market

Throughout the book, we will see that there are three leading actors in our story: workers, firms, and the government.¹

As workers, we receive top casting. Without us, after all, there is no “labor” in the labor market. We decide whether to work or not, how many hours to work, how hard to work, which skills to acquire, when to quit a job, which occupations to enter, and whether to join a labor union.

Each of these decisions is driven by the desire to *optimize*, to choose the best available option from the various choices. In our story, workers will always act in ways that maximize their well-being. Adding up the decisions of millions of workers generates the economy’s labor supply in terms of the number of persons seeking work, and also in terms of the quantity and quality of skills available to employers. As we will see throughout the book, persons who want to maximize their well-being tend to supply more time and more effort to those activities that have a higher payoff. The **labor supply curve**, therefore, is often upward sloping, as illustrated in Figure 1-1.

The hypothetical labor supply curve drawn in Figure 1-1 gives the number of engineers that will be forthcoming at every wage. For example, 20,000 workers are willing to supply their services to engineering firms if the engineering wage is \$40,000 per year. If the engineering wage rises to \$50,000, then 30,000 workers will choose to be engineers. In other words, as the engineering wage rises, more persons decide that the engineering profession is a worthwhile pursuit. More generally, the labor supply curve relates the number of person-hours supplied to the economy to the wage that is being offered. The higher the wage that is being offered, the larger the labor supplied.


Firms co-star in our story. Each firm must decide how many and which types of workers to hire and fire, the length of the workweek, how much capital to employ, and whether to offer a safe or risky working environment to its workers. Firms also have motives. We assume that firms want to maximize profits. From the firm’s point of view, the consumer is king. The firm will maximize its profits by making the production decisions—and hence the hiring and firing decisions—that best serve the consumers’ needs. In effect, the firm’s demand for labor is a **derived demand**, a demand derived from the desires of consumers.

Adding up the hiring and firing decisions of millions of employers generates the economy’s labor demand. The assumption that firms want to maximize profits implies that firms will want to hire many workers when labor is cheap but will refrain from hiring

¹ A fourth actor, trade unions, may have to be added in some countries. Unions may organize a large fraction of the workforce and represent the interests of workers in their bargaining with employers. In the United States, however, the trade union movement has been in decline for several decades. By 2016, only 6.4 percent of private-sector workers were union members.

FIGURE 1-1 Supply and Demand in the Engineering Labor Market

The labor supply curve gives the number of persons willing to supply their services to engineering firms at a given wage. The labor demand curve gives the number of engineers that firms will hire at that wage. Equilibrium occurs where supply equals demand, so that 20,000 engineers are hired at a wage of \$40,000.



when labor is expensive. The relation between the price of labor and how many workers firms are willing to hire is summarized by the downward-sloping **labor demand curve** in Figure 1-1. As drawn, the labor demand curve tells us that firms in the engineering industry want to hire 20,000 engineers when the wage is \$40,000 but will hire only 10,000 engineers if the wage rises to \$50,000.

Workers and firms, therefore, enter the labor market with conflicting interests. Many workers are willing to supply their services when the wage is high, but few firms are willing to hire them. Conversely, few workers are willing to supply their services when the wage is low, but many firms are looking for workers. As workers search for jobs and firms search for workers, these conflicting desires are “balanced out” and the labor market reaches an **equilibrium**. In a free-market economy, equilibrium is attained when supply equals demand.

As drawn in Figure 1-1, the equilibrium wage is \$40,000 and 20,000 engineers will be hired in the labor market. This wage–employment combination is an equilibrium because it balances out the conflicting desires of workers and firms. Suppose, for example, that the engineering wage was \$50,000—above equilibrium. Firms would then want to hire only 10,000 engineers, even though 30,000 engineers are looking for work. The excess number of job applicants would bid down the wage as they compete for the few jobs available. Suppose, instead, that the wage was \$30,000—below equilibrium. Because engineers are cheap, firms want to hire 30,000 engineers, but only 10,000 engineers are willing to work at that wage. As firms compete for the few available engineers, they bid up the wage.

There is one last major player in the labor market, the government. The government can tax the worker’s earnings, subsidize the training of engineers, impose a payroll tax on

firms, demand that the racial and gender composition of engineers hired by firms exactly reflect the composition of the population, enact legislation that makes some labor market transactions illegal (such as paying engineers less than \$50,000 annually), and increase the supply of engineers by encouraging their immigration from abroad. All these actions will change the equilibrium that will eventually be attained in the labor market.

The Trans-Alaska Oil Pipeline

In January 1968, oil was discovered in Prudhoe Bay in remote northern Alaska. The oil reserves were estimated to be greater than 10 billion barrels, making it the largest such discovery in North America.²

There was one problem with the discovery—the oil was located in a remote and frigid area of Alaska, far from where most consumers lived. To solve the daunting problem of transporting the oil to those consumers who wanted to buy it, the oil companies proposed building a 48-inch pipeline across the 789-mile stretch from northern Alaska to the southern (and ice-free) port of Valdez. At Valdez, the oil would be transferred to oil supertankers. These huge ships would then deliver the oil to consumers in the United States and elsewhere.

The oil companies joined forces and formed the Alyeska Pipeline Project. The construction project began in the spring of 1974, after Congress gave its approval in the wake of the 1973 oil embargo. Construction work continued for 3 years and the pipeline was completed in 1977. Alyeska employed about 25,000 workers during the summers of 1974 through 1977, and its subcontractors employed an additional 25,000 workers. Once the pipeline was built, Alyeska reduced its pipeline-related employment to a small maintenance crew.

Many of the workers employed by Alyeska and its subcontractors were engineers who had built pipelines across the world. Very few of those engineers were resident Alaskans. The remainder of the Alyeska workforce consisted of relatively low-skill labor such as truck drivers and excavators. Many of the low-skill workers were resident Alaskans.

The theoretical framework summarized by the supply and demand curves can help us understand the shifts that *should* have occurred in the Alaskan labor market as a result of the Trans-Alaska Pipeline System. As Figure 1-2 shows, the labor market was initially in an equilibrium represented by the intersection of the demand curve D_0 and the supply curve S_0 . A total of E_0 Alaskans were employed at a wage of w_0 in the initial equilibrium.


The construction project clearly led to a sizable increase in the demand for labor. Figure 1-2 illustrates this shift by showing the demand curve moving outward from D_0 to D_1 . The outward shift in the demand curve implies that—at any given wage—Alaskan employers were looking for more workers.

The shift in demand should have moved the Alaskan labor market to a new equilibrium, represented by the intersection of the new demand curve and the original supply curve. At this new equilibrium, a total of E_1 persons were employed at a wage of w_1 . The theory, therefore, predicts that the pipeline construction project should have increased *both* wages

² The discussion is based on William J. Carrington, “The Alaskan Labor Market during the Pipeline Era,” *Journal of Political Economy* 104 (February 1996): 186–218.

FIGURE 1-2 The Alaskan Labor Market and the Construction of the Oil Pipeline

The construction of the oil pipeline shifted the labor demand curve in Alaska from D_0 to D_1 , resulting in higher wages and employment. Once the pipeline was completed, the demand curve reverted back to its original level and wages and employment fell.




and employment. As soon as the project was completed, however, and the temporary need for additional workers disappeared, the demand curve would have shifted back to its original position at D_0 . In the end, the wage should have gone back down to w_0 and E_0 workers would be employed. In short, the pipeline construction project should have led to a temporary increase in both wages and employment during the construction period.

Figure 1-3 shows what *actually* happened to employment and wages in Alaska between 1968 and 1983. Because the state's population was growing steadily for some decades, total employment was rising steadily even before the oil discovery in Prudhoe Bay. The data clearly show, however, that employment “spiked” in 1975, 1976, and 1977 and then went back to its long-run growth trend in 1977. The earnings of Alaskan workers also increased during the relevant period. After adjusting for inflation, monthly earnings rose from an average of \$2,648 in the third quarter of 1973 to \$4,140 in the third quarter of 1976, a surge of 56 percent. By 1979, real earnings were back to the level observed prior to the beginning of the pipeline construction project.

It turns out that the temporary increase in labor supply occurred for two distinct reasons. First, a larger fraction of Alaskans were willing to work when the wage increased. In the summer of 1973, about 39 percent of Alaskans worked. In the summers of 1975 and 1976, about 50 percent of Alaskans worked. Second, the rate of population growth in Alaska accelerated between 1974 and 1976, as workers living in the lower 48 states moved to Alaska to take advantage of the improved economic opportunities (despite the frigid weather conditions there). The increase in the rate of population growth, however, was temporary. Population growth reverted back to its long-run trend soon after the pipeline construction project was completed.

FIGURE 1-3
Wages and
Employment
in the Alaskan
Labor Market,
1968–1984

Source: William J. Carrington, "The Alaskan Labor Market during the Pipeline Era," *Journal of Political Economy* 104 (February 1996): 199.



1-3 Why Do We Need a Theory?

We have just told a simple story of how the Trans-Alaska Pipeline System affected labor market outcomes in Alaska—and how each of the actors in our story played a major role. The government approved the pipeline project despite the potential environmental hazards involved; firms that saw income opportunities in building the pipeline increased their demand for labor; and workers responded to the change in demand by increasing the quantity of labor supplied to the Alaskan labor market.

We have, in effect, constructed a theory or **model** of the Alaskan labor market. Our model is characterized by an upward-sloping labor supply curve, a downward-sloping labor demand curve, and the assumption that an equilibrium is eventually attained that resolves the conflicts between workers and firms. This model predicts that the construction of the pipeline would temporarily increase wages and employment in the Alaskan labor market. Moreover, this prediction is testable—that is, the predictions about wages and employment can be compared with what actually happened. It turns out that the supply–demand model passes the test; the data confirm the theoretical predictions.

Needless to say, the model of the labor market illustrated in Figure 1-2 does not do full justice to the complexities of the Alaskan labor market. It is easy to come up with many variables that our simple model ignored and that could potentially change our predictions. For instance, it is possible that workers care about more than just the wage when they make labor supply decisions. The opportunity to participate in such a challenging or cutting-edge project as the construction of the Trans-Alaska Pipeline could have attracted engineers at wages lower than those offered by firms engaged in more mundane projects—despite the harsh working conditions in the field. The theoretical prediction that the construction of the pipeline project would increase wages would then be incorrect because the project could have attracted more workers at lower wages.

If the factors that we omitted from our theory play a crucial role in understanding how the Alaskan labor market operates, we might be wrongly predicting that wages and employment would rise. If these factors are only minor details, however, our model captures the essence of what goes on in the Alaskan labor market and our prediction would be valid.

We could try to build a more complex model, a model that incorporates every single one of the omitted factors. Now *that* would be a tough job! A completely realistic model would have to describe how millions of workers and firms interact and how these interactions work themselves through the labor market. Even if we knew how to accomplish such a difficult task, this “everything-but-the-kitchen-sink” approach defeats the whole purpose of having a theory. A theory that mirrored the real-world labor market in Alaska down to the minutest detail might indeed be able to explain all the facts, but it would be as complex as reality itself, cumbersome and incoherent, and would not really help us understand how the Alaskan labor market works.

There has been a long debate over whether a theory should be judged by the realism of its assumptions or by the extent to which it helps us understand and predict the labor market phenomena we are interested in. We obviously have a better shot at predicting correctly if we use more realistic assumptions. At the same time, a theory that mirrors the world too closely is too clumsy and does not isolate what *really* matters. The “art” of labor economics lies in choosing which details are essential to the story and which details are not. There is a tradeoff between realism and simplicity, and good economics hits the mark just right.

As we will see throughout this book, the supply–demand framework in Figure 1-1 helps to isolate the key factors that motivate the various actors in the labor market. The model provides a useful way of organizing our thoughts about how the labor market works. It also gives a solid foundation for building more complex and more realistic models. And, most important, the model works. Its predictions are often consistent with what is observed in the real world.

The supply–demand framework predicts that the construction of the Alaska oil pipeline would temporarily increase employment and wages in the Alaskan labor market. This prediction is an example of **positive economics**. Positive economics addresses the relatively narrow “What is?” questions, such as, What is the impact of the discovery of oil in Prudhoe Bay, and the subsequent construction of the oil pipeline, on the Alaskan labor market?

Positive economics, therefore, addresses questions that can, in principle, be answered with the tools of economics, without interjecting any value judgment as to whether the particular outcome is desirable or harmful. This book is devoted to the analysis of such positive questions as: What is the impact of the minimum wage on unemployment? What is the impact of immigration on the earnings of native-born workers? What is the impact of a tuition assistance program on college enrollment rates? What is the impact of unemployment insurance on the duration of a spell of unemployment?

These positive questions, however, beg many important issues. In fact, some would say that these positive questions beg *the* most important issues: *Should* the oil pipeline have been built? *Should* there be a minimum wage? *Should* the government subsidize college tuition? *Should* the United States accept more immigrants? *Should* the unemployment insurance system be less generous?

These questions fall in the realm of **normative economics**, which addresses much broader “What should be?” questions. By their nature, the answers to these normative questions require value judgments. Because each of us probably has different values, our answers to these normative questions may differ *regardless* of what the theory or the facts

tell us about the economic impact of the oil pipeline, the employment effects of the minimum wage, or the impact of immigration on the well-being of native workers.

Normative questions force us to make value judgments about the type of society we wish to live in. Consider, for instance, the impact of immigration on a particular host country. As we will see, the supply–demand framework implies that an increase in the number of immigrants lowers the income of competing workers but raises the income of the firms that hire those workers by even more. On net, therefore, the receiving country gains. Moreover, because immigration is typically a voluntary supply decision, it also makes the immigrants better off.

Suppose, in fact, that the evidence for a particular host country was consistent with the model’s predictions. In particular, the immigration of 10 million workers improved the well-being of the immigrants (relative to their well-being in their country of birth); reduced the income of native workers by \$25 billion annually; and increased the income of employers by \$40 billion. Let’s now ask a normative question: *Should* the country admit 10 million more immigrants?

This normative question cannot be answered solely on the basis of the theory or the facts. Even though total income in the host country has increased by \$15 billion, there also has been a redistribution of wealth. Some persons are worse off and others are better off.

To answer the question of whether the country should continue to admit immigrants, one has to decide whose economic welfare we should care most about: that of immigrants, who are made better off; that of native workers, who are made worse off; or that of employers, who are made better off. One might even bring into the discussion the well-being of the people left behind in the source countries, who are clearly affected by the emigration of their compatriots. It is clear that any resolution of this issue requires clearly stated assumptions about what constitutes the “national interest,” about who matters more.

Many economists often take a fallback position when these types of problems are encountered. Because the immigration of 10 million workers increases the *total* income in the destination country by \$15 billion, it is then possible to redistribute income so that every person in that country is made better off. A policy that can *potentially* improve the well-being of everyone in the economy is said to be “efficient”; it increases the size of the economic pie available to the country. The problem, however, is that this type of redistribution seldom occurs in the real world; the winners typically remain winners and the losers remain losers. Our answer to a normative question, therefore, forces us to confront the tradeoff between efficiency and distributional issues.

As a second example, we will see that the supply–demand framework predicts that unionization transfers wealth from firms to workers, but that unionization also shrinks the size of the economic pie. Suppose that the facts unambiguously support these theoretical implications, unions increase the total income of workers by, say, \$40 billion, but the country as a whole is poorer by \$20 billion. Let’s now ask a normative question: *Should* the government pursue policies that discourage workers from forming labor unions?

Our answer to this normative question again depends on how we balance the gains to the unionized workers with the losses to the employers who must pay higher wages and to the consumers who must pay higher prices for union-produced goods.

The lesson should be clear. As long as there are winners and losers—and government policies inevitably leave winners and losers in their wake—neither the theoretical implications of economic models nor the facts are sufficient to answer the normative question of whether a particular policy is desirable.

Despite the fact that economists cannot answer what many would consider to be the “big questions,” there is an important sense in which framing and answering positive questions is crucial for any policy discussion. Positive economics tells us how particular government policies affect the well-being of different segments of society. Who are the winners, and how much do they gain? Who are the losers, and how much do they lose?

In the end, any informed policy discussion requires that we be fully aware of the price that has to be paid when making particular choices. The normative conclusion that one might reach may well depends on the magnitude of the costs and benefits associated with a particular policy. For example, the redistributive impact of unions (that is, the transfer of income from firms to workers) could easily dominate the normative discussion if unions generated only a small decrease in the size of the economic pie. The distributional impact, however, might be less relevant if unions greatly reduced the size of the economic pie.

Summary

- Labor economics studies how labor markets work. Topics addressed by labor economics include the determination of the income distribution, the economic impact of unions, the allocation of a worker’s time to the labor market, the hiring and firing decisions of firms, labor market discrimination, the determinants of unemployment, and the worker’s decision to invest in human capital.
- Models in labor economics typically contain three actors: workers, firms, and the government. It is typically assumed that workers maximize their well-being and that firms maximize profits. Governments influence the decisions of workers and firms by imposing taxes, granting subsidies, and regulating the “rules of the game” in the labor market.
- A good theory of the labor market should have realistic assumptions, should not be clumsy or overly complex, and should provide empirical implications that can be tested with real-world data.
- The tools of economics are helpful for answering positive questions. The answer to a normative question, however, typically requires that we impose a value judgment on the desirability of particular economic outcomes.

Review Questions

1. What is labor economics? Which types of questions do labor economists analyze?
2. Who are the key actors in the labor market? What motives do economists typically assign to workers and firms?
3. Why do we need a theory to understand real-world labor market problems?
4. What is the difference between positive and normative economics? Why are positive questions easier to answer than normative questions?

Key Concepts

derived demand, 3	labor economics, 1	normative economics, 8
equilibrium, 4	labor supply curve, 3	positive economics, 8
labor demand curve, 4	model, 7	

Appendix

An Introduction to Regression Analysis

Labor economics is an empirical science. It makes extensive use of **econometrics**, the application of statistical techniques to study relationships in economic data. For example, we will be addressing such questions as

1. Do higher levels of unemployment benefits lead to longer spells of unemployment?
2. Do higher levels of welfare benefits reduce work incentives?
3. Does going to school one more year increase a worker's earnings?

The answers to these three questions ultimately depend on a correlation between pairs of variables: the level of unemployment benefits and the duration of unemployment spells; the level of welfare benefits and labor supply; and educational attainment and wages. We also will want to know not only the *sign* of the correlation, but the *size* as well. In other words, by how many weeks does a \$50 increase in unemployment benefits lengthen the duration of unemployment spells? By how many hours does an increase of \$200 per month in welfare benefits reduce labor supply? And by how much do our earnings increase if we get a college education?

Although this book does not use the technical details of econometric analysis in the discussion, the student can better appreciate both the usefulness *and* the limits of empirical research by knowing how labor economists manipulate the available data to answer the questions we are interested in. The main statistical technique used by labor economists is **regression analysis**.

An Example

There are sizable wage differences across occupations. We are interested in determining why some occupations pay more than others. One obvious factor that determines the average wage in an occupation is the level of education of workers in that occupation.

It is common in labor economics to conduct empirical studies of earnings by looking at the logarithm of earnings, rather than the actual level of earnings. There are sound theoretical and empirical reasons for this practice, one of which will be described shortly. Suppose there is a linear equation relating the average log wage in an occupation ($\log w$) to the mean years of schooling of workers in that occupation (s). We write this line as


$$\log w = \alpha + \beta s \quad (1-1)$$

The variable on the left-hand side—the average log wage in the occupation—is called the **dependent variable**. The variable on the right-hand side—average years of schooling in the occupation—is called the **independent variable**. The main objective of regression analysis is to obtain numerical estimates of the coefficients α and β by using actual data on the mean log wage and mean schooling in each occupation. It is useful, therefore, to spend some time interpreting these **regression coefficients**.

Equation (1-1) traces out a line, with intercept α and slope β ; this line is drawn in Figure 1-4. As drawn, the regression line makes the sensible assumption that the slope β is

FIGURE 1-4 The Regression Line

The regression line gives the relationship between the average log wage rate and the average years of schooling of workers across occupations. The slope of the regression line gives the change in the log wage resulting from a one-year increase in years of schooling. The intercept gives the log wage for an occupation where workers have zero years of schooling.



positive, so wages are higher in occupations where the typical worker is better educated. The intercept α gives the log wage that would be observed in an occupation where workers have zero years of schooling. Elementary algebra teaches us that the slope of a line is given by the change in the vertical axis divided by the corresponding change in the horizontal axis or

$$\beta = \frac{\text{Change in log wage}}{\text{Change in years of schooling}} \quad (1-2)$$

Put differently, the slope β gives the change in the log wage associated with a one-year increase in schooling. It turns out that a small change in the log wage approximates the percent change in the wage. For example, if the difference in the mean log wage between two occupations is 0.051, we can say that there is approximately a 5.1 percent wage difference between the two occupations. This property is one of the reasons why labor economists typically conduct studies of salaries using the logarithm of the wage; they can then interpret changes in this quantity as a percent change in the wage. This mathematical property of logarithms implies that the coefficient β can be interpreted as giving the percent change in earnings resulting from one more year of education.


To estimate the parameters α and β , we first need to obtain data on the average log wage and average years of schooling by occupation. These data can be easily calculated using the Annual Social and Economic Supplement of the Current Population Surveys (CPS). These data, collected in March of every year by the Bureau of Labor Statistics, report employment conditions and salaries for tens of thousands of workers. One can use the data to compute the average log hourly wage and the average years of schooling for men working in each of 45 different occupations. The resulting data are reported in Table 1-1. The typical male engineer had a log wage of 3.37 and 15.8 years of schooling. In contrast, the typical construction laborer had a log wage of 2.44 and 10.5 years of schooling.

TABLE 1-1 Characteristics of Occupations, 2001

Source: Annual Demographic Files of the Current Population Survey, 2002.

Occupation	Mean Log Hourly Wage of Male Workers	Mean Years of Schooling for Male Workers	Female Share (%)
Administrators and officials, public administration	3.24	15.7	52.4
Other executives, administrators, and managers	3.29	14.9	42.0
Management-related occupations	3.16	15.4	59.4
Engineers	3.37	15.8	10.7
Mathematical and computer scientists	3.36	15.6	32.2
Natural scientists	3.22	17.4	34.2
Health diagnosing occupations	3.91	19.8	31.2
Health assessment and treating occupations	3.23	16.2	86.2
Teachers, college and university	3.17	18.8	44.7
Teachers, except college and university	2.92	16.5	75.8
Lawyers and judges	3.72	19.7	29.3
Other professional specialty occupations	2.90	15.9	54.0
Health technologists and technicians	2.76	14.2	83.1
Engineering and science technicians	2.97	13.8	26.0
Technicians, except health, engineering, and science	3.30	15.4	48.5
Supervisors and proprietors, sales occupations	2.96	13.9	37.6
Sales representatives, finance and business services	3.39	15.1	44.7
Sales representatives, commodities, except retail	3.14	14.4	25.4
Sales workers, retail and personal services	2.61	13.4	64.0
Sales-related occupations	2.93	14.8	72.4
Supervisors, administrative support	2.94	13.8	61.2
Computer equipment operators	2.91	13.8	57.1
Secretaries, stenographers, and typists	2.75	13.8	98.0
Financial records, processing occupations	2.67	14.2	92.9
Mail and message distributing	2.87	13.2	41.9
Other administrative support occupations, including clerical	2.66	13.4	79.2
Private household service occupations	2.46	10.6	96.0
Protective service occupations	2.80	13.6	18.7
Food service occupations	2.23	11.4	60.0
Health service occupations	2.38	13.2	89.1
Cleaning and building service occupations	2.37	11.2	48.2
Personal service occupations	2.55	13.4	80.4
Mechanics and repairers	2.81	12.6	5.2
Construction trades	2.74	11.9	2.4
Other precision production occupations	2.82	12.3	22.5
Machine operators and tenders, except precision	2.62	11.8	35.2
Fabricators, assemblers, inspectors, and samplers	2.65	12.0	36.2
Motor vehicle operators	2.59	12.1	12.7
Other transportation occupations and material moving	2.68	11.8	6.3
Construction laborer	2.44	10.5	3.9
Freight, stock, and material handlers	2.44	12.0	30.4
Other handlers, equipment cleaners, and laborers	2.42	11.3	28.0
Farm operators and managers	2.52	12.9	20.5
Farm workers and related occupations	2.29	9.9	18.5
Forestry and fishing occupations	2.70	12.0	3.7

FIGURE 1-5 Scatter Diagram Relating Wages and Schooling by Occupation, 2001




The plot of these data in Figure 1-5 is called a **scatter diagram** and describes the relation found between the average log wage and the average years of schooling in the real world. The relation between the two variables does not look anything like the regression line that we hypothesized. Instead, it is a scatter of points. But the points are not randomly scattered on the page; they have a noticeable upward-sloping drift. The raw data, therefore, suggest a positive correlation between wages and schooling, but nothing as simple as an upward-sloping line.

We have to recognize, however, that education is not the only factor that determines the average wage in an occupation. There is probably a great deal of error when workers report their salary to the Bureau of Labor Statistics. This measurement error disperses the points on a scatter diagram away from the line that we believe represents the “true” data. There also might be other factors that affect average earnings in an occupation, such as the average age of the workers or perhaps a variable indicating the “female-ness” of the occupation. It is often argued that jobs that are predominantly done by men (for example, welders) tend to pay more than jobs that are predominantly done by women (for example, kindergarten teachers). All of these factors would again disperse our data points away from the line.

The objective of regression analysis is to find the *best* line that goes through the scatter diagram. Figure 1-6 redraws our scatter diagram and inserts a few of the many lines that we could draw through the scatter. Line A does not represent the general trend very well; after all, the raw data suggest a positive correlation between wages and education, yet line A

FIGURE 1-6 Choosing among Lines Describing the Trend in the Data

There are many lines that can be drawn through the scatter diagram. Lines *A*, *B*, and *C* provide three such examples. None of these lines “fit” the trend in the scatter diagram very well.



has a negative slope. Both lines *B* and *C* are upward sloping, but they are both a bit “off”; line *B* lies above all of the points in the scatter diagram and line *C* is too far to the right.

The **regression line** is the line that best summarizes the data.³ The formula that calculates the regression line is included in every statistics and spreadsheet software program. If we apply the formula to the data in our example, we obtain the regression line.


$$\log w = 0.869 + 0.143s \quad (1-3)$$

This estimated regression line is superimposed on the scatter diagram in Figure 1-7. We interpret the regression line reported in equation (1-3) as follows. The estimated slope is positive, indicating that the average log wage is indeed higher in occupations where workers are more educated. The 0.143 slope implies that each one-year increase in the mean schooling of workers in an occupation raises the wage by approximately 14.3 percent.

The intercept indicates that the log wage would be 0.869 in an occupation where the average worker had zero years of schooling. We have to be very careful when we use this result. After all, no occupation has a workforce with zero years of schooling. In fact, the

³ More precisely, the regression line is the line that minimizes the sum of the square of the vertical differences between every point in the scatter diagram and the corresponding point on the line. This method of estimating the regression line is called *least squares*.

FIGURE 1-7 The Scatter Diagram and the Regression Line



smallest value of s is 9.9 years. The intercept is obtained by extrapolating the regression line to the left until it hits the vertical axis. In other words, we are using the regression line to make an out-of-sample prediction. It is easy to get absurd results when we do this type of extrapolation: After all, what does it mean to say that the typical person in an occupation has no schooling whatsoever? An equally silly extrapolation takes the regression line and extends it to the right until, say, we wish to predict what would happen if the average worker had 25 years of schooling. Put simply, it is problematic to predict outcomes that lie outside the range of the data.

“Margin of Error” and Statistical Significance

If we plug the data reported in Table 1-1 into a statistics or spreadsheet program, we will find that the program reports many more numbers than just the intercept and the slope of a regression line. The program also reports what are called **standard errors**, or a measure of the statistical precision with which the coefficients are estimated. When poll results are reported in the media, we often hear, for instance, that 52 percent of the population believes that tomatoes should be bigger and redder, with a margin of error of ± 3 percent. We use standard errors to calculate the margin of error for our estimated regression coefficients.

In our data, it turns out that the standard error for the intercept α is 0.172 and that the standard error for the slope β is 0.012. *The margin of error that is used commonly in econometric work is twice the standard error.* We can then say that a one-year increase in average

schooling increases the log wage by 0.143 ± 0.024 (or twice the standard error of 0.012). In other words, our data suggest that a one-year increase in schooling increases the average wage in an occupation by as little as 11.9 percent or by as much as 16.7 percent. Statistical theory tells us that the *true* impact of the one-year increase in schooling lies within this range with a 95 percent probability.

The regression program will also report a ***t* statistic** for each regression coefficient. The *t* statistic helps us assess the **statistical significance** of the estimated coefficients. The *t* statistic is defined as

$$t \text{ statistic} = \frac{\text{Absolute value of regression coefficient}}{\text{Standard error of regression coefficient}} \quad (1-4)$$

If a regression coefficient has a *t* statistic above the “magic” number of 2, the regression coefficient is said to be significantly different from zero. In other words, it is very likely that the true value of the coefficient is not zero, so there is some correlation between the two variables that we are interested in. If a *t* statistic is below 2, the coefficient is said to be insignificantly different from zero, so we cannot conclude that there is a correlation between the two variables of interest.

Note that the *t* statistic associated with our estimated slope is 11.9 (or $0.143 \div 0.012$), which is certainly above 2. Our estimate of the slope is significantly different from zero. It is extremely likely that there is indeed a positive correlation between the average log wage in an occupation and the average schooling of workers.

Finally, the statistical software will also report a number called the ***R*-squared**. This statistic gives the fraction of the dispersion in the dependent variable that is “explained” by the dispersion in the independent variable. The *R*-squared of the regression reported in equation (1-3) is 0.762. In other words, 76.2 percent of the variation in the mean log wage across occupations can be attributed to differences in educational attainment across the occupations. Put differently, our very simple regression model seems to do a very good job at explaining why engineers earn more than construction laborers—it is largely because one group of workers has a lot more education than the other.

Multiple Regression

Up to this point, the regression model contains only one independent variable, mean years of schooling. As noted above, the average log wage of men in an occupation will depend on many other factors. The simple correlation between wages and schooling implied by the regression model in equation (1-3) could be confounding the effect of some of these other variables. To isolate the relationship between the log wage and schooling (and avoid what is called “omitted variable bias”), it is important to control for other variables that also might generate wage differences across occupations.

Suppose we believe that occupations that are predominantly held by men tend to pay more—for given schooling—than occupations that are predominantly held by women. We can then write an expanded regression model as

$$\log w = \alpha + \beta s + \gamma p \quad (1-5)$$

where the variable *p* gives the percent of workers in an occupation that are women.

We now wish to interpret the coefficients in this multiple regression model—a regression that contains more than one independent variable. Each coefficient in the **multiple regression** measures the impact of a particular variable on the log wage, *other things being equal*. For instance, the coefficient β gives the change in the log wage resulting from a one-year increase in mean schooling, holding constant the relative number of women in the occupation. Similarly, the coefficient γ gives the change in the log wage resulting from a one-percentage-point increase in the share of female workers, holding constant the average schooling of the occupation. Finally, the intercept α gives the log wage in a fictional occupation that employs only men and where the typical worker has zero years of schooling.

The last column in Table 1-1 reports the values of the female share p for the occupations in our sample. The representation of women varies significantly across occupations: 75.8 percent of teachers below the university level are women, as compared to only 5.2 percent of mechanics and repairers.

Because we now have two independent variables, our scatter diagram is three dimensional. The regression “line” is now the plane that best fits the data in this three-dimensional space. If we plug these data into a computer program to estimate the regression model in equation (1-5), the estimated regression line is given by

$$\log w = 0.924 + 0.150s - 0.003p \qquad R\text{-squared} = 0.816 \qquad (1-6)$$

(0.154) (0.011) (0.001)

where the standard error of each of the coefficients is reported in parentheses below the coefficient.

A one-year increase in the occupation’s mean schooling raises weekly earnings by approximately 15.0 percent. In other words, if we compare two occupations that have the same female share but differ in years of schooling by one year, workers in the better educated occupation earn 15 percent more.

We also find that the female share of the occupation has a statistically significant negative impact on the log wage. In other words, men who work in predominantly female occupations earn less than men who work in predominantly male occupations—even if both occupations have the same mean schooling. The regression coefficient, in fact, implies that a 10-percentage-point increase in the female share lowers the average earnings of an occupation by 3.0 percent.

The multiple regression model can, of course, be expanded to incorporate many more independent variables. As we will see throughout this book, labor economists put a lot of effort into defining and estimating regression models that isolate the correlation between the two variables of interest *after controlling for all other relevant factors*. Regardless of how many independent variables are included in the regression, however, all the regression models are estimated in essentially the same way: The regression line best summarizes the trends in the underlying data.

Key
Concepts

dependent variable, 11	regression analysis, 11	scatter diagram, 14
econometrics, 11	regression coefficients, 11	statistical significance, 17
independent variable, 11	regression line, 15	standard errors, 16
multiple regression, 18	R-squared, 17	t statistic, 17

Chapter 2

Labor Supply

It's true hard work never killed anybody, but I figure, why take the chance?

—Ronald Reagan

Each of us must decide whether to work and, once employed, how many hours to work. At any point in time, the aggregate labor supply in the economy is given by adding the work choices made by all person in the population.

The economic and social consequences of these decisions vary dramatically over time. In 1948, 84 percent of American men and 31 percent of American women aged 16 or over worked. By 2017, the proportion of working men had declined to 66 percent, whereas the proportion of working women had risen to 55 percent. Similarly, the length of the average workweek in manufacturing fell from 55 to 42 hours over the past century.¹ These labor supply trends have surely altered the nature of the American family as well as greatly affected the economy's productive capacity.

This chapter develops the framework economists use to study labor supply decisions. In this framework, individuals seek to maximize their well-being by consuming goods (such as fancy cars and nice homes) and leisure. Goods have to be purchased in the marketplace. Because most of us are not independently wealthy, we must work in order to earn the cash required to buy the desired goods. The economic trade-off is clear: If we do not work, we can consume a lot of leisure, but we have to do without the goods and services that make life more enjoyable. If we do work, we will be able to afford many of these goods and services, but we must give up some of our valuable leisure time.

The economic model of labor–leisure choice isolates the person's wage rate and income as the key variables that guide the allocation of time between the labor market and leisure activities. In this chapter, we initially use the framework to analyze “static” labor supply decisions, the factors that determine a person's labor supply at a point in time. We then extend the basic model to explore how the work decision changes as a person ages.

This economic framework not only helps us understand why women's work propensities rose and hours of work declined, but also allows us to address a number of questions with important policy implications. For example, do welfare programs reduce incentives to work? Or do cuts in the income tax rate increase hours of work?

¹ The Bureau of Labor Statistics website contains a vast collection of employment statistics; see www.bls.gov/data/home.htm.

2-1 Measuring the Labor Force

On the first Friday of every month, the Bureau of Labor Statistics (BLS) releases its estimate of the unemployment rate for the previous month. This statistic is widely regarded as a measure of the overall health of the U.S. economy. The media often interpret minor month-to-month blips in the unemployment rate as a sign of either a precipitous decline in economic activity or a surging recovery.

The unemployment rate is tabulated from the responses to a monthly BLS survey called the *Current Population Survey* (CPS). In this survey, nearly 60,000 households are questioned about their work activities during a particular week of the month (that week is called the reference week). Almost everything we know about trends in the U.S. labor force comes from tabulations of CPS data. The survey instrument used by the CPS also influenced the development of comparable surveys in other countries. In view of the importance of the CPS in the calculation of labor force statistics both in the United States and abroad, it is crucial to review the definitions of labor force activities that are routinely used by the BLS to generate its statistics.

The CPS classifies all persons aged 16 or older into one of three categories: The *employed*, the *unemployed*, and the residual group that is said to be *out of the labor force*. To be employed, a person must have been at a job with pay for at least 1 hour or worked at least 15 hours on a nonpaid job (such as the family farm). To be unemployed, a person must either be on a temporary layoff from a job or have no job but be actively looking for work in the four-week period prior to the reference week.

Let E be the number of persons employed and U the number of persons unemployed. A person participates in the **labor force** if he or she is either employed or unemployed. The size of the labor force (LF) is given by

$$LF = E + U \quad (2-1)$$

The vast majority of employed persons (those who work at a job with pay) are counted as being in the labor force regardless of how many hours they work. The size of the labor force, therefore, does not say anything about the “intensity” of work.

The **labor force participation rate** gives the fraction of the population (P) that is in the labor force and is defined by

$$\text{Labor force participation rate} = \frac{LF}{P} \quad (2-2)$$

The **employment rate** (also called the “employment–population ratio”) gives the fraction of the population that is employed, or

$$\text{Employment rate} = \frac{E}{P} \quad (2-3)$$

Finally, the **unemployment rate** gives the fraction of labor force participants who are unemployed:

$$\text{Unemployment rate} = \frac{U}{LF} \quad (2-4)$$

Note a crucial detail: The number of persons who are out of the labor force does not play *any* role in the calculation of the official unemployment rate.

The Hidden Unemployed

The BLS calculates an unemployment rate based on a subjective measure of what it means to be unemployed. To be considered unemployed, a person must either be on temporary lay-off or claim that he or she has “actively looked for work” in the past 4 weeks. Persons who have given up and stopped looking for work are not counted as unemployed, but are “out of the labor force.” At the same time, some persons who have little intention of working may claim to be actively looking for a job in order to qualify for unemployment benefits.

The unemployment numbers, therefore, can be interpreted in different ways. During the severe recession that began in 2009, for instance, it was argued that the official unemployment rate (that is, the BLS statistic) understated the economic hardships. Because it was so hard to find work, many laid-off workers became discouraged with their futile job search activity, dropped out of the labor market, and stopped being counted *as* unemployed. A more sensible approach would perhaps add this army of **hidden unemployed** to the pool of unemployed workers, making the unemployment rate far higher than it appeared from the BLS data. For example, if the “unemployed” included persons who are out of the labor force because they are “discouraged over job prospects” as well as persons who are only “marginally attached” to the labor force, the unemployment rate in March 2011 would have been 15.7 percent, rather than the official 8.8 percent.

Some analysts believe that a more objective measure of aggregate economic activity may be given by the employment rate. The employment rate gives the fraction of the population at a job. But this statistic has the drawback that it lumps together persons who say they are unemployed with everyone who is out of the labor force. Although the latter group includes the hidden unemployed, it also includes many individuals who have little intention of working, including retirees, some women with small children, and many students enrolled in school.

A decrease in the employment rate could then be attributed to either increases in unemployment or unrelated increases in fertility or school enrollment rates. It is far from clear, therefore, that the employment rate provides a better measure of fluctuations in economic activity than the unemployment rate. We will return to some of these issues in the unemployment chapter.

2-2 Basic Facts about Labor Supply

This section summarizes some of the key trends in labor supply in the United States.² These facts have motivated much of the research in recent decades. Table 2-1 documents the historical trends in the labor force participation rate of men. There was a slight fall in the labor force participation rates of men in the twentieth century, from 80 percent in 1900 to 71 percent by 2010. The decline is particularly steep for men near or above age 65, as more men choose to retire earlier. The labor force participation rate of men aged 45–64, for example, declined by 12 percentage points between 1950 and 2010, while the participation rate of men over 65 declined from 46 to 22 percent. Moreover, the labor force participation

² More detailed discussions of labor supply trends are given by John H. Pencavel, “Labor Supply of Men: A Survey,” in Orley C. Ashenfelter and Richard Layard, editors, *Handbook of Labor Economics*, vol. 1, Amsterdam: Elsevier, 1986, pp. 3–102; and Mark R. Killingsworth and James J. Heckman, “Female Labor Supply: A Survey,” in *ibid.*, pp. 103–204.

TABLE 2-1 Labor Force Participation Rates of Men, 1900–2010

Sources: U.S. Bureau of the Census, *Historical Statistics of the United States, Colonial Years to 1970*, Washington, DC: Government Printing Office, 1975; U.S. Bureau of the Census, *Statistical Abstract of the United States*, Washington, DC: Government Printing Office, various issues.

Year	All Men	Men Aged 25–44	Men Aged 45–64	Men Aged over 65
1900	80.0	94.7	90.3	63.1
1920	78.2	95.6	90.7	55.6
1930	76.2	95.8	91.0	54.0
1940	79.0	94.9	88.7	41.8
1950	86.8	97.1	92.0	45.8
1960	84.0	97.7	92.0	33.1
1970	80.6	96.8	89.3	26.8
1980	77.4	93.0	80.8	19.0
1990	76.4	93.3	79.8	16.3
2000	74.8	93.1	78.3	17.5
2010	71.2	90.6	78.4	22.1

TABLE 2-2 Labor Force Participation Rates of Women, 1900–2010

Sources: U.S. Bureau of the Census, *Historical Statistics of the United States, Colonial Years to 1970*, Washington, DC: Government Printing Office, 1975, p. 133; and U.S. Department of Commerce, *Statistical Abstract of the United States, 2011*, Washington, DC: Government Printing Office, 2011, Table 596.


Year	All Women	Single Women	Married Women	Widowed, Divorced, or Separated
1900	20.6	43.5	5.6	32.5
1910	25.4	51.1	10.7	34.1
1930	24.8	50.5	11.7	34.4
1940	25.8	45.5	15.6	30.2
1950	29.0	46.3	23.0	32.7
1960	34.5	42.9	31.7	36.1
1970	41.6	50.9	40.2	36.8
1980	51.5	64.4	49.9	43.6
1990	57.5	66.7	58.4	47.2
2000	59.9	68.9	61.1	49.0
2010	58.6	63.3	61.0	48.8

rate of men in their prime working years (ages 25–44) also declined, from 97 percent in 1950 to 91 percent in 2010. Note, however, that the labor force participation rate of men in their retirement years has stabilized and even begun to increase in the past two decades.

As Table 2-2 shows, there also has been a huge increase in the labor force participation rate of women. At the beginning of the twentieth century, only 21 percent of women were in the labor force. As late as 1950, even after the social and economic disruptions caused by two world wars and the Great Depression, only 29 percent of women were in the labor force. During the past 50 years, however, the labor force participation rate of women grew dramatically. By 2010, almost 60 percent of all women were in the labor force. It is worth noting that the increase in female labor force participation was particularly steep among married women. Their labor force participation rate almost doubled in recent decades, from 32 percent in 1960 to 61 percent in 2010.

FIGURE 2-1 Average Weekly Hours of Work, 1900–2013

Sources: The pre-1947 data refer to workers in manufacturing and are drawn from Ethel Jones, “New Estimates of Hours of Work per Week and Hourly Earnings, 1900–1957,” *Review of Economics and Statistics* 45 (November 1963): 374–385. The post-1947 data are drawn from U.S. Department of Labor, Bureau of Labor Statistics, *Employment, Hours, and Earnings from the Current Employment Statistics Survey*, “Table B-7. Average Weekly Hours of Production or Nonsupervisory Workers on Private Nonfarm Payrolls by Industry Sector and Selected Industry Detail.”



These dramatic shifts in labor force participation rates were accompanied by a sizable decline in average hours of work per week. Figure 2-1 shows that the typical person employed in production worked 55 hours per week in 1900, 40 hours in 1940, and just under 34 hours in 2010.

There exist sizable differences in the various dimensions of labor supply across demographic groups at a particular point in time. As Table 2-3 shows, men not only have larger participation rates than women, but are also less likely to be employed in part-time jobs. Only 4 percent of working men are in part-time jobs, as compared to 13 percent of working women. The table also documents a strong positive correlation between labor supply and educational attainment for both men and women. In 2017, 90 percent of male college graduates and 81 percent of female college graduates were in the labor force, as compared to only 72 and 46 percent of male and female high school dropouts, respectively. There are also racial differences in labor supply, between whites and minorities as well as within the minority population itself, with blacks tending to have the lowest participation rates and Asian men the highest.

2-3 The Worker's Preferences

The framework that economists typically use to analyze labor supply behavior is called the **neoclassical model of labor–leisure choice**. This model isolates the factors that determine whether a particular person works and, if so, how many hours she chooses to work. The model tells a simple “story” that helps us understand many of the stylized facts discussed above. More importantly, it lets us predict how changes in economic conditions or in government policies will affect work incentives.

TABLE 2-3 Labor Supply in the United States, 2017 (Persons Aged 25–64)

Source: U.S. Bureau of Labor Statistics, *Current Population Survey*, Annual Social and Economic Supplement, March 2017. The average number of hours worked is calculated in the subsample of workers. The percent of workers in part-time jobs refers to the proportion working fewer than 30 hours per week.

	Labor Force Participation Rate		Annual Hours of Work		Percent of Workers in Part-Time Jobs	
	Men	Women	Men	Women	Men	Women
All persons	83.1	71.4	2,170	1,933	4.3	12.9
Educational attainment:						
Less than 12 years	72.1	45.6	2,033	1,753	5.4	19.7
12 years	79.1	63.3	2,124	1,875	4.7	14.1
13–15 years	82.5	73.5	2,166	1,906	4.8	13.4
16 years or more	90.4	80.5	2,235	2,000	3.4	11.2
Age:						
25–34	87.1	75.6	2,101	1,904	5.7	12.0
35–44	89.2	75.1	2,201	1,928	2.7	13.0
45–54	85.3	74.7	2,221	1,978	2.9	12.0
55–64	70.5	60.0	2,160	1,922	6.2	15.2
Race:						
White	83.8	73.1	2,208	1,933	4.1	13.8
Black	74.9	72.0	2,096	1,963	6.0	9.6
Hispanic	85.6	64.7	2,086	1,882	4.0	12.7
Asian	87.5	68.2	2,121	1,961	3.1	11.3

The representative person in our model receives satisfaction both from the consumption of goods (which we denote by C) and from the consumption of leisure (L) in a particular time period. Obviously, the person consumes many different types of goods. To simplify, we aggregate the dollar value of all the goods that the person consumes and define C as the total dollar value of all goods purchased. For example, if the person spends \$1,000 weekly on food, rent, car payments, movie tickets, and other items, the variable C would take on the value of \$1,000. The variable L gives the number of hours of leisure that a person consumes during that same period.

Utility and Indifference Curves

The notion that individuals get satisfaction from consuming goods and leisure is summarized by the **utility function**:


$$U = f(C, L) \quad (2-5)$$

The utility function transforms the person's consumption of goods and leisure into an index U that measures the individual's level of satisfaction or happiness. This index is called *utility*. The higher the index U is, the happier the person will be. We make the sensible assumption that buying more goods or having more leisure hours both increase a person's utility. In the jargon of economics, C and L are "goods," not "bads."

Suppose that a person is consuming \$500 worth of consumption goods and 100 hours of leisure weekly (point Y in Figure 2-2). This particular consumption basket yields a particular level of utility to the person, say 25,000 utils. It is easy to imagine that different combinations

FIGURE 2-2 Indifference Curves

Points X and Y lie on the same indifference curve and yield the same utility (25,000 utils); point Z lies on a higher indifference curve and yields more utility.



of goods and leisure might yield the same level of utility. For example, the person might say that she would be indifferent to consuming \$500 worth of goods and 100 hours of leisure or consuming \$400 worth of goods and 125 hours of leisure. Figure 2-2 illustrates the many combinations of C and L that generate this particular level of utility. The locus of such points is called an **indifference curve**—and all points along this curve yield 25,000 utils.

Suppose that the person was instead consuming \$450 worth of goods and 150 hours of leisure (point Z in the figure). This consumption basket would certainly make the person happier, placing her on the higher indifference curve with 40,000 utils. We can then construct an indifference curve for that level of utility. In fact, we can construct an indifference curve for every level of utility. As a result, the utility function can be represented graphically in terms of a family (or a “map”) of indifference curves.

Indifference curves have four important properties:

1. *Indifference curves are downward sloping.* We assumed that individuals prefer more of both C and L . If indifference curves were upward sloping, a consumption basket with more C and more L would yield the same level of utility as a consumption basket with less C and less L . This clearly contradicts our assumption that the individual likes both goods and leisure. The only way that we can offer a person a few more hours of leisure, and still hold utility constant, is to take away some of the goods.
2. *Higher indifference curves indicate higher levels of utility.* The consumption bundles lying on the indifference curve that yields 40,000 utils are preferred to the bundles lying on the curve that yields 25,000 utils. To see this, note that point Z in the figure

must yield more utility than point X , simply because the bundle at point Z allows the person to consume more goods and more leisure.

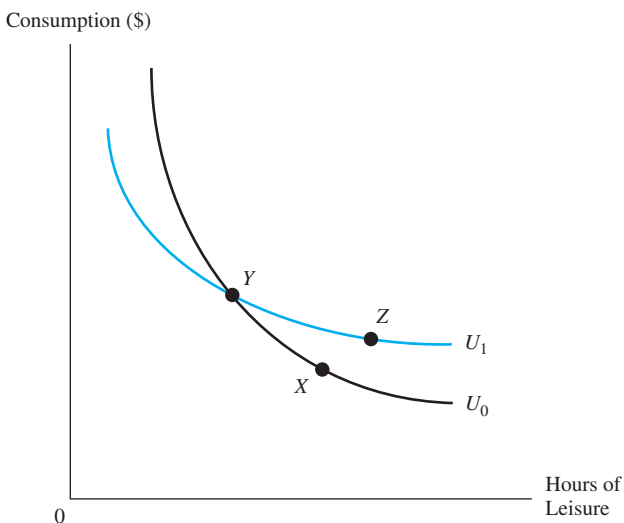
3. *Indifference curves do not intersect.* To see why, consider Figure 2-3, where indifference curves are allowed to intersect. Because points X and Y lie on the same indifference curve, the individual would be indifferent between the bundles X and Y . Because points Y and Z lie on the same indifference curve, the individual would be indifferent between bundles Y and Z . The person would then be indifferent between X and Y , and between Y and Z , so that she should also be indifferent between X and Z . But Z is clearly preferable to X , because Z has more goods and more leisure. Indifference curves that intersect contradict our assumption that individuals like to consume both goods and leisure.
4. *Indifference curves are convex to the origin.* The convexity of indifference curves does not follow from either the definition of indifference curves or the assumption that both goods and leisure are “goods.” The convexity reflects an additional assumption about the shape of the utility function. It turns out (see Problem 2-1 at the end of the chapter) that indifference curves must be convex to the origin if we are ever to observe a person both working and consuming some leisure in the same period.

The Slope of an Indifference Curve

What happens to a person’s utility as she allocates one more hour to leisure or buys an additional dollar’s worth of goods? The **marginal utility** of leisure is defined as the change in utility resulting from an additional hour devoted to leisure activities, holding constant the amount of goods consumed. We denote the marginal utility of leisure as MU_L . Similarly,

FIGURE 2-3 Indifference Curves Do Not Intersect

Points X and Y yield the same utility because they are on the same indifference curve; points Y and Z should also yield the same utility. Point Z , however, is preferable to point X .



the marginal utility of consumption gives the change in utility if the individual consumes one more dollar of goods, holding constant the number of hours of leisure. We denote the marginal utility of consumption by MU_C . Because both leisure and the consumption of goods are desirable activities, the marginal utilities of leisure and consumption must be positive numbers.

As we move along an indifference curve, say from point X to point Y in Figure 2-2, the slope of the indifference curve measures the rate at which a person is willing to give up some leisure time in return for additional consumption, *while holding utility constant*. Put differently, the slope tells us how many additional dollars' worth of goods it would take to "bribe" the person into giving up some leisure time. It can be shown that the slope of an indifference curve equals³

$$\frac{\Delta C}{\Delta L} = -\frac{MU_L}{MU_C} \quad (2-6)$$

The absolute value of the slope of an indifference curve, which is called the **marginal rate of substitution (MRS) in consumption**, is the ratio of marginal utilities.

The assumption that indifference curves are convex to the origin is essentially an assumption about how the marginal rate of substitution changes as the person moves along an indifference curve. Convexity implies that the slope of an indifference curve is steeper when the worker is consuming a lot of goods and little leisure, and flatter when the worker is consuming few goods and a lot of leisure. As a result, the absolute value of the slope of an indifference curve declines as the person "rolls down" the curve. The assumption of convexity, therefore, is equivalent to an assumption of *diminishing* marginal rate of substitution.

Differences in Preferences across Workers


The map of indifference curves presented in Figure 2-2 illustrates the way a *particular* worker views the trade-off between leisure and consumption. Different workers will view this trade-off differently. Some of us may like to devote a lot of time to our jobs, while others would prefer to devote most of their time to leisure. These differences in preferences imply that the indifference curves may look quite different for different workers.

Figure 2-4 shows the indifference curves for two workers, Cindy and Mindy. Cindy's indifference curves tend to be very steep, indicating that her marginal rate of substitution takes on a very high value (see Figure 2-4a). In other words, she requires a sizable monetary bribe (in terms of additional consumption) to convince her to give up an additional hour of leisure. Cindy obviously likes leisure a lot. Mindy, on the other hand, has flatter indifference curves, indicating that her marginal rate of substitution takes on a low value (see Figure 2-4b). Mindy, therefore, does not require a large bribe to convince her to give up an additional hour of leisure.

³ To show that the slope of an indifference curve equals the ratio of marginal utilities, suppose that points X and Y in Figure 2-2 are very close to each other. When going from X to Y , the person is giving up ΔL hours of leisure, and each hour of leisure given up has a marginal utility of MU_L . Therefore, the loss in utility associated with moving from X to Y is given by $\Delta L \times MU_L$. The move from X to Y also involves a gain in utility. After all, the worker is not just giving up leisure time; she is consuming an additional ΔC dollars of goods. Each additional dollar of consumption increases utility by MU_C units. The total gain in utility is given by $\Delta C \times MU_C$. All points along an indifference curve yield the same utility. This implies that the utility loss in moving from X to Y must be exactly offset by the gain, or $(\Delta L \times MU_L) + (\Delta C \times MU_C) = 0$. Equation (2-6) follows by rearranging terms.

FIGURE 2-4 Differences in Preferences across Workers

(a) Cindy's indifference curves are steep; she requires a substantial bribe to give up an hour of leisure. (b) Mindy's indifference curves are flatter; she attaches a much lower value to her leisure time.



(a) Cindy's Indifference Curves

(b) Mindy's Indifference Curves

Interpersonal differences in the “tastes for work” are obviously important determinants of differences in labor supply in the population. For the most part, economic models gloss over these differences in preferences. The reason for this omission is that differences in tastes, although probably very important, are hard to observe and measure. It would be extremely difficult, if not impossible, to conduct surveys that attempt to measure differences in indifference curves across workers. Moreover, the reliance on taste differences provides an easy way out for anyone who wishes to explain why different workers behave differently. One can always assert that the different behavior patterns of two workers arise because worker *A* likes leisure more than worker *B*, and there would be no way of proving whether such a claim is correct.

Economic models instead emphasize the impact of variables that are easily observable—such as wages and incomes—on the labor supply decision. Because these variables can be observed, the predictions made by the model about which types of persons will tend to work more are testable and refutable.

2-4 The Budget Constraint

The person's consumption of goods and leisure is constrained by her income and by the fact there are only 24 hours in a day. Part of a person's income (such as property income, dividends, and lottery prizes) is independent of how many hours she works. We denote this

“nonlabor income” by V . Let h be the number of hours the person will allocate to the labor market during the period and w be the hourly wage rate. The person’s **budget constraint** can be written as

$$C = wh + V \quad (2-7)$$

In words, the dollar value of expenditures on goods (C) must equal the sum of labor earnings (wh) and nonlabor income (V).⁴

The wage rate plays a central role in the labor supply decision. Initially, we assume that the wage rate is constant *for a particular person*, so the person receives the same hourly wage regardless of how many hours she works. In fact, the “marginal” wage rate (that is, the wage rate received for the last hour worked) generally depends on how many hours a person works. Persons who work over 40 hours per week typically receive an overtime premium, and the wage rate in part-time jobs is often lower than that in full-time jobs.⁵ For now, we ignore the possibility that a worker’s marginal wage may depend on how many hours she chooses to work.

It is then easy to graph the budget constraint. The person has two alternative uses for her time: Work or leisure. The total time allocated to each of these activities must equal the total time available in the period, say T hours per week, so that $T = h + L$. We can rewrite the budget constraint as

$$C = w(T - L) + V \quad (2-8)$$

or

$$C = (wT + V) - wL$$

This last equation is in the form of a line, and the slope is the negative of the wage rate (or $-w$).⁶ The **budget line** is illustrated in Figure 2-5. Point E in the graph indicates that if the person decides not to work at all and devotes T hours to leisure, she can still purchase V dollars’ worth of consumption goods. Point E is the *endowment point*. If the person is willing to give up 1 hour of leisure, she can then move up the budget line and purchase an additional w dollars’ worth of goods. In fact, each additional hour of leisure that the person is willing to give up allows her to buy an additional w dollars’ worth of goods. In other words, each hour of leisure consumed has a price, and the price is given by the wage rate. If the worker gives up all her leisure activities, she ends up at the intercept of the budget line and can buy $(wT + V)$ dollars’ worth of goods.

The consumption and leisure bundles that lie below the budget line are available to the worker; the bundles that lie above the budget line are not. The budget line, therefore, gives the frontier of the worker’s **opportunity set**—the set of all the consumption baskets that a particular worker could afford to buy.


⁴ The budget constraint implies that the worker spends all her income in the period, so there are no savings.

⁵ Shelly Lundberg, “Tied Wage-Hours Offers and the Endogeneity of Wages,” *Review of Economics and Statistics* 67 (August 1985): 405–410.

⁶ The equation of a line relating the variables y and x is $y = a + bx$, where a is the intercept and b is the slope.

FIGURE 2-5 The Budget Line Is the Boundary of the Worker's Opportunity Set

Point E is the endowment point, telling the person how much she can consume if she does not work at all. The worker moves up the budget line as she trades an hour of leisure for consumption of goods. The absolute value of the slope of the budget line is the wage rate.



2-5 The Hours of Work Decision

We make one important assumption about the person's behavior: She chooses the particular combination of goods and leisure that maximizes her utility. This means that the person will choose the level of goods and leisure that lead to the highest possible level of the utility index U —given the limitations imposed by the budget constraint.


Figure 2-6 illustrates the solution to this problem. As drawn, the budget line FE describes the opportunities available to a worker who has \$100 of nonlabor income per week, faces a market wage rate of \$10 per hour, and has 110 hours of nonsleeping time to allocate between work and leisure activities (assuming she sleeps roughly 8 hours per day).

Point P gives the optimal bundle of goods and hours of leisure chosen by the utility-maximizing worker. The highest indifference curve attainable places her at point P and gives her U^* units of utility. The worker then consumes 70 hours of leisure per week, works a 40-hour workweek, and buys \$500 worth of goods weekly. The worker would obviously prefer to choose a point on indifference curve U_1 , which provides a higher level of utility. For example, the worker would prefer to be at Y , where she works a 40-hour workweek and buys \$1,100 worth of goods. Given her wage and nonlabor income, however, the worker could never afford this outcome. In contrast, the worker could choose a point such as A , which lies on the budget line, but she would not do so. After all, point A gives her less utility than point P .

The optimal consumption of goods and leisure, therefore, is given by the point where the budget line is tangent to the indifference curve. This type of solution is called an *interior solution* because the worker is not at either corner of the opportunity set (that is, at point F , working all available hours, or at point E , working no hours whatsoever).

FIGURE 2-6 Interior Solution to the Labor–Leisure Decision

A utility-maximizing worker chooses the consumption–leisure bundle at point P , where the indifference curve is tangent to the budget line.

**Interpreting the Tangency Condition**

At the optimal point P , the slope of the indifference curve equals the slope of the budget line. This implies that⁷

$$\frac{MU_L}{MU_C} = w \quad (2-9)$$

At the chosen level of consumption and leisure, the marginal rate of substitution (the rate at which a person is willing to give up leisure hours in exchange for additional consumption) equals the wage rate (the rate at which the market allows the worker to substitute one hour of leisure time for consumption).

⁷ Although the slope of the indifference curve and the slope of the budget line are both negative numbers, the minus signs cancel out when the two numbers are set equal to each other, resulting in equation (2-9).

The economic intuition behind this condition is easier to grasp if we rewrite it as

$$\frac{MU_L}{w} = MU_C \quad (2-10)$$

The quantity MU_L gives the additional utility received from consuming an extra hour of leisure. This extra hour costs w dollars. The left-hand side of equation (2-10), therefore, gives the number of utils received from spending an additional dollar on leisure. Because C is defined as the dollar value of expenditures on consumption goods, MU_C gives the number of utils received from spending an additional dollar on goods. The tangency solution at point P implies that the last dollar spent on leisure buys the same number of utils as the last dollar spent on goods. If this equality did not hold (so that, for example, the last dollar spent on consumption buys more utils than the last spent on leisure), the worker would not be maximizing utility. She could rearrange her consumption plan so as to purchase more of the commodity that yields more utility for the last dollar.

What Happens to Hours of Work When Nonlabor Income Changes?

We want to know what happens to hours of work when the worker's nonlabor income V increases. The increase in V might be triggered by the payment of higher dividends on the worker's stock portfolio or because some distant relatives named the worker as the beneficiary in their will.

Figure 2-7 illustrates what happens to hours of work when the worker has an increase in V , *holding the wage constant*. Initially, the worker's nonlabor income equals \$100 weekly, which is associated with endowment point E_0 . Given the worker's wage rate, the budget line is then given by F_0E_0 . The worker maximizes utility by choosing the bundle at point P_0 . At this point, the worker consumes 70 hours of leisure and works 40 h.

The increase in nonlabor income to \$200 weekly shifts the endowment point to E_1 , so that the new budget line is given by F_1E_1 . Because the worker's wage rate is being held constant, the slope of the new budget line is the same as the slope of the budget line that originated at point E_0 . An increase in nonlabor income that holds the wage constant expands the worker's opportunity set through a parallel shift in the budget line.


The increase in nonlabor income allows the worker to jump to a higher indifference curve, such as point P_1 in Figure 2-7. Increases in nonlabor income necessarily make the worker better off. After all, the expansion of the opportunity set opens up many additional opportunities for the worker. Figure 2-7a draws point P_1 so that the additional nonlabor income increases both purchases on goods and leisure hours. As a result, the length of the workweek falls to 30 hours. Figure 2-7b draws point P_1 so that the additional nonlabor income reduces leisure hours, increasing the length of the workweek to 50 hours. The impact of the change in nonlabor income (holding wages constant) on the number of hours worked is called an **income effect**.

Both panels in Figure 2-7 draw "legal" indifference curves. The indifference curves are downward sloping, do not intersect, and are convex to the origin. We cannot predict how an increase in nonlabor income affects hours of work unless we make an additional restriction on the shape of indifference curves. The additional restriction we make is that leisure is a "normal" good (as opposed to leisure being an "inferior" good).

We define a commodity to be a normal good when increases in income, holding the prices of all goods constant, increase its consumption. A commodity is an inferior good

FIGURE 2-7 The Effect of a Change in Nonlabor Income on Hours of Work

An increase in nonlabor income leads to a parallel, upward shift in the budget line, moving the worker from point P_0 to point P_1 . (a) If leisure is a normal good, hours of work fall. (b) If leisure is an inferior good, hours of work rise.



when increases in income, holding prices constant, decrease its consumption. Low-priced subcompact cars, for instance, are typically thought of as inferior goods, whereas BMWs are typically thought of as normal goods. In other words, we would expect the demand for low-quality subcompacts to fall as nonlabor income increases, and the demand for BMWs to increase.

If we reflect on whether leisure is a normal or an inferior good, most of us would probably conclude that leisure is a normal good. Put differently, if we were wealthier, we would surely demand a lot more time off. We could then visit Aspen in December, Rio in February, and exotic beaches in the summer.

Because it seems reasonable to assume that leisure is a normal good and because there is some evidence (discussed below) supporting this assumption, our discussion focuses on this case. The assumption that leisure is a normal good resolves the conflict between the two panels in Figure 2-7 in favor of the one on the left-hand side. *The income effect, therefore, implies that an increase in nonlabor income, holding the wage rate constant, reduces hours of work.*

What Happens to Hours of Work When the Wage Changes?

Consider a wage increase from \$10 to \$20 an hour, holding nonlabor income V constant. The wage increase rotates the budget line around the endowment point, as illustrated in Figure 2-8. The rotation of the budget line shifts the opportunity set from FE to GE . It should be obvious that a wage increase does not change the endowment point: The dollar value of the goods that can be consumed when one does not work is the same regardless of whether the wage rate is \$10 or \$20 an hour.

The two panels presented in Figure 2-8 illustrate the possible effects of a wage increase on hours of work. In Figure 2-8a, the wage increase shifts the optimal consumption bundle from point P to point R . At the new equilibrium, the individual consumes more leisure (from 70 to 75 hours), so that hours of work fall from 40 to 35 hours.


Figure 2-8b, however, shows the opposite result. The wage increase again moves the worker to a higher indifference curve and shifts the optimal consumption bundle from point P to point R . This time, however, the wage increase reduces leisure hours (from 70 to 65 hours), so the length of the workweek increases from 40 to 45 hours. It seems, therefore, that we cannot make an unambiguous prediction about an important question without making even more assumptions.

The reason for the ambiguity in the relation between hours of work and the wage rate is of fundamental importance and introduces tools and ideas that play a central role in all of economics. Both panels in Figure 2-8 show that, regardless of what happens to hours of work, a wage increase expands the worker's opportunity set. Put differently, a worker has more opportunities when she makes \$20 an hour than when she makes \$10 an hour. We know that an increase in income increases the demand for all normal goods, including leisure. The increase in the wage thus increases the demand for leisure, which reduces hours of work.

But this is not all that happens. The wage increase also makes leisure more expensive. When the worker earns \$20 an hour, she gives up \$20 every time she decides to take an hour off. Leisure time is a very expensive commodity for high-wage workers and is relatively cheap for low-wage workers. High-wage workers would have strong incentives to cut

FIGURE 2-8 The Effect of a Change in the Wage Rate on Hours of Work

A change in the wage rate rotates the budget line around the endowment point E . A wage increase moves the worker from point P to point R , and can either decrease or increase hours of work.



back on their consumption of leisure. A wage increase thus reduces the demand for leisure and increases hours of work.

This discussion highlights the source of the ambiguity in the relation between hours of work and the wage rate. A high-wage worker wants to enjoy the rewards of her high income, and would like to consume more leisure. The same worker, however, finds that leisure is very expensive and that she simply cannot afford to take time off from work.

These two conflicting forces are illustrated in Figure 2-9a. The initial wage rate is \$10 per hour. The worker maximizes her utility by choosing the consumption bundle given by point *P*, where she consumes 70 hours of leisure and works 40 hours per week. Suppose the wage increases to \$20. The budget line rotates and the new consumption bundle is given by point *R*. The worker is now consuming 75 hours of leisure and working 35 h. As drawn, the person is working fewer hours at the higher wage.


It helps to think of the move from point *P* to point *R* as a two-stage move. The two stages correspond exactly to our discussion that the wage increase generates two effects: It increases the worker's income and it raises the price of leisure. In particular, suppose we draw a budget line that is parallel to the old budget line (so that its slope is also $-\$10$), but tangent to the new indifference curve. This budget line (*DD*), also illustrated in Figure 2-9a, generates a new tangency point *Q*.

The move from initial position *P* to final position *R* can then be decomposed into a first-stage move from *P* to *Q* and a second-stage move from *Q* to *R*. It is easy to see that the move from point *P* to point *Q* is an income effect. In particular, the move from *P* to *Q* arises from a change in the worker's income, holding wages constant. The income effect

FIGURE 2-9 Income and Substitution Effects


An increase in the wage rate generates both income and substitution effects. The income effect (the move from point *P* to point *Q*) reduces hours of work; the substitution effect (the move from *Q* to *R*) increases hours of work.

Consumption (\$)



(a) Income Effect Dominates

Consumption (\$)



(b) Substitution Effect Dominates

gives the change in the consumption bundle induced by the additional income resulting from the wage increase. Because both leisure and goods are normal goods, point Q must lie to the northeast of point P (so that more is consumed of both goods and leisure). The income effect increases the demand for leisure (from 70 to 85 hours) and reduces hours of work by 15 hours per week.

The second-stage move from Q to R is called the **substitution effect**. It illustrates what happens to the worker's consumption bundle as the wage increases, holding utility constant. By moving along an indifference curve, the worker's utility or "real income" is held fixed. The substitution effect isolates the impact of the increase in the price of leisure on hours of work, holding real-income constant.

The move from point Q to point R shows a substitution away from leisure and toward goods. In other words, as the wage rises, the worker devotes less time to expensive leisure activities (from 85 to 75 hours) and increases her consumption of goods. Through the substitution effect, therefore, hours of work rise by 10 hours. *The substitution effect implies that an increase in the wage rate, holding real income constant, increases hours of work.*

As drawn in Figure 2-9a, the decrease in hours of work generated by the income effect (15 hours) exceeds the increase in hours of work associated with the substitution effect (10 hours). The stronger income effect thus leads to a negative relationship between hours of work and the wage rate. In Figure 2-9b, the income effect (again the move from point P to point Q) decreases hours of work by 10 hours, whereas the substitution effect (the move from Q to R) increases hours of work by 15 hours. Because the substitution effect dominates, there is a positive relationship between hours of work and the wage rate.

The reason for the ambiguity in the relationship between hours of work and the wage rate should now be clear. As the wage rises, a worker faces a larger opportunity set and that income effect increases her demand for leisure and reduces hours of work. As the wage rises, however, leisure becomes more expensive and the substitution effect encourages the worker to switch away from the consumption of leisure and instead consume more goods. This shift frees up leisure hours and increases hours of work.

To summarize:

- An increase in the wage rate increases hours of work if the substitution effect dominates the income effect.
- An increase in the wage rate decreases hours of work if the income effect dominates the substitution effect.

2-6 To Work or Not to Work?

Our analysis of the relation between nonlabor income, the wage rate, and hours of work assumed that the person worked both before and after the change in nonlabor income or the wage. Hours of work then adjusted to the change in the opportunity set. But what factors motivate a person to work in the first place?

To illustrate the nature of this decision, consider Figure 2-10. The figure draws the indifference curve that goes through the endowment point E . This indifference curve indicates that a person who does not work at all receives U_0 units of utility. The woman, however, can choose to enter the labor market and trade some of her leisure time for earnings that will allow her to buy goods. The decision of whether to work or not boils down to

Theory at Work

DOLLARS AND DREAMS

The fact that our consumption of leisure responds to its price is not surprising. When the wage rate is high, we will find ways of minimizing the use of our valuable time. We will go through a ticket broker and pay high prices for concert and theater tickets, rather than stand in line for hours to buy a ticket at face value. We will hire a nanny or send our children to day care, rather than withdraw from the labor market. And we will consume preprepared meals and order pizza or take-out Chinese, rather than engage in lengthy meal preparations.

It turns out that how we allocate our time responds to economic incentives even when there are no easy substitutes available, such as when we decide how many hours to sleep. Sleeping takes a bigger chunk of our time than any other activity. The typical person sleeps around 57 hours a week. Although most of us believe that how long we sleep is biologically (and perhaps even culturally) determined, there is evidence that hours sleeping

can also be viewed as another activity that responds to economic incentives. As long as some minimum biological threshold for the length of a sleeping spell is met, the demand for sleep time seems to respond to changes in the price of time.

In particular, there is a negative correlation between a person's earnings capacity and the number of hours spent sleeping. More highly educated persons, for example, sleep less—an additional four years of school reduces sleep time by about an hour per week. Similarly, a 20 percent wage increase reduces sleep time by 1 percent, or about 34 minutes per week. Even dreaming of a nice vacation in a remote island becomes expensive when our time is valuable.

Source: Jeff E. Biddle and Daniel S. Hamermesh, "Sleep and the Allocation of Time," *Journal of Political Economy* 98 (October 1990): 922–943.

a simple question: Are the "terms of trade"—the rate at which leisure can be traded for goods—sufficiently attractive to bribe her into entering the labor market?

Suppose initially that the person's wage rate is given by w_{low} so that the woman faces budget line GE in Figure 2-10. No point on this budget line can give her more utility than U_0 . At this low wage, the person's opportunities are quite meager. If the worker were to move from the endowment point E to any point on the budget line GE , she would be moving to a lower indifference curve. For example, at point X the woman gets only U_G utils. At wage w_{low} , therefore, the woman chooses not to work.


In contrast, suppose that the wage rate was given by w_{high} , so that the woman faces budget line HE . Moving to any point on this steeper budget line would increase her utility. At point Y , the woman gets U_H utils. At the wage w_{high} , therefore, the woman is better off working.

In sum, Figure 2-10 indicates that the woman does not work at low-wage rates (such as w_{low}), but does work at high-wage rates (such as w_{high}). As we rotate the budget line from wage w_{low} to wage w_{high} , we will typically encounter a wage rate, call it w^* , that makes her indifferent between working and not working. We call w^* the **reservation wage**. The reservation wage gives the minimum increase in income that would make a person indifferent between remaining at the endowment point E and working that first hour. It is given by the absolute value of the slope of the indifference curve at point E .

The definition of the reservation wage implies that the person will not work when the market wage is less than the reservation wage; but the person will work when the market

FIGURE 2-10 The Reservation Wage

If the person chooses not to work, she can remain at the endowment point E and get U_0 units of utility. At a low wage (w_{low}), the person is better off not working. At a high wage (w_{high}), she is better off working. The reservation wage w^* is given by the slope of the indifference curve at the endowment point.



wage exceeds the reservation wage. The decision to work, therefore, depends entirely on a comparison of the market wage, which indicates how much employers are willing to pay for an hour of work, and the reservation wage, which indicates how much the worker requires to be bribed into working that first hour.

The theory obviously implies that a high reservation wage makes it less likely that a person will work. The reservation wage will typically depend on the person's tastes for work, which helps to determine the slope of the indifference curve, as well on many other factors. For instance, the assumption that leisure is a normal good implies that the reservation wage rises as nonlabor income increases.⁸ Because workers want to consume more leisure as nonlabor income increases, a larger bribe will be required to convince a wealthier person to enter the labor market.

⁸ Try to prove this statement by drawing a vertical line through the endowment point in Figure 2-6. Because of convexity, the indifference curves will get steeper as we move to higher indifference curves.