



MARKETS AND THE ENVIRONMENT

SECOND EDITION

NATHANIEL O. KEOHANE  SHEILA M. OLMSTEAD

MARKETS
AND THE
ENVIRONMENT

Second Edition

Foundations of Contemporary Environmental Studies

Ecology and Ecosystem Conservation

Oswald J. Schmitz

Global Environmental Governance

James Gustave Speth and Peter M. Haas

Coastal Governance

Richard Burroughs

Water Resources

Shimon C. Anisfeld

Ecology and Religion

John Grim and Mary Evelyn Tucker

MARKETS AND THE ENVIRONMENT

Second Edition

Nathaniel O. Keohane

Sheila M. Olmstead

 **ISLANDPRESS**


Washington | Covelo | London

Copyright © 2016 Nathaniel O. Keohane and Sheila M. Olmstead
First Island Press edition, 2007

All rights reserved under International and Pan-American Copyright Conventions.
No part of this book may be reproduced in any form or by any means without
permission in writing from the publisher: Island Press, 2000 M Street, NW,
Suite 650, Washington, DC 20036.

ISLAND PRESS is a trademark of the Center for Resource Economics.

Library of Congress Control Number: 2015939771

Printed on recycled, acid-free paper 

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Keywords: environmental economics, environmental policy, cost-benefit analysis,
externalities, capital assets, natural resources, green accounting, water pricing,
carbon market, limits to growth

For Frances and Eleanor, and for Gau.
—N.O.K.

For Kevin, Finn, and Laurel.
—S.M.O.

Contents

Preface xv

Chapter 1: Introduction 1

Economics and the Environment 1

Global Climate Change 2

Organization and Content of This Book 6

What We Hope Readers Will Take Away from This Book 9

Chapter 2: Economic Efficiency and Environmental Protection 11

Economic Efficiency 12

Efficiency and Environmental Policy 20

Equating Benefits and Costs on the Margin 22

Dynamic Efficiency and Environmental Policy 31

Conclusion 34

Chapter 3: The Benefits and Costs of Environmental Protection 35

Measuring Costs 35

Evaluating the Benefits 44

Benefit–Cost Analysis 55

Conclusion 68

Chapter 4: The Efficiency of Markets 69

Competitive Market Equilibrium 70

The Efficiency of Competitive Markets 74

Conclusion 78

Chapter 5: Market Failures in the Environmental Realm 80

Externalities 81

Public Goods 85

The Tragedy of the Commons 91

Conclusion 97

Chapter 6: Managing Stocks: Natural Resources as Capital Assets 99

- Economic Scarcity 100
- Efficient Extraction in Two Periods 103
- A Closer Look at the Efficient Extraction Path 105
- What about Market Power? 110
- The Critical Role of Property Rights 111
- Conclusion 113

Chapter 7: Stocks That Grow: The Economics of Renewable Resource Management 114

- Economics of Forest Resources 114
- Fisheries 128
- Conclusion 137

Chapter 8: Principles of Market-Based Environmental Policy 139

- The Coase Theorem 140
- The Array of Policy Instruments 143
- How Market-Based Policies Can Overcome Market Failure 147
- Raising Revenues 160
- Setting Prices versus Setting Quantities 162
- Conclusion 167

Chapter 9: The Case for Market-Based Instruments in the Real World 168

- Reducing Costs 169
- Promoting Technological Change 179
- Market-Based Instruments for Managing Natural Resources 184
- Other Considerations 189
- Conclusion 198

Chapter 10: Market-Based Instruments in Practice 199

- The U.S. Sulfur Dioxide Market 200
- Individual Tradable Quotas for Fishing in New Zealand 207
- Municipal Water Pricing 214
- The European Union's Emissions Trading System 217
- Water Quality Trading 221
- Waste Management: "Pay as You Throw" 223
- Habitat and Land Management 224
- Conclusion 230

Chapter 11: Sustainability and Economic Growth 231

- Limits to Growth? 232
- Sustainability, in Economic Terms 238
- Keeping Track: Green Accounting 245

Are Economic Growth and Sustainability Compatible? 252

Conclusion 253

Chapter 12: Conclusion 254

What Does Economics Imply for Environmental Policy? 254

The Roles of Firms, Consumers, and Governments 256

Some Final Thoughts 257

Discussion Questions 259

References 267

Further Reading 291

Index 297

Preface

This book provides a concise introduction to the economic theory of environmental policy and natural resource management. If you have used this book before, you may be asking yourself what is new in the second edition. In the 8 years since the publication of the first edition, although little has changed in economic theory with respect to environmental quality and environmental policy, readers urged us to revise the book for several reasons. First, faculty members using the book to teach undergraduate environmental and resource economics encouraged us to strengthen the links between the material in the book and that covered in a typical introductory microeconomics course, mostly via changes in the language we used in discussing economic concepts. We've done this throughout the book. Also at the recommendation of users, the descriptions of cost and benefit estimation in Chapter 2 have been revised and expanded, and the discussion of environmental taxation in Chapter 8 has been restructured.

Research in the field of environmental economics moves quickly, and we've incorporated a good deal of important new knowledge created since the first edition. Throughout the book, we have updated old examples and added many new examples of market-based environmental policy in action, primarily in the boxes that accompany the text but also in the text itself. In this vein, major updates were made to the coverage of deforestation in Chapter 7, the discussion of market-based instruments and nonuniformly mixed pollutants in Chapter 9, and all sections in Chapter 10.

Finally, we were shocked at how quickly some of the popular culture references in the first edition (those to compact discs and Napster, for example) became dated, so we've done our best to sound current, although we admit that our children are now better sources for this kind

of information than we are. Despite these many changes, this edition preserves the basic structure of the original, with some small exceptions; for example, we have dropped the mathematical appendix on the economics of fishing from Chapter 7.

As in the first edition, our goal is to illuminate the role economic theory—and more broadly economic *thinking*—can play in informing and improving environmental policy. To our minds, noneconomists tend to perceive economics rather narrowly, as being concerned only with money or with national indicators such as exchange rates and trade balances. In fact, economics has a much wider reach. It sheds light on individuals' consumption choices in the face of scarce resources, the interaction between firms and consumers in a market, the extent to which individuals are likely to contribute toward the common good or ignore it in the pursuit of their own self-interest, and the ways government policies and other institutions shape incentives for action (or lack thereof). As we explain in the first chapter of the book, economics is central to understanding why environmental problems arise and how and why to address them. As concerned citizens as well as economists, we think it is vital for anyone interested in environmental policy to be conversant in the language of economics.

The approach we have taken here draws on our own experience teaching environmental and natural resource economics to master's students and undergraduates. It also draws on our experiences in the real world of environmental policy, in the public and nonprofit sectors. The emphasis is on intuition rather than algebra; we seek to convey the underlying concepts through words and graphs, presenting mathematical results only when necessary. We have also included a wealth of real-world examples, from the conservation of the California condor, to mitigation of global climate change, to using markets to manage fisheries in New Zealand and elsewhere.

The book was written with university students in mind, but its informal style and the importance of the subject make it suitable for a wide range of professionals or other concerned readers seeking an introduction to environmental economics. We have tried to make the language accessible to someone without any prior knowledge of economics. At the same time, the treatment is comprehensive enough that even an economics major with little experience in environmental policy could learn a great deal from the book. The lack of mathematical notation does not reduce the rigor of the underlying analysis.

In our teaching, we have noticed a gap between short articles on how economists think about the environment and textbooks filled with algebra and detailed information on the history of U.S. federal environmental legislation. In addition, most textbooks on the subject of markets and the environment treat either the economics of pollution control or the economics of natural resource management. At an introductory level there is little integration of these two “halves” of the discipline of environmental and resource economics. This book aims to fill these gaps. It can be used as a primer for a core course in environmental studies, at either the undergraduate or master’s level. In that context, this book would be the sole economics text, used alongside several other books representing different perspectives on environmental studies from the social, natural, and physical sciences. The book is also well suited to a semester-long course in environmental or natural resource economics, either as a main text (supplemented with more mathematical lecture notes and problem sets) or as a complement to another, more detailed (but perhaps less intuitive) textbook. Finally, the book could be used (as we ourselves have used the notes from which it grew) as an introduction to environmental economics in a course with a different focus. For example, a course on business strategy can use this book to explain the basic logic and practice of market-based policies to regulate pollution. Similarly, a principles of microeconomics course could use this book to show how economic theory can be applied to real-world problems and illuminate the market failures aspect of the course.

At the end of the volume, readers will find a list of references, including works cited in the text and other recommended readings of possible interest. We have also provided a set of study questions for each chapter, designed to be thought provoking and open-ended rather than simply reiterating the material.

We thank Karen Fisher-Vanden for providing thoughtful comments on the first edition and Robert Stavins, Elizabeth Walker, and Louise Marshall for their extensive input on what to fix in the second. We are also grateful to the book’s many other users who have e-mailed us comments, suggestions, and corrections over the years. Please keep that information coming. Our editors at Island Press for both editions, Todd Baldwin and Emily Davis, patiently moved us through the process of writing and revising the book. We thank our spouses, Todd Olmstead and Georgia Levenson Keohane, for their support and encouragement. Finally, we both owe a great deal to Robert Stavins, whose passion for teaching environmental

economics and communicating its principles to policymakers—and unrivaled ability to do so—continues to inspire us.

Nathaniel O. Keohane
New York, New York

Sheila M. Olmstead
Austin, Texas

1

Introduction

This book is a primer on the economics of the environment and natural resources. The title, *Markets and the Environment*, suggests one of our central themes. An understanding of markets—why they work, when they fail, and what lessons they offer for the design of environmental policies and the management of natural resources—is central to an understanding of environmental issues. But even before we start thinking about how markets work, it is useful to begin with a more basic question: What *is* environmental economics?

Economics and the Environment

“Environmental economics” may seem like a contradiction in terms. Some people think that economics is just about money, that it is preoccupied with profits and economic growth and has nothing to do with the effects of human activity on the planet. Others view environmentalists as being naive about economic realities or “more concerned about animals than jobs.”

Of course neither stereotype is true. Indeed, not only is “the environment” not separate from “the economy,” but environmental problems cannot be fully understood without understanding basic economic concepts. Economics helps explain why firms and individuals make the decisions they do—why coal (despite generating significant local air pollution and carbon dioxide emissions) still generates almost 40 percent of electricity in the United States, or why some people drive large sport utility vehicles instead of Priuses. Economics also helps predict how those same firms and individuals will respond to a new set of incentives—for example, what investments electric utilities will make in a carbon-constrained world and

how high gas prices would have to rise before people stopped buying enormous cars.

At its core, economics is the study of the allocation of scarce resources. This central focus, as much as anything else, makes it eminently suited to analyzing environmental problems. Let's take a concrete example. The Columbia and the Snake Rivers drain much of the U.S. Pacific Northwest, providing water for drinking, irrigation, transportation, and electricity generation and supporting endangered salmon populations. All these activities—including salmon preservation—provide *economic* benefits to the extent that people value them.

If there is not enough water to meet all those needs, then we must trade off one good thing for another (less irrigation for more fish habitat, for example). How should we as a society balance these competing claims against each other? To what lengths should we go to protect the salmon? What other valued uses should we give up? We might reduce withdrawals of water for agricultural irrigation, remove one or more hydroelectric dams, or implement water conservation programs in urban areas. How do we assess these various options?

Economics provides a framework for answering these questions. The basic approach is simple enough: Measure the costs and benefits of each possible policy, including a policy of doing nothing at all, and then choose the policy that generates the maximum net benefit to society as a whole (that is, benefits minus costs). This is easier to say than to do, but economics also provides tools for measuring costs and benefits. Finally, economic theory suggests how to design policies that harness market forces to work *for* rather than against environmental protection.

To illustrate how economic reasoning can help us understand and address environmental problems, let's take a look at perhaps the most pressing environmental issue today: global climate change.

Global Climate Change

There is overwhelming scientific consensus that human activity—primarily the burning of fossil fuels and deforestation caused by agriculture and urbanization—is responsible for a sharp and continuing rise in the concentration of carbon dioxide (CO₂) and other heat-trapping gases in the earth's atmosphere. The most direct consequence is a rise in average global surface temperatures, which is why the phenomenon is known widely as global warming. (Globally averaged surface temperatures have already increased by 0.85°C, or about 1.5 degrees Fahrenheit, since the late nineteenth century.)¹ But the consequences are much broader than

warming, which is why the broader term *climate change* is more apt. Expected impacts (many of which are already measurable) include sea level rise from the melting of polar ice caps; regional changes in precipitation; the disappearance of glaciers from high mountain ranges; the deterioration of coastal reefs; increased frequency of extreme weather events such as droughts, floods, and major storms; species migration and extinction; and spatial shifts in the prevalence of disease. The worst-case scenarios include a reversal of the North Atlantic thermohaline circulation, better known as the Gulf Stream, which brings warm water northward from the tropics and makes England and the rest of northern Europe habitable. Although there has been much international discussion about the potential costs and benefits of taking steps to slow or reverse this process, little progress has been achieved.

What are the causes of climate change? A natural scientist might point to the complex dynamics of the earth's atmosphere—how CO₂ accumulating in the atmosphere traps heat (the famous greenhouse effect) or how CO₂ gets absorbed by ocean and forest sinks. From an *economic* point of view, the roots lie in the incentives facing individuals, firms, and governments. Each time we drive a car, turn on a light, or use a computer, we are indirectly increasing carbon emissions and thereby contributing to global climate change. In doing so, we impose a small cost on the earth's population. However, these costs are invisible to the people responsible. You do not pay for the carbon you emit. Nor, indeed, does the company that provides your electricity (at least if you live in most of the United States) or the company that made your car. The result is that we all put CO₂ into the atmosphere, because we have no reason not to. It costs us nothing, and we receive significant individual benefits from the energy services that generate carbon emissions.

Economics stresses the importance of incentives in shaping people's behavior. Without incentives to pay for the true costs of their actions, few people (or firms) will voluntarily do so. You might think at first that this is because the "free market" has prevailed. In fact, that gets it almost exactly backward. Very often, as we shall see in this book, the problem is not that markets are so pervasive but that they are not pervasive *enough*—that is, they are incomplete. There is simply no market for clean air or a stable global climate. If there were, then firms and individuals who contributed to climate stabilization (by reducing their own carbon emissions or offsetting them) would be rewarded for doing so, just as firms that produce automobiles earn revenue from selling cars. This is a key insight from economics: Many environmental problems would be alleviated if proper

markets existed. Because those markets usually don't arise by themselves (for reasons we shall discuss later on in the book), governments have a crucial role in setting them up—or in creating price signals that mimic the incentives a market would provide.

If this is such a problem, you may have asked yourself, why haven't the world's countries come together and designed a policy to solve it? After all, the consequences of significant climate change may be dire, especially for low-lying coastal areas and countries in which predicted changes in temperature and precipitation will marginalize much existing agricultural land. If you have been following the development of this issue in the global media, and you know of the difficulty experienced by the international community in coming to agreement over the appropriate measures to take in combating climate change, it will not be terribly surprising that economics predicts that this is a difficult problem to solve. Carbon emission abatement is what economists would call a global *public good*: Everyone benefits from its provision, whether they have contributed or not. If a coalition of countries bands together to achieve a carbon emission abatement goal, all countries (including nonmembers of the coalition) will benefit from their efforts. So how can countries be induced to pay for it if they will receive the benefit either way? This is a thorny problem to which we will return in later chapters.

As a starting point, we must understand just what the benefits of carbon emission abatement are. They may be obvious to you. Put simply, slowing climate change can help us avert damages. For example, rising seas may inundate many coastal areas. If it is possible to slow or reverse this process, we might avoid damages including the depletion of coastal wetlands, the destruction of cultural artifacts, and the displacement of human populations. Warming in Arctic regions may lead to the extinction of the polar bear and other species; the benefits from slowing or reversing climate change would include the prevention of this loss. Climate change may exacerbate local pollution (such as ground-level ozone) and boost the spread of disease (such as malaria in the tropics and West Nile virus in North America); we would want to measure the benefits from avoiding those damages as well. Policies to mitigate climate change may also bring “co-benefits,” as when a shift away from burning fossil fuels results in lower levels of local and regional air pollution from sulfur dioxide or particulate matter.

All these benefits (even the intangible ones such as species preservation) have economic value. In economic terms, their value corresponds to what people would be willing to pay to secure them. Measuring this

value is easy when the losses are reflected in market prices, such as damages to commercial property or changes in agricultural production. But economists also have developed ways to measure the benefits of natural resources and environmental amenities that are not traded in markets, such as the improvements in human health and quality of life from cleaner air, the ecosystem services provided by wetlands, or the existence value of wilderness.

The economic cost of combating global climate change, meanwhile, is the sum of what must be sacrificed to achieve these benefits. Economic costs include not just out-of-pocket costs but also (and more importantly) the forgone benefits from using resources to slow or reverse climate change rather than for other objectives. Costs are incurred by burning cleaner but more expensive fuels or investing in pollution abatement equipment; by changing individual behavior, say by turning down the heat or air conditioning; by sequestering carbon in forests, oceans, depleted oil reservoirs, and other sinks; and by adapting to changing climatic conditions, for example by switching crops or constructing seawalls. Costs arise from directing government funds for research and development into climate-related projects rather than other pursuits. And of course the implementation, administration, monitoring, and enforcement of climate policy incurs some costs, as with almost any public policy.

Sound public policy decisions require an awareness of these costs and benefits and some ability to compare them in a coherent and consistent fashion. Economics provides a framework for doing so. In practice, as you will see through the theory and examples in this book, implementing the framework requires taking account of a number of other wrinkles. For example, we must worry about how to weigh near-term costs against benefits that accrue much later.

Rigorous consideration of economic benefits and costs can help answer the questions, “How much should we reduce greenhouse gas emissions in order to limit future climate change? How stringent should policies to address climate change be?” Economics can also shed light on a distinct but equally important question: “*How* should those policies be designed?”

For example, under the Copenhagen Accord, signed in 2009, the United States committed to reduce greenhouse gas emissions by 17 percent below 2005 levels. Although a large number of economic analyses informed the debate about this target, it was ultimately the result of political decisions rather than any explicit calculation of economic efficiency. Even so, economics can help inform the design of policies to meet the target. Emissions can be reduced in myriad ways: by requiring polluters

to install and operate specific abatement technologies or to meet specific standards of performance at their facilities, by mandating tough energy efficiency standards for consumer appliances and tightening fuel economy requirements for vehicles, by levying a tax on greenhouse gas emissions, or by capping emissions and allowing emitters to trade allowances under that cap. (And that is hardly an exhaustive list!) As we will discuss at length in this book, especially in Chapters 8 through 10, both economic theory and experience provide compelling arguments for market-based policies, such as emission taxes and cap-and-trade policies, that harness market forces to achieve regulatory goals at less overall cost than traditional approaches.

In sum, economics offers quite a different approach than other disciplines to the problem of global climate change—and to a range of other environmental issues we will explore in this book. You will find that the economic approach sometimes arrives at answers that are compatible with other approaches and sometimes at answers that conflict with those approaches. Regardless of such agreement or disagreement, economics provides a set of tools and a way of thinking that anyone with a serious interest in understanding and addressing environmental problems should be familiar with.

Organization and Content of This Book

This book provides an introduction to the application of economic reasoning to environmental issues and policies. In each chapter, we draw heavily on a range of real-world examples to illustrate our points.

Chapter 2 begins by asking, “Why compare benefits and costs?” Here we introduce the central concept of economic efficiency, meaning the maximization of the net benefits of a policy to society. We illustrate the key points by discussing the abatement of sulfur dioxide at U.S. power plants, and many other examples. We introduce the key concepts of marginal costs and benefits, showing how they relate to total costs and benefits and how they inform the analysis of efficiency. We also extend the concept of efficiency to the dynamic context, in which policies are defined by streams of benefits and costs occurring over time. In doing so, we introduce the concept of discounting, the process by which economists convert values in the future to values today, and explain its usefulness in a dynamic setting.

Chapter 3 follows up on the same themes. We discuss at length how economists define and measure the costs and benefits of environmental protection. We then consider how benefit–cost analysis has been used to evaluate policies in the real world. Finally, we explore the philosophical

justification for benefit–cost analysis and consider some of the most frequent criticisms lodged against its use. In particular, benefit–cost analysis focuses on the net benefits from a policy rather than its distributional consequences. Partly for this reason, economists do not advocate using a simple cost–benefit test as the sole criterion for policy decisions. Although it is a valuable source of information, benefit–cost analysis is just one of a number of tools to use in assessing policies or setting goals.

We then turn our attention more explicitly to markets: how they function, what they do well and what they do poorly, and how they can be designed to achieve desirable outcomes. We begin Chapter 4 with a key insight from economics: Under certain conditions, competitive markets achieve efficient outcomes. That is, they maximize the net benefits to society from the production and allocation of goods and services. This is a powerful result, and it helps explain the wide appeal of markets. It also aids understanding of the root causes of environmental problems: To an economist, they stem from well-defined failures in how unregulated markets incorporate environmental amenities. Moreover, it lays the groundwork for designing policies that rely on market principles to promote environmental protection.

The notion of “market failure” is the focus of Chapter 5. We discuss three ways of framing the types of market failure most common in the environmental realm: externalities, public goods, and the tragedy of the commons. In each case we offer a range of motivating examples. We then unify the discussion by showing how each of the three descriptions of market failure captures the same underlying divergence between individual self-interest and the common good.

In Chapter 6, we apply the concept of dynamic efficiency to the problem of the optimal rate of extraction of a nonrenewable natural resource, such as petroleum. We define scarcity in economic terms, which leads naturally to the concept of rent, the extra economic value imparted by scarcity. We illustrate the underlying similarities between nonrenewable resources and other capital assets and emphasize the powerful market incentives that encourage private owners of nonrenewable resources to account for scarcity in their extraction decisions.

Chapter 7 applies the same reasoning to two renewable resources, forests and fish. We develop bioeconomic models to demonstrate the efficient level of fishing effort and the efficient rotation period for a forest stand, both graphically and conceptually. In both cases, we include noncommercial benefits in an economic approach to efficient use of the resource.

Chapter 8 discusses the design of policies to overcome market failures

in the provision of environmental amenities and the management of natural resources. We start by considering a central debate in economics: Should the government intervene to solve market failures? After satisfying ourselves that the answer is yes, at least in many cases of real-world concern, we go on to review the various tools a government regulator has at her disposal, ranging from conventional command-and-control policies such as technology standards to market-based instruments such as taxes on pollution or resource use and tradable allowances. We discuss the intuition behind how these latter approaches can restore the efficient workings of the market. We close by contrasting the two market-based instruments, asking when prices or quantities are the preferable tool for governments to wield.

Chapter 9 continues our discussion of policy design but focuses more broadly on cases where efficiency may not be the objective. Even so, market-based instruments have two strong advantages: They can (in theory) achieve a desired level of environmental protection at the lowest total cost while spurring the development and diffusion of new technologies over the long run. We briefly consider a range of other factors relevant to the design of policy. Market-based instruments are not the solution to every problem, and we show when conventional command-and-control approaches are preferable even on strictly economic grounds. But the main conclusion is that market-based instruments are a crucial component of the regulatory toolkit.

Chapter 10 reviews the real-world performance of market-based instruments in regulating pollution and managing natural resources. We consider three cases in careful detail: the market for sulfur dioxide (SO_2) emissions from power plants in the United States, the tradable individual fishing quota (IFQ) system for New Zealand's fisheries, and municipal drought pricing of water resources in the United States. In each of these cases, we discuss the performance of the market-based approach, consider the implications for distributional equity, and assess the ease of monitoring and enforcement. We go on to review a longer catalog of examples, each in less detail than the initial case studies. Our aim is to equip readers to think broadly and creatively about the ways in which prices and markets can be injected into the regulatory process, aligning the incentives of firms and consumers with those of society in achieving environmental and resource management goals.

Chapter 11 addresses the links between economic growth and the natural environment—topics grouped under the heading of *macroeconomics*, in contrast to the *microeconomic* reasoning (based on the behavior of

individuals and firms) that characterizes most of the book. We begin by reviewing the debate over the limits imposed on economic growth by natural resource scarcity, focusing on the critical importance of two often overlooked factors: substitutability and technological change. The same key issues arise in our discussion of sustainability in economic terms. We highlight the insights of economic definitions of sustainability for current natural resource management and environmental protection. We end with a discussion of green accounting, emphasizing the need to incorporate natural resource depletion and changes in environmental quality into traditional measures of economic growth.

In the concluding chapter, we reflect on the relative roles of firms, consumers, and governments in the creation and mitigation of environmental and resource management problems. We then offer some final thoughts about the role of economic analysis as one of many important tools at the disposal of decision makers in environmental policy.

What We Hope Readers Will Take Away from This Book

If this is your first and last exposure to economics, and your interests lie in other areas of environmental studies, we offer three good reasons to use this text. First, many of the causes and consequences of environmental degradation and poor natural resource management are economic. That is, they arise from the failure of an unregulated market to give firms and individuals adequate incentives to promote environmental quality. Second, so-called market-based approaches to environmental regulation and natural resource management are increasingly common at local, national, and global levels. Prominent examples include the cap-and-trade policies used to limit sulfur dioxide pollution from U.S. power plants between 1995 and 2010, and CO₂ emissions in Europe, California, and elsewhere, and tradable fishing quotas to manage commercial fisheries. Third, economic arguments play an important role in some environmental policy debates, such as management of public lands and the structure of international approaches to counter global climate change. Without an understanding of basic economic principles, it is difficult to formulate an economic argument—or to refute one.

Thinking systematically about benefits, costs, and tradeoffs can improve your ability to tackle real-world environmental problems, even when it is not possible to estimate benefits and costs explicitly. The theory we introduce and the applications we discuss are meant to demonstrate this. Of course, our treatment of individual topics in this text is necessarily brief;

our intention is to give you just a basic grounding in the field. But we hope the information we do present will pique your interest and prompt you to explore environmental and resource economics in greater depth.

Reading this book will not make you an economist. Nonetheless, we hope to convince you that despite its reputation as a “dismal science,” economics can make vital contributions to the analysis of environmental problems and the design of possible solutions.

2

Economic Efficiency and Environmental Protection

Imagine that you are planning a spring break trip to the Bahamas, and you are choosing from among four vacation packages you have found on the web. The “Bahamas on a budget” trip, a 3-day affair staying in tent cabins, costs \$200. Suppose you would be *willing* to pay up to \$550 for that trip but no more. In other words, you wouldn’t care if you paid \$550 for the trip or spent the money on something else. The next step up is a trip that costs \$500. This trip includes 4 days’ lodging in beachfront cabanas, and the setting is so beautiful that you would be willing to pay up to \$900 for it. An even pricier 5-day trip, with a few extras thrown in, would cost \$850 and be worth \$1,100 to you. Finally, a deluxe week-long package is available for \$1,250, which on your student’s budget is just about the maximum you would be willing to pay for any vacation, although this package is so breathtaking, you might just be willing to pay that much for it.

Faced with these possibilities, which trip should you choose? At first glance, you might think that the deluxe trip is the best one to take; after all, you value it the most and are willing to pay the cost (even if only just barely). But in that scenario, you end up with zero *net* benefits. Indeed, because we have defined your “willingness to pay” as the amount for which you would be indifferent between paying for the trip and staying home, going on (and paying for) the week-long trip would make you no better off than if you didn’t take a vacation at all. Choosing the deluxe trip on the grounds that you would be willing to pay the most for it amounts to ignoring the costs of the vacation completely.

Instead of choosing the trip with the highest *gross* value to you,

regardless of cost, you would be better off choosing the trip that gives you the greatest *net* benefit—that is, the difference between the benefit of taking the trip (measured by your willingness to pay) and the cost (measured by its price). On these grounds, the best option turns out to be the 4-day \$500 trip, which you value at \$900, for a net benefit of \$400. This is greater than the net benefit from the more expensive \$850 trip: the added cost (+\$350) outweighs the increase in value (+\$200), so that net benefits decline to \$250. The \$500 trip is also better (from a net-benefit perspective) than the “budget” trip. Although that trip is cheaper, it is also worth less to you, and the drop in value is greater than the cost savings.

So how does this resemble an environmental problem? Well, imagine that, instead of taking a trip to the Bahamas, you are evaluating the possibilities for reducing pollution in your community, and there are a number of different options and price tags. As in the case of the vacation, a reasonable criterion for making decisions is maximizing net benefits. The net benefits of controlling air pollution, for example, are the difference between the total benefits of cleaner air and the total costs of reducing emissions. Maximizing the net benefits of a policy corresponds to the notion of economic efficiency. And as we’ll see in Chapter 3, willingness to pay is indeed at the heart of how economists conceive of and measure the value of environmental protection and natural resources.

You may be surprised to learn that if we accept economic efficiency as a reasonable goal for society, then the optimal level of pollution will in general be greater than zero. The reason for this will become clear as we proceed, but it can be summed up as follows: Although there would certainly be benefits from eliminating pollution completely, the costs would (in most cases) be much higher. We could get nearly the same benefit, at much lower cost, by tolerating some pollution.

Economic Efficiency

To an economist, answering the question “How much environmental protection should society choose?” is much like answering the question “Which vacation package is best?” in the simple example above (albeit on a much larger scale): It depends on comparing benefits and costs and finding where their difference is greatest.

This comparison between benefits and costs leads to a central concept in economics: that of *economic efficiency*. To an economist, an efficient policy or outcome is one that achieves the greatest possible net benefits. You should note that *efficiency* has a precise meaning here, which differs somewhat from common usage. In other contexts, *efficiency* connotes a

minimum of wasted effort or energy. For example, the energy efficiency of a home appliance is the amount of electricity the appliance uses per unit of output—for example, the amount of electricity used by an air conditioner to cool a room of a certain size. The less energy an appliance uses to produce a given outcome, the more energy-efficient it is. Similarly, the efficiency of a generator in an electric power plant measures how much useful energy a turbine generates, relative to the energy content of the fuel burned to drive the turbine. In both of these examples, efficiency is a function only of inputs and processes. The goal (cooling a room of a given size or generating a certain amount of electricity) is taken as given, and efficiency measures how little energy is used to achieve it. In other words, energy efficiency does not relate benefits and costs—the comparison at the heart of the concept of economic efficiency.

To illustrate this contrast, suppose you are choosing between a top-of-the-line air conditioner that costs \$500 and a model that uses more electricity but costs only \$150. The more expensive air conditioner is certainly more energy efficient. However, whether it is more efficient from an economic point of view—that is, whether the net benefits are greater—depends on how often you will use the air conditioner, how much more electricity the lower-end model uses, and the price of electricity.

To understand what economic efficiency means for environmental policy, let's start by considering a real-world environmental issue: sulfur dioxide (SO_2) emissions from fossil-fueled electric power plants. Burning oil or coal to generate electricity creates SO_2 as a byproduct, because those fuels contain sulfur. In downwind areas, SO_2 emissions contribute to urban smog, particulate matter, and acid rain. For these reasons, the control of SO_2 emissions from power plants has been a focus of air pollution legislation in the United States and many other countries.

From an economic perspective, we can frame this issue in terms of the efficient level of SO_2 emissions abatement. (It is often easier to think in terms of abatement, or pollution control, which is a “good,” rather than pollution, which is a “bad.”) Suppose we observe the amount a firm or industry would pollute in the absence of any regulatory controls. Abatement is measured relative to that benchmark. If a firm would emit a thousand tons of pollution in the absence of regulation but cuts that to six hundred tons of pollution (for example, by installing pollution control equipment), it has achieved four hundred tons of abatement.

What level of sulfur dioxide abatement will maximize net benefits to society? To answer this question, of course, requires thinking systematically about the costs and benefits of pollution control.

The Costs of Sulfur Dioxide Abatement

Typically, a minor amount of abatement can be achieved at very little cost simply by improving how well a power plant burns coal, because a cleaner-burning plant will emit less pollution for any given amount of electricity generated. (One reason the resulting abatement is cheap is that a cleaner-burning plant will also use less fuel to produce the same amount of electricity, saving money for its managers.) At a somewhat higher cost, power plants can increase their abatement by burning coal with slightly less sulfur than they would otherwise use. The abatement cost increases further as the power plant burns coal containing less and less sulfur that is more and more expensive. For example, a power plant in Illinois can burn cheap high-sulfur coal from mines in the southern part of the state. To reduce SO_2 pollution, such a plant might switch to coal from eastern Kentucky with half the sulfur content but a slightly higher transportation cost. Still greater reductions could be achieved, at still greater cost, by switching to very low-sulfur coal from Wyoming. Finally, achieving reductions of 90 percent or more from baseline levels typically requires investment in large end-of-pipe pollution control equipment, such as flue gas desulfurization devices (better known as scrubbers) that remove SO_2 from the flue gases. Such equipment is often very expensive, making high levels of abatement much more costly than low levels. Moreover, the cost is typically driven by the percentage reduction achieved, so that removing the first 90 percent of pollution costs about the same as going from 90 to 99 percent removal.

The costs we just described trace out a particular pattern. Costs rise slowly at first, as abatement increases from zero. As abatement continues to increase, however, costs rise more and more rapidly. This pattern

*To an economist, being efficient
means maximizing net benefits.*

is reinforced when we consider the costs of abatement at the level of the industry rather than the individual power plant. Some power plants (those located close to low-sulfur

coal deposits, for example) can abate large amounts of pollution at low cost, whereas others may find even small reductions very expensive. As we increase pollution control at the industry level, we must call on plants where abatement is more and more expensive.

Figure 2.1 depicts a stylized *abatement cost function* that corresponds to this pattern of rising cost. By *abatement cost function* we mean the total cost of pollution control as a function of the amount of control achieved. In

The Energy Efficiency Gap

The difference between what economists mean by *efficiency* and what engineers and others often mean is illuminated if we think about the concept of energy efficiency. Many studies have estimated significant private net benefits to technical energy efficiency investments by households and firms, including things such as switching from incandescent lightbulbs to compact fluorescent lamps (CFLs), installing more effective insulation, and buying more efficient appliances. Outside economics, analysts often wonder why these investments don't happen on a larger scale, identifying an energy efficiency "gap" between what would appear to be cost-minimizing and actual energy efficiency investments. The solution, according to these analyses, is a broad effort by the public sector to reduce barriers to the adoption of energy efficient technologies, through education or information provision, subsidies, and other policies.

In response, economists point to several problems with this perspective. We'll discuss a few here.¹ First, analyses that identify this gap usually rely on engineering estimates of the potential energy cost savings associated with efficiency investments, and real-world savings often differ from potential savings. As we will explore in greater detail in Chapter 3, economic costs are opportunity costs, which would include perceived risks from new technologies (for example, if your usual plumber is not willing or able to install a tankless hot water heater), changes in the quality of the produced service (as with the change from incandescents to CFLs), and other costs—not simply the dollars spent on your energy bill. These costs, though hard to quantify, are real economic costs not accounted for in technical efficiency studies. Second, energy use behavior changes when households and firms purchase more efficient technologies; a rebound effect of increased usage due to lower operating cost has been observed for many energy technologies. Thus, both energy savings and cost savings in the real world will differ from engineering estimates of potential savings. Third, the rate at which energy consumers are willing and able to trade the future benefits of reduced energy costs for current investment costs is poorly understood; in particular, low-income households may face significant credit constraints and steeper consequences for this tradeoff than others. In addition, to the extent that energy and cost savings from efficient technologies have been estimated from households and firms that have adopted these technologies, the results of these studies may not be generalized to nonadopters. The inherent bias could go either way: Those who adopt energy efficient technologies may be "conservation-oriented," or they may be energy "hogs" who purchase efficient technologies to support increased use (at lower cost).

The Energy Efficiency Gap *continued*

The point is not that households and firms in the real world always make economically efficient decisions about energy technology investments. Consumers may lack the information necessary to understand how energy efficiency varies between different appliances or how that translates into potential savings; other characteristics of those appliances may seem more salient at the time of purchase. Incentives may not be properly aligned: For example, renters will lack sufficient incentive to install energy-efficient technologies, knowing that some of the benefits will accrue to landlords and future occupants. But it is difficult to tell from data on the *technical* efficiency of these investments—both how much energy they would save if operating according to engineering specifications and how much these savings would reduce the total cost of energy consumption—how large the *economic* energy efficiency gap might be.

the figure we have used X to represent the amount of pollution control and $C(X)$ to denote the total cost (in dollars) as a function of X . A function with this bowed-in shape is called a convex function.

The Benefits of Sulfur Dioxide Abatement

Recall that in Chapter 1 we described the benefits from reducing greenhouse gas emissions as corresponding to the avoided damages from global climate change. In the same way, the benefits of SO_2 abatement are simply the avoided damages from pollution.

How do these damages vary with pollution? As the air gets dirtier, pollution damages tend to increase more and more rapidly. At low concentrations, SO_2 corrodes buildings and monuments. Higher concentrations lead to acid rain, with the attendant damages to forest ecosystems from the acidification of lakes and soils. In urban areas, the adverse effects of SO_2 increase from eye and throat irritation, to difficulty breathing, and ultimately to heart and respiratory ailments. These effects are felt first by the most vulnerable members of society: infants, older adults, and asthmatics. But as concentrations rise, the affected population grows.

This pattern of damages corresponds to total *benefits* from pollution control that increase rapidly when abatement is low (and pollution is high) and increase more slowly when abatement is high (and pollution is low). This is illustrated by the curve in figure 2.2, where we have used $B(X)$ to represent the *abatement benefit function*. A function with the bowed-out shape of $B(X)$ is called a concave function.

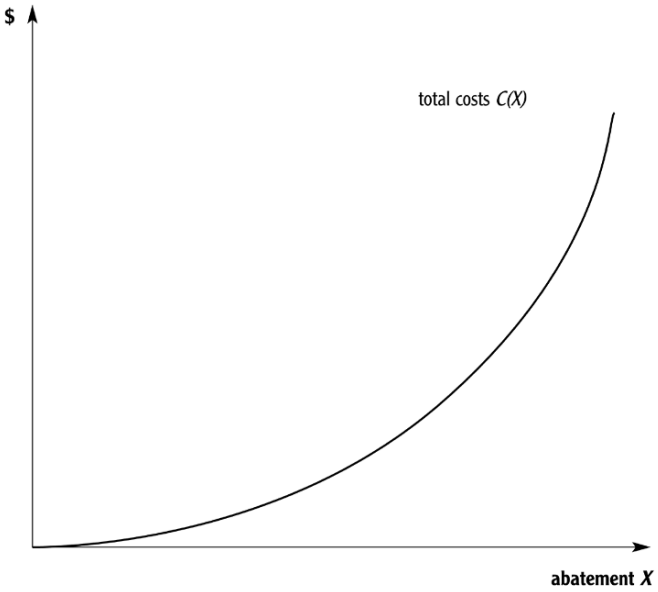


Figure 2.1 Total costs of pollution abatement, as a function of the level of abatement.

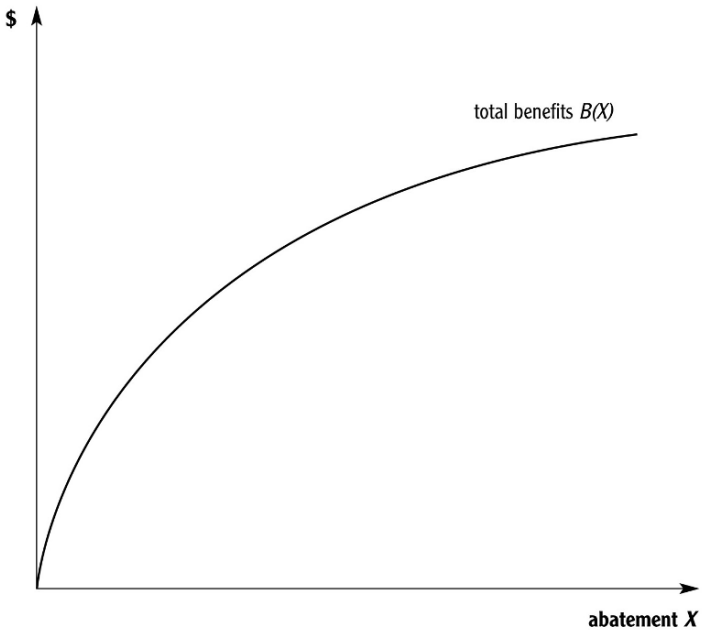


Figure 2.2 Total benefits of pollution abatement, as a function of the level of abatement.

Putting Costs and Benefits Together: Economic Efficiency

We are now ready to answer the question we posed earlier: What is the efficient level of sulfur dioxide abatement? To answer this question, we must compare benefits to costs and find where the difference between them—net benefits—is greatest.

Figure 2.3 places the cost and benefit curves in figures 2.1 and 2.2 on a single pair of axes. As in the previous figures, abatement increases as we move along the horizontal axis from left to right; pollution increases as we move from right to left. We have denoted *maximum abatement*—equivalent to zero pollution—by X^{MAX} .

Recall that net benefits are simply benefits minus costs. Thus on the figure, the net benefit from a given level of pollution control is measured by the vertical distance from the benefit curve down to the cost curve. At low levels of pollution control, net benefits are small. As abatement increases from a low level, the benefits increase more rapidly at first than do the costs, so that net benefits increase. As more and more abatement is done, however, the benefits rise less rapidly, while the costs of abatement increase. Eventually, the benefits increase more slowly than costs, and net benefits *fall* as more and more abatement is done.

In between those two extremes, of course, the difference between benefits and costs must reach a maximum. On our graph, this happens at level X^* . By definition, this is the efficient level of pollution control. You can see from the figure that X^* is greater than zero but less than the maximum possible abatement. Accordingly, the efficient level of pollution must also be less than its maximum (unregulated) level but greater than zero. We come right away to the point that we mentioned at the outset of the chapter:

- In general, the economically efficient level of pollution is not zero.

Zero pollution is not efficient (in general), because the gains from achieving it are not worth the extra cost required. Consider increasing abatement from the level X^* to the level X^{MAX} . In our real-world example, this might correspond to installing expensive scrubbers on every power plant. This much abatement would certainly bring benefits, such as reductions in acid rain and improvements in urban air quality. On the graph, the increase in benefits is shown by the fact that curve $B(X)$ increases as we move to the right, so that $B(X^{\text{MAX}}) > B(X^*)$.

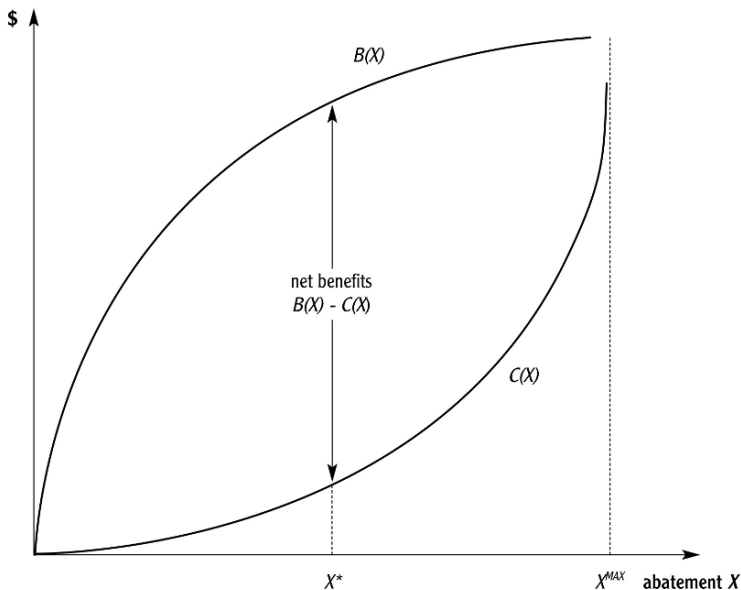


Figure 2.3 The efficient level of pollution abatement, denoted X^* , achieves the greatest possible net benefit.

However, those extra benefits from maximizing abatement are outweighed by the extra costs of achieving them. While benefits increase, costs rise even faster. As a result, the gap between benefits and costs shrinks dramatically as we increase abatement from X^* to X^{MAX} . In the real world, requiring scrubbers on all power plants would raise costs by an order of magnitude, and the boost in benefits would be much smaller.

Therefore, zero pollution is generally not desirable—at least not if we measure the success of our policy by the magnitude of its net benefits. Of course, it is equally true (although perhaps less surprising) that zero abatement is also not efficient. Abating less than X^* would reduce costs, but the cost savings would be less than the forgone benefits. On balance, net benefits would fall.

If you find these results surprising or counterintuitive, it may help to recall the distinction between economic and technical notions of efficiency. Pollution is sometimes described as “inefficient” when the pollution represents a form of wasted inputs. For example, a key component of water pollution from paper mills or textile factories is excess chemicals used in the production process—bleach in the case of paper mills, dye in

the case of textile factories. Although such pollution may be “inefficient” in a technical sense, it is a mistake (albeit a common one) to conclude that it is also necessarily inefficient in an *economic* sense. Economic efficiency depends on the costs as well as the benefits of controlling pollution. If it is extremely costly to clean up pollution completely, zero pollution is unlikely to be a reasonable goal if we aim to maximize net benefits.

Efficiency and Environmental Policy

In our example of SO_2 pollution from power plants, the benefits from abatement rise rapidly at first and then tail off, while costs rise much more slowly at first before becoming steep. Put them together, and we find that net benefits are greatest somewhere in the middle. Because the shapes of the cost and benefit curves are critical in driving the results, it is worth discussing them in a bit more detail.

The pattern of “increasing costs at an increasing rate” is common. The costs of producing most goods—for example, steel or shoes—typically increases with production at an increasing rate (at least in the short run and over some range of quantities). In the case of pollution control, you can think of “clean air” as the good that is being produced: Clean air is costly, and the costs rise more and more steeply as the air gets cleaner and cleaner. Removing the last few ounces of pollution from a waste stream is likely to be prohibitively expensive.

On the benefit side, meanwhile, assuming a concave benefit function corresponds to the simple idea that although we would usually like more of a good thing, the amount we are willing to pay for something is likely to decline as we get more of it. You would probably pay more for one pair of designer shoes or one pair of tickets to a rock concert than you would pay for the second, third, or tenth pair of the same item.

These characteristics of costs and benefits apply in a wide range of cases in the environmental realm—not just other forms of air pollution but also water pollution, biodiversity preservation, endangered species protection, the management of natural resources such as fisheries, and so on. For example, consider the protection of habitat for an endangered species such as the red-cockaded woodpecker, which lives in old-growth stands of longleaf pine forest in the southeastern United States. Habitat protection requires managing forests to maintain suitable old-growth conditions. The cost of such management varies widely between different parcels of land, depending on ownership, suitability for intensive timber production, soil conditions, and so on. If we arrange lands from least to greatest expense, we can construct an increasing cost-of-protection function similar to the

one in figure 2.1. Similarly, on the benefits side, an increase in the woodpecker population from one hundred birds to two hundred birds is likely to yield much greater benefits than from one thousand birds to eleven hundred birds, leading to a benefit-of-protection function much like the curve in figure 2.2.

Accordingly, although we will continue to discuss our model in terms of pollution control or abatement, you should keep in mind that it is much more general than that. For convenience, we will continue to refer to X as pollution control or abatement, but you could substitute any other dimension of environmental quality, such as habitat protection, and the arguments that follow would still apply. The crucial assumptions underlying our model are that costs increase at an increasing rate and that benefits increase at a decreasing rate—in other words, that the total cost function $C(X)$ is convex and the total benefit function $B(X)$ is concave, like those drawn in figures 2.1 through 2.3.

In some cases, these assumptions do not hold. For example, think of litter along a hiking path in a wilderness area. One piece of trash may ruin an otherwise pristine area nearly as much as ten or twenty pieces would. In this case, the marginal benefit of environmental quality does not fall as the amount of trash gets smaller (until the trash goes away completely). Hence the efficient level of litter might well be zero.

A particularly important exception to the conventional rule “equate marginal benefit and cost” arises when the marginal cost of cleanup falls (instead of rising) as more cleanup is done. Cost functions with this characteristic are said to exhibit *economies of scale*. For example, cleaning up hazardous waste sites typically requires digging up the soil and incinerating it to remove the pollution. The cost of such a cleanup depends mostly on the area of the site rather than how contaminated it is or how much pollution is removed. In such a case, pollution control may be an all-or-nothing exercise: If it makes sense to clean up a site at all, then it makes sense to clean it up completely. Over time, this policy would look very different from that of the standard case of increasing marginal cost. Rather than seeking to maintain environmental quality at the level where marginal cost and benefit are equal, the optimal policy would let quality decline over time and then periodically clean things up to a very high level of quality.²

Even if the cost and benefit functions have their typical shapes, of course, one can draw particular examples in which the maximum level of abatement is reached before net benefits start declining—or, conversely, in which net benefits are highest when abatement is zero. (Imagine taking

A very useful way to describe the costs and benefits of pollution control is in terms of marginal—that is, incremental—costs and benefits.

the curves drawn in figure 2.3 and shifting them rightward or leftward while holding the axes and the location of maximum abatement fixed.) However, there are good reasons to view these instances as special cases, as we have already seen. The model

of convex costs and concave benefits presented here is widely accepted as the conventional general model of the costs and benefits of pollution control (and of environmental protection more generally).

Equating Benefits and Costs on the Margin

So far, we have discussed the *total* costs and benefits of pollution control. An alternative and very useful way to describe the costs and benefits of pollution control is in terms of *marginal* costs and benefits. By *marginal cost* we simply mean the cost of an incremental unit of abatement. If we have abated one hundred tons, the marginal cost is the cost of the one-hundredth ton. (Note the contrast with average cost, which takes into account all of the abatement done rather than only the last unit.) Likewise, *marginal benefit* refers to the benefit from the last unit of abatement. Recall that efficiency corresponds to maximizing the difference between total benefits and costs. It turns out that this difference is greatest when marginal benefit and marginal cost are equal.

Marginal Costs and Marginal Benefits

Let's start by considering the relationship between total cost and marginal cost. Because marginal cost measures the cost of one more unit of abatement, it corresponds to the *slope* of the total cost function. To see why this makes sense, consider the cost function depicted in figure 2.1. At low levels of abatement, where the total cost function is nearly flat, the height of the curve changes little as pollution control increases. Therefore, each additional unit of pollution control adds a small amount to the total cost. In other words, the marginal cost of pollution control is small. At higher levels of abatement, the total cost function is steep, so that the cost rises rapidly as abatement increases. This means that the incremental cost of pollution control—the marginal cost—is high.³

Figure 2.4 plots the *marginal cost function* corresponding to a total cost function like that in figure 2.1. As before, abatement is on the horizontal axis, but now the vertical axis measures marginal rather than total cost. Thus the height of the curve $MC(X)$ at any given point represents the

Thinking on the Margin: Pollution Abatement at Aracruz Celulose, S.A.⁴

One of the mainstays of economic reasoning is learning to think in terms of marginal changes when making decisions. To find the level of production that maximizes its profits, for example, a firm needs to compare the revenue from selling one more unit of the good with the cost of making it. Similarly, to find the amount of abatement that maximizes net social benefits, we must compare the marginal benefit from controlling another ton of pollution with the marginal cost.

To make the concept of marginal cost (in particular) more concrete, consider the case of pollution abatement at pulp mills owned by Aracruz Celulose, S.A., a leading Brazilian pulp producer and exporter. Among the major pollutants in effluent from pulp mills are chlorinated organic compounds, known as adsorbable organic halides (AOX). These compounds—dioxin is among the most infamous—are produced when chlorine-containing chemicals used in bleaching react with wood fiber.

In the early 1990s, Aracruz was considering whether to upgrade its environmental controls in order to market its pulp to environmentally conscious customers in Europe. The company had three primary options: continuing to produce standard pulp using chlorine, switching to “elemental chlorine free” (ECF) methods using chlorine dioxide, and eliminating chlorine entirely (“totally chlorine free” [TCF]) by using peroxide as a bleaching agent. These were cumulative efforts: The investments needed to produce ECF pulp were a necessary prerequisite to TCF bleaching. The following table shows the pollution level associated with each option, the corresponding abatement, the total cost, and the marginal cost—that is, the cost per additional unit of abatement. As the table shows, switching to ECF pulp cuts pollution by 80 percent at a fairly low cost. Converting to TCF could cut pollution by an additional 95 percent, but the cost per ton increases significantly.

Alternative	Pollution (AOX, in kg/year)	Incremental abatement (kg/year)	Total annual cost	Increase in total annual cost	Marginal cost (per kg additional abatement)
1. Standard pulp (baseline)	2,000,000	No reduction	\$0	\$0	\$0
2. ECF pulp	400,000	1,600,000	\$575,000	\$575,000	\$0.36
3. TCF pulp	20,000	380,000	\$5.325 million	\$4.75 million	\$12.50

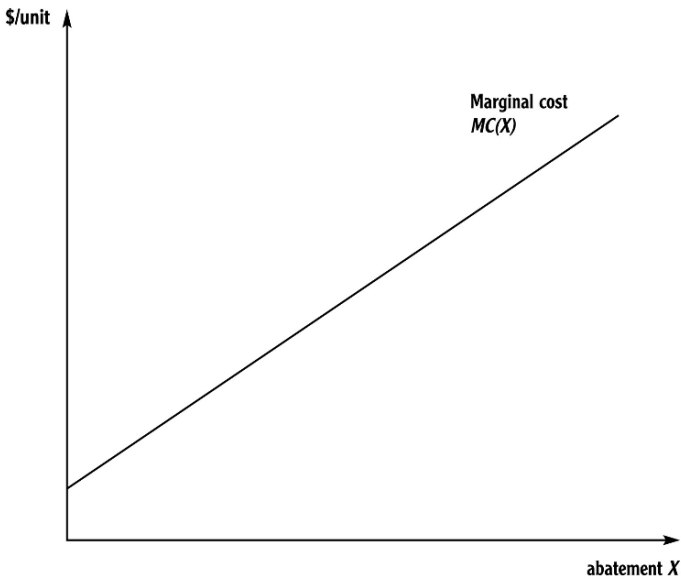


Figure 2.4 Representative marginal abatement cost function.

cost of each additional unit of abatement. Saying that abatement cost increases at an increasing rate is the same thing as saying that abatement has *increasing marginal costs*: Each ton of pollution abatement costs slightly more than the one that preceded it. As a result, the marginal cost function in figure 2.4 slopes upward.

In a similar fashion, we can derive a *marginal benefit function* that corresponds to the incremental benefits of additional abatement. Marginal benefit corresponds to the slope of the total benefit function. If the benefit function is concave, as in figure 2.2, then the marginal benefit function will be downward sloping: Each additional ton of abatement brings smaller additional benefits. We have drawn a representative function, labeled $MB(X)$, in figure 2.6.

Efficiency and the Equimarginal Rule

Let's take another look at figure 2.3, where we plotted the benefit and cost functions and found the efficient level of abatement X^* . Notice that as abatement increases up to X^* , the benefits of pollution control rise faster than the costs. That is, the $B(X)$ curve is steeper than the $C(X)$ curve. As a result, net benefits increase with each additional ton of pollution control over this range. On the other hand, beyond the efficient point, the costs rise faster than the benefits, so that net benefits diminish. Putting these

The Costs of Protecting the California Condor

With a wingspan of nine-and-a-half feet, the California condor (*Gymnogyps californianus*) is the largest bird in North America.⁵ Until the mid-nineteenth century the condor's range extended as far north as the Columbia River Gorge and south into Baja California. Indeed, the diaries of Meriwether Lewis and William Clark report several sightings of the "Buzzard of the Columbia" in 1805 and 1806. Throughout the twentieth century the wild population declined precipitously, falling from approximately one hundred birds in the 1940s to only nine by 1985. The decline appears to have been caused by reduced reproduction (perhaps a result of DDT) and human-created mortality, including lead poisoning from bullets in game carcasses, shooting of the condors themselves, and hazards from human-made structures such as power lines.

In the late 1980s, the U.S. Fish and Wildlife Service captured the remaining wild birds and embarked on a captive breeding program, with the hopes of eventually reintroducing the species into the wild. In 1992, the first two captive-bred juveniles were released into the Sespe Condor Sanctuary in Los Padres National Forest. By October 2003, the wild population had climbed to eighty-three birds, including one chick hatched in the wild.

With the condors back in the wild, measures must be taken to protect the condor populations from threats. From an economic point of view, we can think of these protective measures as "abatement"—in this case, abatement of the causes of condor mortality. Abatement measures include the protection of suitable habitat, provision of food carcasses such as stillborn calves (to prevent lead exposure), promotion of alternatives to lead ammunition, prohibitions on shooting the condors, and modification of power lines and other human structures to reduce injuries to condors.

One study has estimated the costs of abatement using information contained in the Recovery Plan written by the U.S. Fish and Wildlife Service. For each abatement action, the number of condors saved per year was estimated taking into account historical rates of decline in the condor population and the priority accorded that action by the U.S. Fish and Wildlife Service. Unit cost (per condor per year) was then calculated by dividing the cost by number of condors saved. Arranging the unit costs in increasing order produces a marginal cost function, as illustrated by figure 2.5.

The figure illustrates two key points. First, note the wide range in the marginal costs of various techniques: from as little as \$7 per condor saved per year to protect habitat in low-lying areas to more than \$200 per condor per year to modify power lines and step up law enforcement. Second, note that it is the marginal cost, rather than the total cost, that determines which measures should be pursued first. Thus, although the annual cost of heightened law enforcement is only a quarter of the cost of removing contaminants (\$5,000 versus \$20,000), contaminant removal would save more than thirty times as many condors and hence is a much more cost-effective means of protecting the species.

Costs per condor
per year (\$000s)

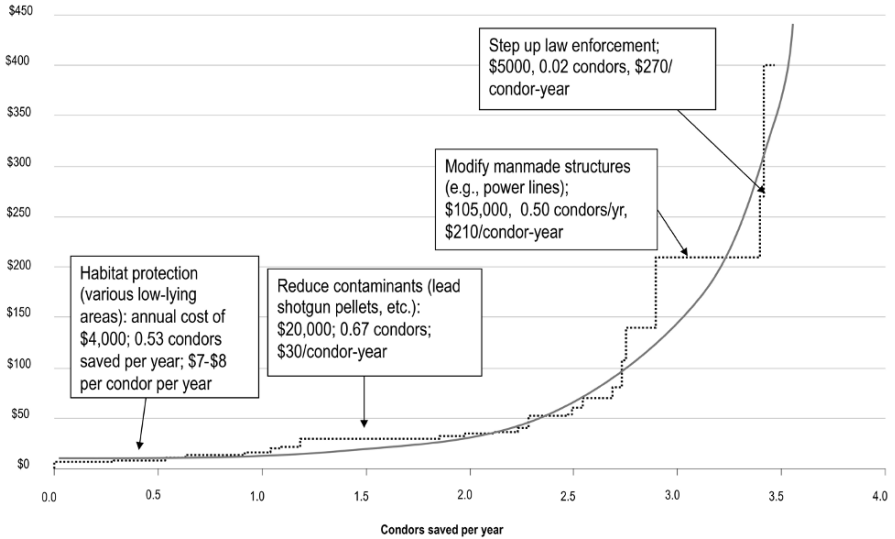


Figure 2.5 Marginal cost graph for condor example. Each “step” on the dashed line corresponds to a specific protection measure, arranged from lowest to highest unit cost. The boxes highlight four specific actions among over two dozen considered. The solid line represents a smooth approximation to the “staircase” function.

observations together, we conclude that at the efficient level of abatement, the benefit and cost curves must have the same slope.

This suggests a way to find the efficient level of pollution control by looking at the marginal benefits and costs. In particular, we can state the *equimarginal rule*:

- The efficient level of abatement X^* occurs where marginal benefit equals marginal cost, that is, $MB(X^*) = MC(X^*)$.

In plain English, this says that the efficient level of pollution control is where the extra benefit of the last unit of abatement done equals its extra cost. Beyond that point, the incremental costs of any further abatement will outweigh the incremental benefits. This result is illustrated by figure 2.7. The top panel is the same as figure 2.3. The bottom panel draws the corresponding marginal benefit and cost curves. The efficient point X^* is easily identified: It is where the MB and MC curves cross.

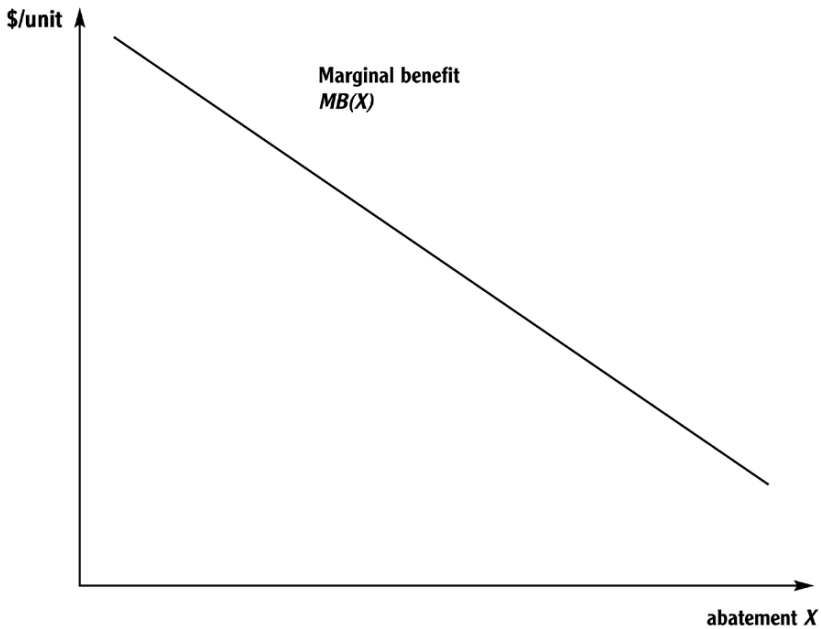


Figure 2.6 Representative marginal abatement benefit function.

This equimarginal condition will show up again and again in our analyses of markets and policy design, so it is worth going over the intuition behind the result. Suppose we pick a low level of abatement, where MB is greater than MC —say the point X_L on figure 2.7. Now let's imagine increasing abatement by 1 ton. What happens to net benefits? Because the resulting increase in benefits (equal to the marginal benefit) is greater than the increase in cost (= marginal cost), net benefits would increase. Thus at X_L efficiency increases with more abatement, as indicated by the arrow on the figure.

Now suppose that we increase abatement all the way to some high level, such as X_H on the figure, where MB lies below MC . Here, one more ton of abatement increases costs by more than it increases benefits; therefore, the incremental net benefit is negative. Indeed, at such a point we could increase net benefits by reducing abatement by one unit, because costs would fall by MC , but benefits would decline by only MB . Thus, at X_H we have overshot the efficient level of abatement.

Of course, we could repeat these arguments for any values of abatement above or below the point where the marginal curves cross. Only

The Benefits of Mitigating Stratospheric Ozone Depletion

Whereas tropospheric or ground-level ozone is a local air pollutant that causes human health damages including respiratory and cardiovascular ailments, stratospheric ozone (the “ozone layer” in the atmosphere, from 6 to 30 miles above the earth’s surface) protects the earth from some of the sun’s harmful ultraviolet radiation. In 1974, two chemists published research suggesting that the ozone layer could be destroyed by the release of chlorofluorocarbons (CFCs), ubiquitous chemicals used (at the time) in applications as diverse as air conditioning, asthma inhalers, hairspray, and styrofoam coffee cups. Damaging effects of this phenomenon included increased incidence of skin cancer and cataracts and reductions in the productivity of farms and fisheries. Some countries restricted consumption and production of CFCs in the late 1970s and early 1980s, but a 1985 report by the British Antarctic Survey that observed 40 percent thinning of the ozone layer over Antarctica between 1977 and 1985 stunned the world and led to the 1987 Montreal Protocol on Substances That Deplete the Ozone Layer.⁶

Countries including the United States and Canada performed their own analyses of the domestic benefits and costs of compliance with the Montreal Protocol, as well as other CFC abatement choices. Independent U.S. analyses were performed by the Environmental Protection Agency (EPA) and the Council of Economic Advisors under President Ronald Reagan, with similar results. The EPA study considered the benefits and costs to the United States of a global freeze on CFC production and consumption, as well as global reductions of 20 percent, 50 percent, and 80 percent.⁷ The EPA study monetized the value of avoided cases of cataracts and fatal and nonfatal skin cancers, avoided crop damage, avoided reductions in commercial fish harvests, and other anticipated impacts, although about 98 percent of the monetized benefits were associated with avoided skin cancer mortality. In Chapter 3, we’ll describe the methods used for monetizing these types of benefits in detail. But an examination of the benefits EPA estimated at varying levels of CFC reduction, described in the following table, provides a helpful illustration of the relationship between total and marginal benefits.

Policy Alternative	Total U.S. benefits (billions of \$1985)	Incremental abatement (percentage change)	Marginal benefit (billions of \$1985)
1. Global CFC freeze	5,995	—	—
2. Global CFC 20% reduction	6,132	20	6.85
3. Global CFC 50% reduction	6,299	30	5.57
4. Global CFC 80% reduction	6,400	30	3.37

The Benefits of Mitigating Stratospheric Ozone Depletion *continued*

Notice that the total benefits of CFC abatement increase monotonically as we move from zero to 80 percent. It is also clear that these benefits increase at a decreasing rate; as we reduce emissions more and more, the incremental benefit of further reductions shrinks. This is seen more clearly in the marginal benefit column, which simply divides the change in total benefit by the percent change in abatement for each alternative. Like the representative function in figure 2.6, the marginal benefit of CFC reductions decreases with abatement.




Figure 2.7 The efficient level of abatement, represented in terms of total costs and benefits (top panel) and marginal costs and benefits (bottom panel).

at the efficient point X^* , where $MB = MC$, is the difference between benefits and costs at its maximum.

Relating Marginal Benefits and Costs to Total Benefits and Costs

We have just seen how marginal benefits and costs correspond to the slopes of the total benefit and cost functions. Conversely, total benefits and costs can be represented as the *areas under the marginal benefit and cost curves*. Recall that the height of the marginal benefit curve (for example) at a given level of abatement represents the additional benefit derived from that unit of abatement. Imagine drawing a rectangle with a width equal to one unit of abatement and height equal to the height of the MB curve. The area of that rectangle would be equal to the marginal benefit of the corresponding unit of abatement.

Now imagine drawing a series of such rectangles, one for each unit of abatement, starting from zero and going up to X_L . Because the area of each rectangle represents the additional benefit from a certain unit of abatement, their areas must sum to the total benefit from X_L units of abatement. But the sum of the areas of the rectangles is also equal to the area under the curve.⁸ Thus the area under the marginal benefit curve from zero to any point equals the total benefit from that amount of abatement. Similarly, the area under the marginal cost curve from zero to any point is the corresponding total abatement cost.

This relationship between marginals and totals can give us another perspective on the equimarginal condition for efficiency. Let's return to the bottom panel of figure 2.7. At the efficient level of abatement (the point X^*), total benefits equal the area under the MB curve, and total costs are the area under the MC curve. Subtracting costs from benefits leaves total net benefits (the shaded triangle to the left of the intersection of the two curves). You can see right away that no other level of abatement provides as much net benefit as X^* . Less abatement leaves some net benefits unrealized. At X_L , for example, net benefits are smaller than at X^* by the area of the triangle labeled abc on the figure. Beyond X^* , the extra costs outweigh the extra benefits. At X_H , net benefits are smaller than they are at X^* by the triangle cde .

Dynamic Efficiency and Environmental Policy

So far, we have discussed the efficiency rule—set marginal costs equal to marginal benefits—in terms of maximizing the net benefits of a resource

(such as clean air or water) at a particular point in time. But projects and policies often have streams of benefits and costs occurring at many different points in time. For example, if we choose to set aside a large tract of land, such as the Arctic National Wildlife Refuge in Alaska, disallowing commercial uses in favor of wilderness and recreation, society will receive benefits and incur costs from this designation over many years, or even in perpetuity.

When benefits and costs vary over time, economic analysis must apply the rules of *dynamic efficiency*.⁹ For example, so-called stock pollutants that accumulate in the environment—such as carbon dioxide in the earth’s atmosphere or polychlorinated biphenyls (PCBs) in a riverbed—involve streams of benefits and costs over a very long period of time. Dynamic efficiency plays a particularly important role in the management of natural resources. Some resources, such as petroleum, do not regenerate at all (at least over time scales relevant to human activity); for others (such as fisheries), natural regeneration must be balanced against extraction and consumption. In both cases, the limited availability of the resources means that the amount available tomorrow depends on what we consume today. In order to apply the concept of efficiency in a dynamic setting, we must introduce the concept of *discounting*.

Discounting and Present Value

The introduction of a time dimension requires an additional step in thinking about efficiency. In the static analysis earlier in this chapter, we maximized net benefits. In a dynamic setting, an efficient policy maximizes the *present value* of net benefits to society. That is, we must convert all the benefits and costs of a potential environmental policy, no matter when they occur, into their dollar value *today* before summing them up. In this way, we use a common yardstick to measure benefits and costs occurring at different points in time.

To see why the value of a dollar today is not the same as the value of a dollar received next year, consider the following thought experiment. Suppose we offered you the choice of being paid \$100 today or the same amount a year from now. Which option would you choose? What if the choice were between \$100 today and \$105 a year from now? \$110? If you are like most of our students, you would take \$100 today over the same amount a year from now. You would

When benefits and costs vary over time, economic analysis must apply the rules of dynamic efficiency.

probably also prefer \$100 today to \$105 in a year, although as the future amount increased, you would find it more attractive to wait.

Now ask yourself: *Why* do you prefer money today to the same or even a slightly larger amount in the future? You can probably come up with several reasons. First, you might prefer money today because you can get the immediate benefit of spending it on something you value (a ticket to a concert or the theater, a piece of clothing, a meal at a good restaurant). Second, you might prefer money today because you anticipate having more money in the future, making an extra dollar today worth more. (Although that might not be a factor in our simple thought experiment of getting paid now or in a year, it probably *is* relevant to how much you would value money now rather than in 10 or 20 years.) Third, you might prefer money today because you could invest it today and earn a rate of return, whether from a savings account or by investing in the stock market. Each of these reasons illustrates a different facet of the *time value of money*.¹⁰

The time value of money is the reason that we discount costs and benefits expected to occur in the future. You are probably familiar with the power of compound interest. Discounting entails thinking in reverse. To see how this works, consider a simple example. Suppose you invested \$100 at an annual interest rate of 5 percent; how much would that investment be worth in 50 years? We can calculate the future value (*FV*) as follows, where *PV* is the present value, *r* is the interest rate, and *t* is the year.

$$FV = PV(1+r)^t = 100(1+.05)^{50} = \$1,146.74.$$

This equation simply says that \$100, growing at an annual rate of 5 percent, will yield \$1,146.74 in 50 years. Applying that logic in reverse, if we asked how much we needed to invest today at a 5 percent interest rate to have \$1,146.74 in 50 years, the answer would be \$100. That suggests that \$1,146.74 in 50 years, given a discount rate of 5 percent, has a *present value* of \$100.

You have probably already realized that the choice of discount rate is crucial. The discount rate reflects how much weight we put on future costs and benefits relative to those that occur today: The higher the discount rate, the *less* weight is put on the future. There is a rich literature in economics, with a wide range of views, on the correct discount rate to use in assessing public policies, especially those with long time horizons (such as policies to reduce greenhouse gas emissions to mitigate future climate change).¹¹

The Incredible Shrinking PV: The Influence of the Discount Rate

The choice of discount rate can have a surprisingly large effect on the present value (PV) of future costs or benefits, especially when those costs or benefits come many years in the future. The following table illustrates this point. For example, the PV of \$1,000 received 100 years from now is \$138 using a discount rate of 2 percent but barely more than a dollar using a discount rate of 7 percent.

Discount rate	Present value of \$1,000			
	T years from now			
	T = 10	T = 50	T = 100	T = 200
1%	\$905	\$608	\$370	\$137
2%	\$820	\$372	\$138	\$19
3%	\$744	\$228	\$52	\$2.7
5%	\$614	\$87	\$7.6	\$0.06
7%	\$508	\$34	\$1.2	\$0.001
10%	\$386	\$8.5	\$0.07	\$0.00001

Broadly speaking, one school of thought holds that in evaluating public policies, analysts should use discount rates based on the returns to alternative investments that could, in principle, be made instead. In this view, discounting effectively asks whether the returns to a project, policy, or other investment, such as a greenhouse gas emission regulation, the establishment of a new national park, or the decision to pump groundwater from a nonrenewable aquifer, are greater or less than the returns to investing in education, building a new hospital, or simply placing an equivalent amount of funds in an interest-bearing asset such as Treasury bills. If the answer to this question is “no,” we can do better by choosing that alternative investment today and letting future generations decide how to invest the returns.

Another school of thought views the discount rate as a *normative* decision that should take into account deliberative judgments about appropriate rates of time preference, equity between present and future generations, and so on. We do not take a position here, except to note that good practice in policy analysis is to apply a range of discount rates rather than to choose a single one.

The Equimarginal Rule in a Dynamic Setting

The equimarginal rule we discussed previously still applies in the dynamic setting, although we must convert marginal benefits and marginal costs into present value terms in order to compare the magnitude of streams of benefits and costs over time. In a dynamic context, efficient environmental policy equates the *present value of marginal benefits with the present value of marginal costs*. We will explore real-world applications of the dynamic equimarginal rule in Chapter 6, when we approach the problem of non-renewable resource extraction, and in Chapter 7, when we discuss the economics of forests and fisheries.

Conclusion

This chapter has laid the groundwork for everything that follows. When economists talk of efficiency, they have something very specific in mind: maximizing net benefits. As we have seen, in a static setting net benefits are largest (in general) when the benefits and costs of environmental protection are equal *on the margin*. In a dynamic setting, net benefits are largest when we equate marginal benefits and marginal costs in present value. This equimarginal condition is a powerful tool for making decisions. In many instances, the benefits of taking some action (controlling pollution, say, or providing habitat for endangered species) are increasing at a decreasing rate, while the costs rise more and more rapidly. If so, the proper response—at least if we want to maximize net benefits—is to act until the benefit of one more unit of environmental quality just equals the incremental cost. In many cases, moreover, the benefits from pursuing “perfect” policies—such as zero pollution—often do not outweigh the costs. As a result, zero pollution is typically not an efficient outcome (although the same can also be said for zero pollution control).

The discussion in this chapter has abstracted from many of the challenges in using efficiency as a guide to policy. In particular, we have assumed that the costs and benefits of environmental protection are known, and we have taken for granted that maximizing net benefits is a reasonable goal to pursue. The next chapter tackles these challenges head-on.

3

The Benefits and Costs of Environmental Protection

The previous chapter proposed a destination—economic efficiency—for our journey, but it didn't give us a road map or even a compass. Imagine you are a policymaker deciding whether to approve construction of a hydroelectric dam on a wild river. Even if you embrace the idea of maximizing net benefits to society, how can you measure the costs and benefits of the project? How can you weigh a cheap, clean source of electricity against the damage to fish populations and the loss of rapids for rafting? How should you decide whether to build the dam or let the river run wild?

A first step is to define the costs and benefits of each option. To compare these costs and benefits, you need to measure them on a common yardstick. Then you can decide which option offers the greatest net benefit, although you may still want to ask why maximizing net benefits is what you should care about in the first place.

This chapter tackles these issues. We start by considering how economists think about the costs and benefits of environmental protection and how those might be measured. Although determining costs is relatively straightforward, measuring benefits takes extra effort, as we shall see. We then consider how efficiency is implemented in practice, through benefit-cost analysis. That discussion culminates in an investigation into why efficiency might (or might not) be a desirable goal for policy.

Measuring Costs

How are costs defined and measured? In economic terms, the true costs of any activity are the *opportunity costs*—what you give up by doing one thing instead of another. For example, the true cost of going to graduate school

is not simply the tuition plus the cost of room and board but also (and crucially) the forgone income from 2 or more years out of work. More broadly, the prices of inputs such as capital, labor, and materials reflect their values in alternative uses. To produce electricity requires capital to pay for the construction of the generating unit (money that could have gone into alternative investments), labor to operate the plant (workers who could earn wages in other jobs), and fuel to produce steam (fuel that could have been used by other companies and that required expending other resources in extraction and transportation). The same principle applies to reducing pollution: Scrubbing sulfur dioxide out of flue gases requires capital to build the scrubber and labor and materials to operate it. Devoting these resources to pollution control leaves less to spend on other opportunities, such as improving the plant's operation or increasing output.

Economics offers another valuable insight into the costs of environmental protection: They are ultimately borne by individuals, whether taxpayers, shareholders, or consumers. It is tempting to think that the benefits of clean air are enjoyed by society as a whole, whereas the costs of pollution control are paid out of corporate profits. In reality, of course, the costs of pollution control—even when they are “paid for” by corporations or electric utilities—end up being borne largely by consumers of the goods and services that cause the pollution. For example, electric utilities typically recover much of the cost of pollution control by charging higher rates for electricity. Even if abatement costs also reduce the utilities' profits, much of that loss is felt by shareholders, who include retired pensioners as well as wealthy investors.

Several categories of costs are important in assessing environmental regulations: private compliance costs, government sector costs, social welfare costs, and transitional effects. What do these terms mean? Private compliance costs include most of the costs we have mentioned thus far in the book: capital costs for pollution control equipment and other infrastructure required to comply with a new regulation, changes in inputs (such as a power plant's increased costs of switching from high-sulfur to low-sulfur coal to reduce sulfur dioxide emissions or from coal to natural gas to reduce carbon dioxide emissions), and the costs of capturing regulated waste products for treatment and disposal.¹ Estimating compliance costs can be a challenge in a market economy, because most of this information is private, and firms

In economic terms, the true costs of any activity are the opportunity costs—what you give up by doing one thing instead of another.

are reluctant to share it because it directly affects competitiveness. In a pinch, analysts can use published estimates of costs for standard technologies and processes—the “engineering cost” approach.²

Outside of pollution control, the “compliance costs” analogous to abatement costs are often less obvious but no less important to consider. For example, let’s think about endangered species protection. The costs of protecting an endangered species might include money spent on preserving habitat, enforcing prohibitions on hunting or poaching, and educating landowners and the public at large. As we saw in the example of the California condor in Chapter 2, some of these costs involve public expenditures—such as increased law enforcement—not just private expenditures, as for the pollution abatement compliance costs described earlier. These kinds of government expenditures make up the second category of regulatory costs that must be included in weighing the costs of a regulation against its benefits. Relevant government costs include those for training, monitoring and reporting, permitting, and litigation.

The third category—social welfare costs—is more complicated. These costs are incurred when regulations increase the prices of goods and services in the regulated sector and beyond. If you have taken an introductory microeconomics course, you will be used to thinking about these costs as changes in consumer and producer surplus. If not, just think about the ways in which regulating emissions from power plants might affect prices. When firms face higher production costs from switching fuels or removing constituents from their waste stream, they may raise electricity prices in response.³ This price increase hurts consumers; they must now pay more for each kilowatt of electricity they purchase than they paid before the regulation, and this negative effect can be an important regulatory cost. In reaction to an electricity price increase, some consumers will purchase less electricity: they may change their behavior (turning down their thermostat in the winter, for example, and wearing a sweater indoors) or invest in energy efficiency measures that cut their electricity use. This substitution effect—spending less on electricity and more on sweaters or insulation—cushions consumers from the impact of the price change and thus dampens the social welfare costs of the regulation. That is, if we calculated the reduction in consumer surplus from the new power plant pollution regulation but failed to take into account how consumers would react to the increased price of electricity, we would overestimate the costs of the regulation.⁴ Thus, the first step in estimating a regulation’s social welfare costs is to consider the impacts of any expected price increase for the regulated good or service on consumers, accounting for consumers’

Estimating the Costs of Mitigating Greenhouse Gas Emissions

In the mid-2000s, as global interest in climate change was increasing, the prominent consulting firm McKinsey & Company carried out analyses of the costs of reducing greenhouse gas emissions, both at a global level and in specific countries. In preparing their studies, teams of consultants identified hundreds of potential abatement options across virtually every sector of the economy. For example, McKinsey's analysis for the United States considered more than 250 options, including energy-efficient lighting and heating systems in residential and in commercial buildings, fuel economy standards for various classes of vehicles, changes in industrial processes, no-till agriculture, active forest management, solar and wind power, coal-to-gas switching at power plants, and carbon capture and sequestration at fossil-fired power plants.

The consultants estimated the number of tons of emissions that could be reduced by each option, relative to an assumed "reference case" of what emissions would otherwise be in the year 2030. Drawing on assessments of current technology, assumptions about cost reductions and technological change, and projections of energy prices and other key parameters, McKinsey then estimated the cost of each option per ton of greenhouse gases abated in 2030. Finally, the options were arrayed from least to most cost to produce a marginal abatement cost curve (figure 3.1).⁵

The right-hand part of the McKinsey cost curve looks much like the marginal cost curves we discussed in Chapter 2: a range of options with increasing abatement costs. The left-hand part of the curve, on the other hand, lies *below* the horizontal axis, implying that the costs of abatement are *negative*.

A number of observers pointed to this feature of the McKinsey curve as evidence that firms and consumers could save money by reducing emissions. Many economists were skeptical. They argued that firms and consumers would not systematically be "leaving money on the table." Instead, they argued that McKinsey's analysis failed to account for all the costs. For example, consumers might not install energy-efficient fluorescent lightbulbs, even though doing so would (according to McKinsey's calculations) save them money, because they preferred the light from conventional incandescent ones. McKinsey, for its part, explained the "negative costs" by alluding to "market barriers." (Recall our discussion of the "energy efficiency paradox" in Chapter 2.)

McKinsey's analysis is perhaps the best-known example of what economists call a bottom-up cost model: To estimate the aggregate cost of reducing emissions, they identified a range of concrete options and the cost and abatement potential of each one and added them up. Top-down models take the opposite approach. Rather than specify particular abatement options, they attempt to project how the economy would respond to a policy change. Instead

Estimating the Costs of Mitigating Greenhouse Gas Emissions

continued

of relying on estimates about costs and abatement potential of different technologies, such models rely on estimates about key economic relationships. For example, data on energy prices and household energy use can be used to estimate the elasticity of energy demand, which captures how sensitive changes in energy consumption are to changes in price. A top-down model uses the elasticity of energy demand to project how changes in energy prices resulting from a policy such as a carbon tax or cap-and-trade system would translate into lower energy use and therefore lower emissions.⁶

In both types of models, the role of assumptions—about “business-as-usual” emissions in the absence of policy, future energy prices, technological change, and so on—is critical. Different models (whether top-down or bottom-up) can produce widely divergent cost estimates. For that reason, a good rule is to look at a range of estimates of the costs of greenhouse gas mitigation rather than relying on a single model.

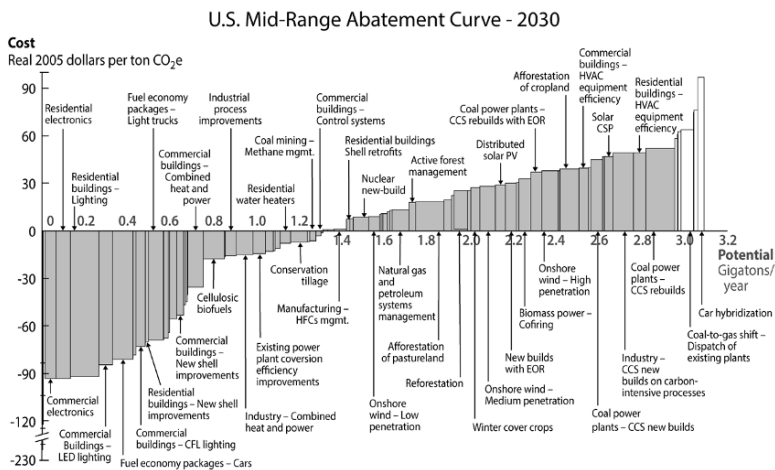


Figure 3.1 U.S. marginal abatement cost curve for greenhouse gases, developed by McKinsey & Company.

propensity to substitute away from the affected good or service. When we restrict the estimation of these social welfare costs to the market or markets directly affected by a regulation—the electricity market, in our example—we are conducting a partial equilibrium analysis.

However, social welfare costs, and some social welfare benefits, may accrue outside the market targeted by the regulation. A general equilibrium

analysis is needed to elucidate the potential spillover effects of the regulation in other sectors of the economy. Consider, again, what might happen if a new air quality regulation increases production costs at a power plant. Earlier we discussed the fact that the plant may raise electricity prices in response, causing changes in consumer behavior (and consumer and producer surplus) in the market for electricity. But higher electricity prices may also cause the owners of commercial office space to raise rents, universities to increase room and board fees for dormitories, and the manufacturers of energy-intensive consumer products (such as iPhones) to raise prices for those products. Firms that sell pollution control equipment may benefit from the regulation, because more power plants will purchase their products. A general equilibrium model captures these interactions between markets and measures a regulation's aggregate effects on consumer and producer surplus in the economy. These extra costs and benefits may be very small, not even worth estimating, for regulations with minimal effects outside the regulated sector. In other cases, such as the regulation of carbon dioxide emissions, the effects of which would ripple strongly through all kinds of product and service markets, general equilibrium models may be absolutely necessary to understanding a regulation's costs.

In our discussion about regulatory costs, you may have noticed an omission: As yet, we have said nothing about jobs. Political debates about environmental regulation are often dominated by competing views about the impact of such regulation on employment. Opponents of regulation talk of “job-killing regulation.” Supporters point to “green jobs.” How much of an impact do environmental regulations actually have on jobs?

Economic analysis suggests that the answer is “Not much.” As a starting point, it is useful to note a few general reasons why we might not expect environmental regulations to have significant effects on jobs at a national level. In general, overall employment in the economy is determined by macroeconomic factors such as investment, labor supply, and technological progress. Under normal conditions, a job “lost” in one part of the national labor market is “gained” elsewhere. Moreover, the local and regional impacts of many government interventions—from environmental regulation to highway construction—are temporary. A burst in highway construction may create jobs only for a short period of time. Jobs lost when a factory closes may be made up when another moves to town. Quite apart from the impacts of government policies, the economy as a whole is remarkably dynamic. Indeed, as much as one-fifth of jobs in the U.S. manufacturing sector are gained or lost each year.⁷

Valuing the Damages from Climate Change: The Social Cost of Carbon

Climate change is a global phenomenon caused by emissions of tens of billions of tons of carbon dioxide and other greenhouse gases each year. The potential future damages from unchecked climate change are several percentage points of global economic output, amounting to trillions of dollars annually. Although these aggregate figures offer a compelling argument for taking action on climate change, they provide little guidance in estimating or comparing the benefits achieved by particular regulations or in assessing the benefits and costs of more stringent policies.

For those questions, we need a measure of the damages from much smaller changes in emissions. Consider the impacts of 1 metric ton of carbon dioxide. (For perspective, 1 ton of CO₂ is roughly the same as the emissions per passenger on a full roundtrip flight from Washington, D.C., to Paris, or the emissions from a typical passenger vehicle in the United States over a 10-week period.) That ton will remain in the atmosphere for hundreds of years, on average, contributing ever so slightly to increased greenhouse gas concentrations and thus climate change. Consider the extra damages attributable to that ton over time: the cost of rising seas, extreme weather, drought, and wildfires; the damages to public health from heat waves, exacerbated local air pollution, and the spread of malaria; the impacts on food production; the welfare change from hotter summers and warmer winters; the value of extinct species; and so on. Finally, discount that stream of damages back to the present.

The result of that thought exercise is the *social cost of carbon* (SCC). In short, the SCC represents the present value of the marginal social damages from carbon dioxide emissions, equivalent to the marginal social benefit from emissions abatement.

As you might imagine, coming up with an estimate of the SCC is a challenge.⁸ You won't be surprised to learn that the choice of discount rate is crucial. As we saw in Chapter 2, seemingly small differences in discount rates can cause the present value of distant costs and benefits to vary by two or three orders of magnitude.

Other challenges lurk in the models used to estimate damages (called integrated assessment models [IAMs] because they integrate models of the climate system with models of the economy to assess the potential damages from climate change under various emission scenarios). A key issue is specifying the "damage function" that relates the rise in global average temperatures to aggregate damages. For small temperature increases, on the order of 1°C to 3°C, researchers may be able to come up with reasonable estimates (based on observations of how actual variation in temperatures affects crop yields, say, or

Valuing the Damages from Climate Change *continued*

human health). Because unchecked climate change could raise temperatures well beyond that point, climate modelers must make strong assumptions about the damages from even higher temperatures. Similarly, it is difficult to incorporate the impact of unlikely but potentially catastrophic events, such as the melting of the Greenland ice sheet or the reversal of the North Atlantic ocean currents that keep northern Europe habitable. Nonetheless, despite their flaws, the comprehensive approach taken by IAMs makes them a valuable tool to inform our thinking about the economic impacts of climate change.

In 2009, the Obama administration set out to establish a common value for the SCC to use in analyzing the impact of proposed regulations. (Up to that point, different agencies used different values for the SCC or didn't use one at all; indeed, no U.S. government agency included the SCC in regulatory impact analysis before 2008.) Meeting over the course of several months, a group of analysts from various White House offices, the Environmental Protection Agency, the Department of Energy, and other government agencies agreed on a set of assumptions (such as future emissions and economic growth scenarios, values of key parameters relating emissions to temperature increases, and discount rates). These assumptions were used as inputs into the three most commonly cited IAMs, which were run hundreds of thousands of times, resulting in different distributions of SCC estimates. The results were published in 2010 and then updated (using the same assumptions but revised versions of the three IAMs) in 2013.⁹

Figure 3.2 illustrates the results for the year 2020. (The SCC rises over time in real terms, because future emissions will impose greater costs on the margin; recall from our discussion in Chapter 2 that the damages from pollution typically rise more and more rapidly as pollution increases.) Results are shown separately for each of the three discount rates considered: 2.5, 3, and 5 percent. The results are shown as *distributions* of the SCC, because different model runs rely on different assumptions about parameters. (For example, a model run that assumed rapid emission growth and a high climate sensitivity would yield greater future damages and thus a higher SCC than a model run with slow emission growth and lower climate sensitivity.) The height of each bar is the fraction of model runs that yielded a particular value for the SCC. For example, using a 3 percent discount rate, the estimated SCC was around \$20 per ton in just over 10 percent of the model runs. The figure also highlights the average estimates across all model runs, for each discount rate.

These distributions, especially at the lower discount rates, have a distinctive feature: They have a long right-hand "tail." Intuitively, there is a small but nonzero probability that damages from CO₂ emissions are much higher than

Valuing the Damages from Climate Change *continued*

the average value. Given a 3 percent discount rate, for example, the *average* SCC estimate was \$43 in 2007 dollars, but roughly 10 percent of the model runs yielded an SCC of \$60 or greater (corresponding to the sum of the areas of the bars at or above \$60). One way of capturing this “tail effect” is to calculate the 95th percentile value of a distribution, the value that is greater than 95 percent of the model runs for a given discount rate. As the figure shows, the 95th percentile SCC for a 3 percent discount rate is \$129/ton.

For the purposes of regulatory analysis, the working group recommended using the four estimates highlighted in the figure, corresponding to the three average values plus the 95th percentile under a 3 percent discount rate. These values are now used by agencies across the U.S. government in assessing the benefits and costs of policies that are expected to reduce CO₂ emissions, including fuel economy standards for cars and trucks, energy efficiency standards for appliances, and emission standards for power plants.

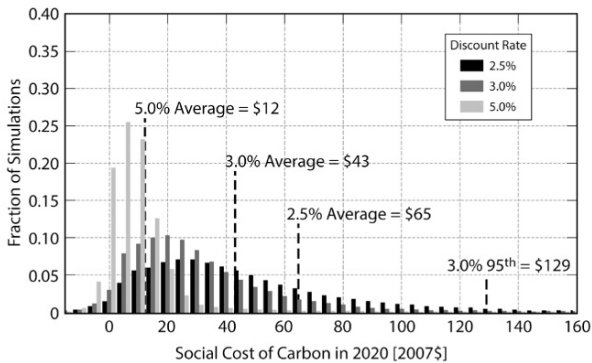


Figure 3.2 Distributions of the U.S. government’s estimated “social cost of carbon” — the monetized damages from a ton of CO₂ emitted in 2020 — under three different discount rates.

Source: United States Government Interagency Working Group on the Social Cost of Carbon, “Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866” (May 2013).

If job impacts at the level of the economy as a whole are small, what about those in regulated industries? From a theoretical point of view, the effect of environmental regulation on employment in the regulated industry could be either positive or negative, because competing effects push in different ways. Some of the money firms spend to comply on regulation will be spent on labor (a positive effect on jobs); some of the

cost of compliance will be passed on to the firm's consumers in the form of higher prices, reducing demand for output (a negative effect); and regulation may lead firms to use more or less labor relative to other inputs such as energy and capital (an ambiguous effect).

Meanwhile, empirical studies have generally failed to find a significant effect. For example, pollution expenditures in four U.S. manufacturing industries are not associated with significant long-run employment changes, and there is no evidence of long-run adverse effects of environmental regulation on employment in the United Kingdom. Similarly, environmental regulations do not appear to be a major driver of decisions about where to locate manufacturing plants; we discuss this issue further in Chapter 11, when we discuss interactions between trade and environmental regulation.¹⁰

Of course, even job impacts that are transitory and limited from a national perspective can be keenly felt by affected workers. One analysis estimated that workers in manufacturing plants newly regulated under the 1990 Clean Air Act Amendments experienced more than \$5 billion in forgone earnings in subsequent years, because of temporary unemployment and lower wages in later jobs.¹¹ These costs were temporary and very small relative to the human health and other benefits of the amendments. But such transitional employment impacts can play an outsized role in national debates over environmental regulation.

This example suggests a broader conclusion. The impacts of environmental regulation on employment (whether positive or negative) should indeed be considered in economic analyses of environmental regulations. But on the basis of the available evidence, those impacts are likely to be transitory and small relative to the other costs and benefits of regulation.¹²

Evaluating the Benefits

Before we think about how to measure benefits, we must first think about how to define them. Let's start with an example. Suppose a parcel of open space near where you live—a wetland or a woodland or a beachfront site—is up for sale and likely to be converted into a housing development. How would you think about your own value for preserving the open space? It might seem self-evident that the value from open space is higher than the value from development. But of course that might not be so obvious to the developer, or the people who would like to buy houses in the new development, or other residents of the town who care less about the open space than the influx of new residents and the addition to the tax base. Indeed, your own house might once have been part of

a new development on what used to be open space. In order to balance the views of people who favor preserving the open space against those in favor of development, we need to have some way of thinking about the benefits from preservation.

How, then, should we think about benefits? From an economic point of view, a person's value for a particular good can be determined by what he or she would willingly give up in exchange. Economists call this measure of benefits *willingness to pay*. It captures the basic truth that in a world of limited resources, we always have to give up a good thing in order to get another good thing. (We'll come back to this point again when we discuss the rationale behind benefit–cost analysis.) Put differently, if you are not willing to give up anything to preserve the open space, then we might reasonably conclude that your value for it is zero.

Although this is a straightforward way of defining *value*, several important points should be made. First, by someone's "willingness to pay" we mean the amount that the person would just barely be willing to pay. If my willingness to pay to preserve a parcel of open space is \$100, it means I would be just as happy if I paid \$100 and the open space was preserved as if I paid nothing and the open space was developed. In this sense, willingness to pay really measures the *maximum* a person would be willing to pay.

Second, nothing in this concept requires that any payment actually be made. My willingness to pay exists as a measure of value, independently of whether I actually write a check for \$100 or not. We are concerned here with willingness to pay "in one's heart of hearts," so to speak.

Third, this concept of value is unabashedly anthropocentric, or human centered. To economists, the value of anything depends on the satisfaction (or "happiness" or "utility") that humans derive from it. Note that this is not the same thing as saying that only "useful" things have value. On the contrary, the economic notion of willingness to pay is as expansive as people's imaginations. A value for open space or biodiversity or endangered species is perfectly consistent with it, because humans have shown (through their actions) that they are willing to sacrifice much to preserve them.

Fourth, in addition to being human centered, the economic notion of value centers on the preferences of the individual rather than of the group. Economic analysis emphasizes the values and preferences of individuals *as they perceive them*. One might

A person's value for a particular good can be determined by what he or she would willingly give up in exchange.

Economists call this measure of benefits willingness to pay.

say that value is in the eye of the beholder. Economic theory is rigorously neutral with respect to what people decide is valuable to them. Most people tend to think of one set of preferences as “right” or “superior” because it seems obvious to them. For example, it might be tempting to think that everyone ought to conserve water or gasoline as much as possible. But from the viewpoint of economics, we must consider the preferences of those who would rather keep an enormous lawn green or drive a gas guzzler. This does not mean that people who prefer to drive Hummers rather than Priuses should not bear the consequences of the extra carbon dioxide, nitrous oxide, and particulate matter they emit. To the contrary, as we shall see in Chapter 5, economics prescribes that people and firms pay the full costs to society for their actions. However, the individualistic perspective of economics does hold that people should ultimately be allowed to make their own decisions about what they value and what they do not.

A drawback of this approach is that it does not necessarily take account of how well informed people are about the environment. For example, if people do not understand the ecosystem services that wetlands provide (e.g., flood control, habitat, and water quality), they may place less value on those resources than they would if they knew more. A natural solution is to educate people about the benefits of environmental protection. For example, when economists conducted surveys in the aftermath of the 1989 *Exxon Valdez* oil spill in Prince William Sound, they took care to provide some background information on the spill and its likely environmental effects before asking people to estimate the damages from the spill (we’ll discuss this particular valuation exercise later in the chapter).

Fifth, a subtle but potentially crucial distinction arises between willingness to pay (WTP) and *willingness to accept* (WTA). The difference between the two hinges on the implicit assignment of property rights. If I ask you what you are willing to pay to preserve open space from development, I implicitly assign a property right to the developer. If, instead, I ask what you are willing to *accept* from the developer in compensation for developing the very same parcel of land, I have given you the implicit property right. In practice, these implicit property rights are often determined by the status quo. For example, if we contemplate a prospective action that has not yet taken place (“Should we as a society allow oil drilling in the Arctic National Wildlife Refuge?”), then we might naturally consider society’s willingness to accept compensation for loss of the wilderness area. On the other hand, in the aftermath of the 1989 *Exxon Valdez* accident or the 2010 Deepwater Horizon oil spill in the Gulf of Mexico,

the relevant issue for valuation was society's willingness to pay to clean up the spill. A related way of stating the distinction is that WTA is a natural measure for your losses, whereas WTP corresponds more readily to gains.

In many cases, WTP and WTA should coincide, at least if we can measure them accurately (i.e., if we can set aside reasons why people might intentionally understate their WTP or overstate their WTA). For example, suppose you head to Fenway Park in Boston to see the Red Sox play a baseball game. Let's assume you would be willing to pay up to (but no more than) \$50 for a grandstand ticket. If you show up on the day of the game without a ticket and find someone willing to sell you one for \$50 or less, you will buy it. On the other hand, if you already have a ticket, and someone offers you \$50 or more for it, you will sell. If you were not willing to pay \$80 to buy a ticket, you will not be willing to forgo \$80 if you can sell it for that much.

As the stakes get larger and larger, however, a gap may open up between what you would be willing to pay and what you would be willing to accept. This is simply because willingness to pay depends in part on *ability to pay*. To see this, consider the Great Barrier Reef in Australia, which is threatened by the rising ocean temperatures associated with global climate change. How much you would be willing to pay to prevent the reef from being destroyed? Although the amount might be large, it is constrained by your need to have money to buy other goods and ultimately by your wealth. Now consider how much you would be willing to accept to allow the reef to be destroyed. Your answer now might well be "infinity," or equivalently "no amount of money." In that case WTA far exceeds WTP. But it helps illustrate an important point. WTA will always be at least as large as WTP. The two measures may diverge considerably, however, when large income effects are involved. Keeping this relationship between WTP and WTA in mind, we will refer to WTP rather than WTA when we consider the benefits of environmental protection.

Finally, the use of WTP as a measure of value is not confined to environmental goods: On the contrary, the concept is fundamental to economic theory. The same principle applies to shoes or coffee or automobiles as to open space: The benefits you derive from a particular good can be measured by what you would be willing to give up for it. Of course, in a market setting what you would be willing to give up for something can be inferred from your purchase decisions. If you pass up a pair of shoes when they cost \$100 but buy them on sale for \$80, then I can infer that the value you place on the shoes is between \$80 and \$100. In contrast, no such observable data exist for most environmental amenities. Indeed,

as we shall see later, the main practical hurdle to valuing environmental goods is that clean air, endangered species, open space preservation, and other environmental amenities are not traded in markets.

A Taxonomy of Values

Defining the benefits of environmental quality in terms of willingness to pay still leaves us with a range of kinds of values—that is, reasons people might be willing to pay. Let's return to the example of open space preservation. Some people would be willing to pay for open space because they are eager to visit it—to take walks in the woods or look for birds in a wetland. Other people without current plans to visit the area might nonetheless be willing to pay something to preserve it now, in order to have the option to use it in the future. Still others might not plan on using it themselves but would want to preserve it to pass on to their children or grandchildren. Finally, some people might be willing to pay something to preserve a parcel of open space simply because they take pleasure in knowing that it exists.

As the example makes clear, we can identify several distinct types of value. Today's visitors value open space because they will use it for recreation, and prospective visitors want to preserve the option of using it in the future. Both of those values fall under the heading of *use value*. Use values involve direct enjoyment or consumption of an environmental good. For another example, consider the gains from reducing smog in Los Angeles. The benefits range from attenuated adverse health effects (eye irritation, asthma, difficulties breathing) to aesthetic values (better views from homes high in the canyons or hiking trails in the mountains nearby). Both the health effects and the aesthetic values constitute use values, because they represent direct consumption or enjoyment of the cleaner air. Use values also arise from recreation—as for water quality in trout streams or at the beach or water volume in river rapids enjoyed by rafters.

The other two types of value mentioned earlier—the desire to preserve a resource for future generations and the pleasure taken from the knowledge that something exists—are, naturally enough, called *nonuse values*. They involve benefits derived from the existence of an environmental amenity *but not from its direct use*. Existence value is of particular importance for endangered species preservation in the real world: How many donors to wildlife conservation societies actually believe they will see a Bengal tiger or a polar bear? Rather, much of the value of endangered species or habitats is simply knowing that they are there.

Measuring Benefits

To apply these concepts, of course—and to carry out benefit–cost analysis—we must be able to measure how much people are actually willing to pay for a given environmental amenity. An in-depth discussion of valuation methods is beyond the scope of this book. However, we can sketch out the basic intuition behind the major approaches.

For most consumer goods, measuring benefits is straightforward. To determine how much consumers are willing to pay for dress shirts, or iPhones, or home appliances, analysts can gather market data on the prices of goods and the quantities purchased. In the environmental realm, this is possible for some goods and services. For example, the impacts of water pollution abatement on the local commercial fish catch can be valued using market prices and quantities, as can the impacts of reducing acid rain on commercial timber harvests. This conceptually simple approach does not work for most aspects of environmental quality, however, because most environmental goods and services are not traded in markets.

Economists have developed two basic strategies to circumvent this lack of price and quantity data in order to estimate the value of environmental amenities. The first approach is to observe behavior indirectly, in related markets, and use that information to infer willingness to pay for environmental quality. Economists call this the *revealed preference* approach, because it treats actual behavior as revealing the true underlying preferences of individuals. The second basic approach to estimating value is simply to ask people how much they would be willing to pay to protect a given environmental resource (e.g., an area of habitat, a population of endangered species). This is the *stated preference* approach. We describe several revealed preference methods, and one stated preference method, in the next two sections. For reasons we describe later, economists tend to favor revealed preference methods, although for some values stated preference methods are the only alternative.

Revealed Preference Methods for Estimating the Value of Environmental Quality

One revealed preference method, the *travel cost method*, infers individual marginal willingness to pay for environmental quality from decisions about where to travel for recreation.¹³ Although many national parks and other major natural resource recreation sites have an entrance fee, that fee is often a very small component of the total cost of a visit; you

would spend much more on an airline ticket or a road trip to Yellowstone National Park than you would to gain entry to the park itself. The travel cost method statistically exploits the relationship between travel cost and visitation rates, permitting the estimation of a demand function for recreation. If the economic value of recreation at a site depends on site characteristics, then multiple-site models can be used to value those characteristics. For example, if we observe an angler traveling to a remote fishing lake, we can infer that he values the experience of fishing there at least as much as the trip's cost (in time as well as money). By comparing the frequency and duration of anglers' visits to pristine lakes with those to more polluted ones, one can estimate the marginal benefits of clean water to recreational fishers. Travel cost models are among the most widely applied valuation methods and have become a very useful tool for estimating recreational demand.

Another revealed preference method uses observed market prices to infer the *implicit* prices for environmental amenities that are bundled with other (private) goods. A common application is to housing markets. Consider the following thought experiment: Imagine two houses located in the same metropolitan area that are identical in every respect, except that one of the houses is in a neighborhood with poorer air quality. The house with cleaner air will command a higher price, with the size of the premium being determined by the difference in air quality and by people's marginal willingness to pay for clean air. In effect, the housing market does set a price on, or capitalize the value of, clean air.

Of course, in the real world houses are never identical on every dimension except for air quality: One is located on a bigger lot and in a better school system, and the other house has a shorter commute and is close to a city park. Instead of being directly observable, the price of clean air in the housing market must be teased out from other attributes. Dealing with these other factors, however, is a methodological rather than a conceptual challenge, and sophisticated statistical techniques have been developed to isolate the effect of air quality from other (potentially correlated) factors. Thus the *hedonic property method*, as it is called, has been used widely to estimate marginal willingness to pay for reductions in a variety of air pollutants and for other local environmental amenities and disamenities, such as toxic waste sites. For example, a recent study demonstrates that proximity to natural gas wells in Pennsylvania's Marcellus Shale region increases home values, perhaps because of actual or expected benefits from mineral royalty payments. But homes that obtain their drinking water from their own private groundwater wells, rather than municipal

piped water, experience decreases in property values of 10 to 22 percent when near shale gas wells, suggesting that local housing markets capitalize the potential risks to groundwater associated with energy development.¹⁴

A similar method can be applied to labor markets. Observed wages vary with the amount of education, training, and expertise needed to perform a job well, the cost of living and desirability of the job location, and other factors. One of those other factors is a job's embodied risks to life and health. Economists apply the *hedonic wage method* to estimate the wage premium workers receive for accepting a high-risk job. The estimated values of these small changes in risks to life and health in job settings are used frequently in benefit–cost analysis of environmental regulations to monetize the value of avoided premature deaths and avoided illness from exposure to pollution. For example, one might observe that workers appear to be willing to give up \$14 in annual wages to reduce the risk of dying on the job from 1 in 400,000 to 1 in 500,000 (a reduction in the probability of death of 0.000002). These marginal valuations can be used directly to estimate the benefits of an air pollution reduction resulting in similarly small changes in risk for the exposed population as a whole. In many cases, regulatory agencies go one step further, linearly extrapolating from these values the value of a single (statistically) avoided premature death from pollution exposure. This value has come to be known as the value of a statistical life (VSL). For example, if we make the following calculation from the change in occupational risk discussed earlier, the marginal willingness to pay for risk reduction divided by the marginal change in risk is $\$14/0.000002$. If we set this fraction equal to the marginal willingness to pay for a change in the risk of dying from 1 to 0, then $\$14/0.000002 = \$X/1.0$. The VSL is the numerator of the right-side fraction ($\$X$), $\$14/0.000002$, or \$7 million.

This (\$7 million) is, in fact, approximately the VSL currently used by the Environmental Protection Agency (EPA) in benefit–cost analyses of pollution control regulations that reduce the risk of premature death from pollution exposure; it is the average value obtained from twenty-one different hedonic wage studies and five stated preference studies (more on these later) that estimate the value of risk reductions. A VSL of \$7 million does not imply that an individual would be willing to pay \$7 million to avoid certain death this year. It does imply that those exposed to a 1 in 400,000 risk of dying would pay about \$14 to reduce that risk to 1 in 500,000, or for that matter, that those exposed to a 1 in 10,000 risk of death would pay about \$700 to eliminate that risk (because $\$700/0.0001 = \7 million). It is not the “value of a life” in either economic or ethical

terms, although it is often a lightning rod for controversy among those who misunderstand the concept or simply oppose its use in the design of environmental policy.

Assigning monetary value to avoided death and illness makes some people uncomfortable. A key thing to take away from this discussion, however, is that spending public and private funds on pollution control regulation, and risk mitigation more broadly, involves tradeoffs. And these tradeoffs are either made implicitly, and can then be quantified *ex post*, or made explicitly *ex ante* using benefit–cost analysis. Avoiding monetization of risks to life and health does not erase the tradeoffs that are made when standards are set.

The travel cost method and hedonic property and wage methods of environmental benefit valuation are all revealed preference methods in that they estimate willingness to pay from actual market transactions. Even though there are not markets for clean air, peoples’ willingness to pay for clean air leaves footprints in housing markets and labor markets; we can exploit this information to estimate the implicit value of nonmarket environmental goods. Similarly, even though park entrance fees do not capture visitors’ willingness to pay for outdoor recreation, visitors’ travel expenditures give us some information to estimate these values. However, some of the values discussed in the introduction to this section—the value of simply knowing that the Arctic National Wildlife Refuge remains pristine or knowing that the giant panda and the polar bear exist—leave no footprints in markets. To estimate these kinds of values, stated preference methods are needed.

Stated Preference Methods for Estimating the Value of Environmental Quality

The *stated preference* approach consists of several methods, although the most common method is *contingent valuation* (CV). CV is conceptually straightforward: Carefully structured surveys are administered to a group of people to obtain information on their willingness to pay. The great advantage of this approach is its broad applicability. Revealed preference methods, by their very nature, can be applied only to estimate use values. Existence value is purely internal rather than behavioral: It leaves behind no “paper trail” of observed behavior that can be used to infer someone’s value for a resource. In contrast, CV can in principle be used to measure willingness to pay for any good, simply by asking people.

Of course, this broad applicability does not come cheap. The disadvantage of CV methods is that the respondents lack strong incentives to tell the truth. If they feel they have little stake in the outcome—thinking, for

example, “this is just a survey”—then they may have little reason to think carefully about their choice and provide a thoughtful response. Things may be even worse if respondents believe that their answers to the survey will affect policy, because in that case they may have strategic incentives to misrepresent their true valuations. For example, suppose I have a moderate willingness to pay for clean air, and I live in an area with a large population. If I think that I will be asked to pay a tax on gasoline or electricity (to finance pollution control) based on how I respond to a CV survey, I will have an incentive to understate my true valuation, in the hopes of “free riding” off of the contributions of others. After all, I might reason, my own tax payment is likely to have very little effect on air pollution, but I would be much better off personally if that tax payment were as small as possible. On the other hand, if I believe that my preferences for clean air will help secure strong regulation but that I will not be asked to pay on the basis of my response, then I may well overstate my willingness to pay, in the knowledge that the cost of clean air will be borne largely by others.¹⁵

Economists who specialize in CV have devised a number of approaches to mitigate these and other potential sources of bias. As a result, a range of meta-analyses that have compiled results from multiple studies have concluded that CV methods and revealed preference approaches yield similar estimates of willingness to pay for environmental amenities whose value can be assessed using either method. This is taken as evidence in support of these methods. Of course, because revealed preference methods cannot be used to estimate nonuse or existence value, CV estimates of these values cannot be similarly calibrated. In addition, estimates produced by CV are generally much more variable across studies. Economists tend to put greater trust in the revealed preference approaches, just as you might wisely attach more weight to what another person actually does rather than what he or she says. Nonetheless, CV remains invaluable for its ability to provide estimates of existence value. The method has been used in high-profile cases of natural resource damage assessment, to estimate lost passive use (or nonuse) values from oil spills and other environmental harms.

Both the revealed preference and stated preference approaches to estimating the value of environmental amenities tend to focus on marginal willingness to pay. Studies that infer value from behavior (such as the hedonic property studies) necessarily focus on marginal willingness to pay, because that is what can be observed directly. The price of clean air is essentially a measure of what people are willing to pay for an incremental improvement in air quality, just as the price of an iPhone is the price of

Contingent Valuation and Natural Resource Damage Assessment

In March 1989, the oil tanker *Exxon Valdez* ran aground on Bligh Reef in Prince William Sound off the coast of Alaska with 50 million gallons of crude oil in the hold. To that date, it was the largest tanker spill in U.S. history; 11 million gallons spilled in less than 5 hours. By August, the oil slick from the spill covered 10,000 square miles in Prince William Sound and the Gulf of Alaska. Known environmental impacts included the deaths of hundreds of thousands of seabirds, thousands of sea otters, hundreds of bald eagles, and many orcas. Following a 1989 U.S. federal court opinion allowing compensation for lost “passive use” (or nonuse) values under U.S. environmental statutes, the Oil Pollution Act of 1990 stated that these values could be included in natural resource damage assessment, the practice used to support litigation on behalf of entities seeking compensation from those found responsible for environmental harms.

When the State of Alaska took Exxon to court over the incident, it hired a team of economists to estimate the nonuse damages from the spill using contingent valuation. Nonuse value was potentially important, because many Americans may have experienced damages from the spill, knowing that a pristine area that was home to many charismatic species had been harmed. Exxon hired its own team of economists to highlight the weaknesses of CV methods and thus cast doubt on Alaska’s damage estimates.¹⁶ The team working on behalf of Alaska estimated total nonuse damages (to all U.S. citizens) from the spill of about \$2.8 billion. This estimate was over and above any losses in use value, such as damages to commercial and recreational fishing and tourism (and was, in fact, many times larger than these estimated losses). After lengthy litigation, Exxon agreed to pay about \$1 billion in damages; it also spent \$2 billion of its own funds on response and restoration. Some have pointed out that the magnitude of estimated damages is approximately equal to the sum of Exxon’s settlement and restoration expenditures.

Although court cases since the Exxon incident have continued to reinforce the validity of CV in litigation over natural resource damages under the Oil Pollution Act, CERCLA (or Superfund), and the Clean Water Act, it has been used infrequently in this context. One additional application of CV supported claims by the federal government of long-term damages to birds and fish caused by exposure to polychlorinated biphenyls and DDT, deposited off the southern California coast for many decades by seven firms. After the Deepwater Horizon accident in the Gulf of Mexico in 2010, the largest marine oil spill in U.S. history (involving almost twenty times the amount of oil spilled by the *Exxon Valdez*), no comprehensive CV analysis was done to estimate nonuse damages. The incident did prompt renewed debate about the validity of CV, however.¹⁷ In the 21 years between the Exxon and Deepwater Horizon spills, thousands of academic articles and books were published on this method of estimating nonuse values, but the role of CV in litigation and public policy decisions remains controversial.

a single iPhone. CV studies also tend to focus on marginal willingness to pay, in large part because policy decisions tend to be made on the margin. That is, the relevant question for pollution control policy is not “Should we reduce air pollution?” but rather “Should we reduce pollution by 8 million tons a year or by 10 million tons?” Similarly, debates over endangered species laws focus on the level of protection that should be provided rather than on whether endangered species should be protected at all.

Benefit–Cost Analysis

Thus far, we have discussed efficiency as a “best” or “maximal” outcome— for example, a particular level of pollution control that maximizes net benefits. The concept of efficiency is also useful in comparing alternatives, neither of which necessarily maximizes net benefits relative to all other possible policies. The basic principle is straightforward: Policy A is more efficient than policy B if the net benefits are greater under policy A. For example, Congress may consider whether to require electric power plants to install pollution control devices to reduce mercury emissions into the atmosphere, comparing a proposed regulation with the status quo. A regional development agency may decide whether to spend transportation funds on highway expansion or on construction of a light-rail network. In cases such as these, a systematic comparison of benefits and costs can be a useful aid to government policy.

In the United States, the major environmental statutes differ in their relative emphasis on benefit–cost analysis as a policymaking tool. As economist and former EPA official Richard Morgenstern has noted, the statutes alternately “forbid, inhibit, tolerate, allow, invite, or require the use of economic analysis in decisionmaking.”¹⁸ Since 1996, the Safe Drinking Water Act has contained the most significant requirement for benefit–cost analysis of any environmental statute; it is required for all new maximum contaminant level (MCL) standards, and revisions to any existing MCLs, for treated drinking water. At the other end of the spectrum, the Clean Air Act forbids the consideration of cost in setting the national ambient air quality standards that govern local and regional air pollutants. But some of the most significant benefit–cost analyses that have been performed for U.S. environmental policy have focused on Clean Air Act regulations; we discuss some of them in this chapter and others later in the book. If the Clean Air Act forbids consideration of costs in standard setting, why has the EPA performed these economic analyses? The answer is that benefit–cost analysis of all environmental, health, and safety regulations expected to have significant costs to the U.S. economy has been required by executive orders signed by every U.S. president (from both political

What Is the Value of Global Ecosystem Services? A Cautionary Tale

In a well-known article published in *Nature* magazine, Robert Costanza and a group of colleagues set about estimating the total worldwide value of renewable ecosystem services.¹⁹ They focused on seventeen types of ecosystem services, including pollination, nutrient cycling, and regulation of the composition of gases in the atmosphere, as well as more mundane goods and services produced by nature, such as food production, raw materials (e.g., timber), and recreational opportunities. Their goal was to demonstrate the economic significance of such natural capital, most of which is not traded in the economy.

Their conclusion? The study estimated the annual value of such ecosystem services worldwide to be \$33 trillion. That might sound like a lot of money. After all, as the authors point out, it is just shy of twice the global gross national product at the time of the study. On reflection, however, this amount starts to look fairly small. After all, the world economy grew at about 2 percent per year in 1995. At that rate, the world economy would surpass the value of ecosystem services by the year 2025 or so. The elimination of major global ecosystem services would obviously have devastating effects on human welfare. Indeed, as one analyst quipped, \$33 trillion is “a serious underestimate of infinity.”²⁰ Where did the analysis go wrong?

An important part of the answer is that they use estimates of marginal willingness to pay in computing the total values of ecosystem services. Of course, these concepts are closely related, as we saw in Chapter 2. But we also saw that marginal willingness to pay often depends on the level of environmental quality. For example, to answer the question “What is the value of pollination services?” Costanza and his colleagues figured out the value of pollination for the marginal hectare and then multiplied it by the total land area that is pollinated. This is akin to the difference between asking “What is the value of a particular bee colony?” and asking “What is the value of all of the bees in the world?” We can’t answer the latter question by scaling up the value of a single beehive. As beehives became scarcer, they would become more valuable.

A related point is that the complete elimination of vital ecosystem services would affect society in fundamental and interconnected ways. When economists measure marginal benefits, they evaluate small changes. The loss of all wild pollinators worldwide—or even more dramatically the shutdown of processes that regulate Earth’s atmosphere—would cause drastic shifts in the demand and supply of all kinds of goods and services (including the very amenities and services being measured).

Does this mean that the benefit techniques we have discussed are not useful? On the contrary, they are eminently suited to assessing the effect of specific policies, such as reducing concentrations of air pollution or setting aside a wilderness area for habitat protection. They are simply inadequate measures of the value of global ecosystem services in their entirety.

parties) since Richard Nixon. Since the Reagan administration, the level of expected cost that has triggered this requirement (for what has become known as Regulatory Impact Analysis) is \$100 million per year. Neither the benefit–cost analysis requirement under the Safe Drinking Water Act nor the executive order requirement covering all environmental policy involves strict benefit–cost tests. That is, EPA and other agencies may still propose and enact rules for which the monetized costs exceed the monetized benefits. An important and controversial example of this was a revision to the rule governing allowable concentrations of arsenic in U.S. drinking water, proposed by EPA in 2000 and enforced beginning in 2006. Nonetheless, the benefit–cost analyses of U.S. environmental regulations produced because of these requirements have provided very important information to the environmental policy process.²¹

Critiques of Benefit–Cost Analysis

Although it is simple to describe, benefit–cost analysis has attracted considerable controversy, especially in the environmental arena. Critics of benefit–cost analysis typically advance four main arguments against it.²² First, basing decisions simply on whether benefits outweigh costs omits important political and moral considerations, such as fundamental rights or duties. Second, discounting benefits that will occur in the distant future privileges current generations at the expense of future ones. Third, goods such as clean air or clean water are devalued and cheapened when their worth is expressed in monetary terms. Finally, focusing on the net benefits to society as a whole ignores the identities of the winners and the losers—that is, an emphasis on efficiency obscures a consideration of distributional equity. We consider each of these criticisms in turn. Doing so allows us to probe more deeply the usefulness of economic efficiency and the limitations of economic analysis.

“Benefit–Cost Analysis Should Not Be the Only Criterion for Decision Making”

Critics of benefit–cost analysis often contend that economic theory prescribes the use of benefit–cost analysis to the exclusion of other considerations. The emphasis that economists place on the concept of efficiency may initially leave the mistaken impression that they view net benefits as the only criterion needed for public policy. In fact, most economists reject such a narrow view. The consensus among economists is that benefit–cost analysis should be viewed as a means of improving the information available to decision makers, not as the sole guide to decision making. In the

Benefit–Cost Analysis in the Real World: EPA’s Study of Lead in Gasoline

In the fall of 1983, the U.S. Environmental Protection Agency initiated an internal study into the benefits and costs of reducing lead in gasoline.²³ The use of lead as a fuel additive to increase octane levels had been restricted since the early 1970s, but high levels of lead persisted in the environment in the early 1980s, partly because older leaded-fueled cars remained on the road and partly because some owners of newer cars “misfueled” their cars by using leaded gasoline.

The agency’s Regulatory Impact Analysis (RIA), published in 1985, identified four main benefits of phasing out lead. Two were linked to reductions in the main adverse human health effects from lead: damages to cognitive and physiological development in children and exacerbation of high blood pressure in men. A third benefit was a reduction in emissions of other pollutants from cars, because burning leaded gasoline destroyed the effectiveness of catalytic converters. The final benefit was lower costs of engine maintenance and related increases in fuel economy. Meanwhile, the primary costs of phasing out leaded gasoline were the costs of installing new equipment at refineries and producing alternative additives to boost octane levels.

The study found that the benefits of reducing lead would substantially outweigh the costs. For a tenfold reduction in lead content (from 1.1 to 0.1 grams per gallon of gasoline), the net benefits were a little over \$7 billion dollars annually (in 1983 dollars): \$7.8 billion in benefits minus \$600 million in increased refining costs. These findings helped support the EPA’s decision to accelerate the required removal of lead from gasoline. That reversed the trend set a few years earlier, when the agency, citing costs to refineries, had settled on a much weaker rule than preferred by environmental and public health advocates.

That the benefit–cost analysis succeeded in bolstering the case for tighter regulations was all the more notable because of the gaps in the analysis, gaps readily acknowledged by the study’s authors. EPA staff were able to include only the avoided costs of medical care and remedial education for children with blood lead levels above a certain threshold, leaving out the lion’s share of benefits, such as the willingness to pay to avoid lasting health and cognitive impacts.

The leaded gasoline example illustrates the usefulness of computing benefits in dollar terms to make them commensurable with other benefits and with costs. Indeed, the case for removing lead from gasoline would have been even stronger had the analysts been able to estimate the dollar value of improving children’s health and cognitive abilities. Quantifying benefits (as well as costs) can focus regulatory efforts on the areas likely to yield the greatest net benefits to society. As one of the study’s authors has pointed out, there was little political pressure on EPA at the time to tighten lead standards. Indeed, much more attention was being devoted to issues such as hazardous air pollutants and uranium mill tailings, even though the damages from such problems (and thus the benefits from cleanup) were orders of magnitude smaller than those at stake with leaded gasoline.

words of a blue-ribbon panel of economists led by Nobel laureate Kenneth Arrow,

Although formal benefit–cost analysis should not be viewed as either necessary or sufficient for designing sensible public policy, it can provide an exceptionally useful framework for consistently organizing disparate information, and in this way, it can greatly improve the process and, hence, the outcome of policy analysis.²⁴

“Discounting Is Unfair to Future Generations”

A second criticism concerns the use of discounting to compare costs and benefits over time. It may seem grossly unfair to weigh costs to people living today more heavily than benefits to future generations. Indeed, economists recognize full well the thorny issues of intergenerational equity raised by discounting. William Nordhaus has illustrated the dilemma with a particularly striking example. Suppose we discover that Florida will be entirely destroyed 200 years from now by an asteroid impact. Suppose that we could prevent this catastrophe by launching a missile today to intercept the asteroid. How much should we be willing to pay for the missile? Using the value of land and capital in Florida and the 7 percent discount rate mandated by the U.S. Office of Management and Budget for use in government benefit–cost analyses, Nordhaus calculates that preventing Florida’s annihilation two centuries hence would be worth only about \$3 million. If launching the missile cost more than that, a strict benefit–cost test (using the government discount rate) would advise against its launch.²⁵

This thought experiment highlights the importance of the choice of discount rate. After all, if Nordhaus applied a 2 percent discount rate—putting more weight on the future—the present value of preventing the destruction of Florida would be around \$43 billion—still small, perhaps, but several orders of magnitude more than \$3 million.

At a deeper level, however, this thought experiment illustrates the limitations of using efficiency as the only criterion for decision making. Sometimes, benefit–cost analysis may suggest a course of action that we might still choose to reject on the basis of ethical considerations—for example, concerns about intergenerational equity. But that does not mean that we should not do the analysis in the first place. In making decisions about how to trade off current costs and future benefits, we as a society are surely better off when we have more information about the choices we face. Carrying out a benefit–cost analysis does not commit us to any

particular course of action; it simply helps us clarify the stakes. This argument echoes the one made by the blue-ribbon panel cited earlier.

“Putting Benefits in Dollar Terms Cheapens the Worth of the Environment”

A third common criticism of benefit–cost analysis is that it relies on monetary measures of the benefits of environmental amenities such as clean air or endangered species preservation. Critics ask how we can attach monetary values to such “priceless” resources without devaluing them. The problem with this critique is that we need a common yardstick for comparing costs and benefits. Expressing them in dollar terms is often convenient, simply because the costs of implementing a government policy are often naturally expressed in monetary terms, such as the increase in electricity bills that would result from more stringent pollution controls on power plants.

Of course, a critic might respond that any attempt to make benefits and costs commensurate is problematic, regardless of the metric used. In a world of limited resources, however, tradeoffs must be made between competing demands. Weighing costs and benefits is simply a way of assessing those tradeoffs. As Nobel laureate Robert Solow has argued,

Cost–benefit analysis is needed only when society must give up some of one good thing in order to get more of another good thing. . . . The underlying rationale of cost–benefit analysis is that the cost of the good thing to be obtained is precisely the good thing that must or will be given up to obtain it.²⁶

A related critique holds that there is a moral imperative to protect the environment, making the monetary value of no consequence. A problem with such arguments is that they overlook other conflicting but equally valid appeals to morality. If saving the endangered spotted owl in the U.S. Pacific Northwest is the morally right thing to do, what about saving the jobs of loggers whose livelihoods depend on cutting down trees? If protecting the Amazonian rainforest from slash-and-burn agriculture is a moral imperative, what about making sure that people clearing that land to farm it have enough to eat? At the very least, expressing benefits and costs in monetary terms can help inform these tradeoffs.

To make this concrete, let’s consider the control of mercury emissions from electric power plants. Some might argue that the health risks from mercury pollution are so great that mercury should be controlled at all costs. If this argument sounds appealing, ask yourself how much you

would be willing to pay in higher electricity bills each month in order to help pay for the costs of installing and operating mercury controls. Would you be willing to pay \$5 a month? \$10? \$100? \$1,000? How would your answer depend on the benefits to you and your family (and society at large) from controlling mercury pollution? Presumably, if the benefits were small enough and the costs were high enough, you might conclude that the emission controls would not be worthwhile. Spending more money on higher electricity bills to pay for pollution control necessarily entails spending *less* money on something else that you value. It is in this sense that money is a useful measure of benefits—not because we think there ought to be an intrinsic dollar value put on everything but because tradeoffs must be made.

Finally, it is worth pointing out that whether we realize it or not, we end up putting implicit values on environmental quality through the decisions we make every day. If you buy conventionally grown fruit and vegetables rather than the more expensive organic alternatives, you are putting a value on the environment. The same is true if you drive rather than bike to work or to school. The question is not whether we put a value on the environment. Rather, the question is whether or not we make that value explicit.

“Benefit–Cost Analysis Ignores the Losers from a Policy”

A fourth major criticism of benefit–cost analysis is directed more broadly at the notion of using economic efficiency as a criterion for social welfare. As we have seen, efficiency is concerned with the overall net benefits to society from a policy, not with who gains and who loses. For example, the allowance trading program of the 1990 Clean Air Act Amendments reduced sulfur dioxide emissions from electric power plants by roughly 50 percent in the latter half of the 1990s, at much lower cost than had been expected. (We’ll learn more about the trading program in Chapter 10, when we discuss market-based environmental policies.) Careful analysis has found that the benefits from that program—principally reduced mortality and morbidity from air pollution in cities of major power plants—far outweighed the costs of reducing SO₂ emissions.²⁷

One contributing factor was the availability of very low-sulfur coal from the Powder River Basin in Wyoming. In the congressional debates during the run-up to the passage of the amendments, Senator Robert Byrd of West Virginia led the opposition to the new law, largely out of a concern that it would threaten the jobs of miners of high-sulfur coal in his home state. Although the benefits were much larger than the costs in

What Happens When Costs and Benefits Are Not Considered Systematically?

A compelling argument for benefit–cost analysis is that when such analysis is not performed, actual policy may reflect implicit biases rather than reasoned considerations.²⁸ A study of endangered species management by the U.S. Fish and Wildlife Service (USFWS) illustrates this point. Andrew Metrick and Martin Weitzman gathered data on the amount of money spent on endangered species protection, along with data on a variety of species characteristics, including “scientific” attributes such as genetic uniqueness and degree of endangerment, along with “visceral” attributes such as the size of the animal. They found that the USFWS tends to lavish money on protecting animals that are physically appealing, even though they have reasonably large breeding populations and close genetic relatives in no danger of extinction. At the same time, little money is spent on species that face far greater risk of extinction (such as the monitor gecko or the Choctawahatchee beach mouse) or that are genetically unique (such as the Red Hills salamander or the Alabama cave fish). Indeed, the amount of money spent to protect listed species was strongly and positively correlated with size but *negatively* correlated with an objective measure of endangerment.

In short, spending by the USFWS on endangered species is biased toward “charismatic megafauna” such as grizzly bears rather than adhering to the scientific principles and priority setting that ostensibly guide decisions. We are not arguing that society should elevate scientific standards over other reasons people might value endangered species. Rather, the point is that the USFWS is not meeting the criteria it sets for itself as its goals, in part because of an absence of a calculation of the benefits and costs of protecting different species. As the authors conclude in another article on the topic, good stewardship requires “confronting honestly the core problem of economic tradeoffs—because good stewardship of natural habitats, like almost everything else we want in this world, is subject to budget constraints. The evidence suggests that our actual behavior may not reflect a reasoned cost–benefit calculation.”

aggregate, West Virginia coal miners may well have ended up worse off as Wyoming coal displaced some West Virginia coal. The costs of sulfur dioxide control also fell heavily on electric utilities in the Midwest, which passed the costs on to their ratepayers in the form of higher electric bills. The benefits from clean air, meanwhile, accrued to residents of downwind cities and to anglers and hikers who gained from the reduction in acid rain in the Adirondacks and elsewhere.

A program can be economically efficient, therefore, and still not make everyone better off. In simple terms, efficiency is about “maximizing the size of the pie,” whereas distributional equity is about who gets what share of the pie. This potential conflict between efficiency and distributional equity is fundamental, and it is an issue that merits careful consideration in each instance. If efficiency ignores something as critical as distributional equity, you may ask, why should we use it as a benchmark for policy? The answer is twofold: Efficiency provides a welfare criterion that is both fundamentally sound and implementable in practice.

The philosophical foundations of efficiency date back to the work of Italian economist Vilfredo Pareto at the end of the nineteenth century. Pareto proposed that one policy is superior to another if at least one person is strictly better off under the first policy and nobody is worse off. A policy is Pareto efficient if—and only if—no member of society could be made better off by an alternative policy without making at least one person worse off.

As a criterion for comparing a proposed policy with the status quo, Pareto efficiency has a certain appeal. If policy A makes at least one person better off than policy B, without harming anyone, would we not always want to adopt policy A?²⁹ When we try to apply this criterion in the real world, however, a drawback becomes clear: It is much too strict. Using it as a guide to making policy would almost always favor the status quo. Think of nearly any medical or technological breakthrough that was undoubtedly beneficial to society as a whole, and you can come up with at least one group of people who were hurt by the change. The rapid expansion of the railroad in the second half of the nineteenth century spurred the development of the western United States, brought fresh meat and produce to urban populations, and led to the rise of Chicago as a great city, but it also put bargemen out of business on now-obsolete canals and contributed to the overexploitation of natural resources such as the American bison. The development of the personal computer, for all its undeniable benefits, also devalued the skills of professional typists and shuttered countless old typewriter stores.

How to reconcile the obvious benefits to society from such changes with their smaller but very real costs? A way out of this conundrum was proposed separately by Nicholas Kaldor and John Hicks in

A policy is Pareto efficient if—and only if—no member of society could be made better off by an alternative policy without making at least one person worse off.

the 1930s. Under the Kaldor–Hicks or “potential Pareto” criterion, policy A is chosen over policy B if policy A would make at least one person better off without making anyone worse off, *provided that suitable transfers were made from the winners to the losers*. Crucially, the Kaldor–Hicks criterion does not require that those transfers are actually carried out. In other words, the modified criterion can be thought of as satisfying the strict Pareto criterion in principle, given appropriate (but possibly unfulfilled) transfers. Of course, compensating the losers from a given policy is much easier to describe than to carry out. Nonetheless, it remains crucial that such transfers be possible, even if only in theory, because the possibility of such compensation ensures that there is a net surplus from the policy. The winners from a policy can only compensate the losers (even in principle) if the benefits from the policy are greater than the losses. Thus, a policy satisfies the modified criterion if and only if it produces greater net benefits than the alternative. Satisfying the Kaldor–Hicks criterion is equivalent to maximizing net benefits.³⁰

The Kaldor–Hicks approach also helps clarify the relationship between efficiency and distributional equity. Because the compensating transfers need not actually take place, an efficient policy may lead to one group gaining much at the expense of another. If we are interested in distributional equity, this indifference to whether the transfers occur is a crucial flaw.

Three main responses are possible to such a critique. First, one can simply point out that efficiency and distributional equity are competing goals that both deserve consideration when we are shaping environmental policy, or policy in any other sphere. Second, we might argue that over time and across a broad range of policies, the gains and losses enjoyed by any particular group of people in any particular case tend to cancel each other out—so that in the long run, pursuing efficient policies will improve everyone’s lot. If today’s winners are tomorrow’s losers, then we need not concern ourselves too much with the distributional implications of particular policies. However, such a sanguine view of affairs is clearly

A policy satisfies the potential Pareto criterion if it would be Pareto efficient provided that the winners from the policy compensated the losers, but such transfers do not actually have to take place.

naive when it comes to the impacts of much public policy in market economies such as the United States, where in the absence of redistributive programs the pursuit of economic efficiency is likely to further concentrate wealth in the hands of a few.

A third response to the critique

is to find creative ways to spread the gains from efficient policies. A real-world example of just such a compensation program is the Northwest Forest Plan, enacted in 1993 by the Clinton administration. The forest plan grew out of concern over declining populations of northern spotted owls and the old-growth habitat they depended on. To protect the owl and its habitat, the plan sharply reduced the allowable timber harvest on 24.5 million acres of federally owned lands in Washington, Oregon, and northern California. Loggers in the local timber industry were among the most vocal opponents of such conservation measures; pickups and logging trucks sported bumper stickers reading “Save a Logger: Eat a Spotted Owl” and “Spotted Owl Tastes Like Chicken.” In response to such concerns, the Clinton administration proposed a Northwest Economic Adjustment Initiative. That program provided \$1.2 billion over 6 years in additional funding to cushion the blow to logging communities, through job retraining, rural development assistance, and direct payments.

Benefit–Cost Analysis Under Uncertainty

In many cases, the benefits and costs of a policy may depend on key parameters that are uncertain. For example, the benefits of fuel economy standards will depend in part on the future price of gasoline. Or the source of uncertainty may be the challenges inherent in estimating the health effects of pollution; in estimating the benefits from reducing concentrations of soot and other fine particles, EPA uses two empirical estimates of the effect of particulate matter on mortality, derived from two different peer-reviewed epidemiological studies.

If the uncertainty is limited to one or two key parameters, as in those examples, then it can be dealt with by performing a sensitivity analysis. Benefits and costs are calculated under a range of values for a given parameter (e.g., high or low gasoline prices) and displayed along with those from the base case, allowing the analyst to determine how sensitive the main results are to alternative assumptions. This is similar to estimating benefits and costs under a range of discount rates, another common practice.

In more complex situations, with multiple sources of uncertainty, analysts may use a more sophisticated method called a Monte Carlo analysis. In this approach, the analyst specifies the likelihood of various parameter values (for example, an oil price of \$80 per barrel in the year 2020 may be deemed more likely than prices of \$150 per barrel or \$20 per barrel). Benefits and costs are then calculated hundreds or even thousands of times, each with a different set (or “draw”) of parameter values. The result

Analyzing Uncertain Benefits and Costs of Regulation: EPA's Carbon Pollution Standards for New Power Plants*

In early 2012, the U.S. Environmental Protection Agency was preparing a proposal for the first-ever performance standards for carbon dioxide emissions from new coal and natural-gas-fired electric power plants. (Power plants that already existed at the time the rule was proposed would be covered by a subsequent regulation.) The agency intended to tie the standard to the emissions rate from a state-of-the-art, natural-gas-fired plant. Because natural gas emits roughly half as much CO₂ as coal per megawatt-hour of electricity generation, the proposed standards would effectively prohibit the construction of “conventional” coal-fired power plants (those without equipment to capture and store carbon emissions).

EPA planned to argue that the regulation would have zero costs and zero benefits. This was because under a business-as-usual scenario without the new standards, no conventional coal-fired power plants were expected to be built. That expectation was partly due to the cost of complying with regulations on other pollutants such as mercury. But the main reason that no new coal plants were planned was historically low natural gas prices, the result of the vast supplies of natural gas that could be tapped by the newly perfected technologies of horizontal drilling and hydraulic fracturing, or fracking. If the current low price of natural gas continued into the future, conventional coal-fired power plants would remain uneconomical regardless of the new carbon standards.

EPA's argument of zero costs and zero benefits raised eyebrows at the Office of Information and Regulatory Affairs, part of the White House Office of Management and Budget, which has responsibility for reviewing regulations. How could such a significant and even historic regulation—one that would effectively end the construction of conventional coal-fired power plants in the United States—have zero costs? Others asked, “If the regulation genuinely had zero benefits, why bother pursuing it when there was sure to be fierce opposition to the regulation on Capitol Hill?”

These questions pointed to the need for a more sophisticated approach. The key was to realize that whether coal-fired plants would be built in the absence of the regulation depended critically on the price of natural gas, which was uncertain. To correctly assess the expected benefits and costs of the regulation, the analysis needed to at least account for the possibility that future natural gas prices would be high enough to make coal plants economically attractive again.

*The analysis presented in this discussion is taken from U.S. Environmental Protection Agency, *Regulatory Impact Analysis for the Proposed Standards of Performance for Greenhouse Gas Emissions for New Stationary Sources: Electric Utility Generating Units*, EPA-452/R-12-001 (March 2012), pp. 5-31–5-35. Figure 3.3 is based on numbers in table 5-7 and footnote 51 of the RIA.

Analyzing Uncertain Benefits and Costs of Regulation *continued*

EPA's resulting calculation for a representative new coal-fired plant, reported in the Regulatory Impact Analysis (RIA) for the proposed rule, is illustrated in figure 3.3. For natural gas prices below \$9.60 per mmBtu, coal would be uneconomical, and the costs and benefits of the standard would be zero. That was the business-as-usual case that EPA had assumed at first. For natural gas prices above that level, however, a utility would prefer to build a coal-fired plant. Note that immediately above this price threshold, the cost savings from coal would be essentially zero, but the social damages from increased pollution would be significant. Therefore, for prices in some range above the threshold, the proposed standards would yield positive net benefits, because the extra cost of having to burn natural gas instead of coal would be outweighed by the benefits from lower emissions. As the natural gas price rose higher, the extra costs of natural gas would increase, while the benefits from lower pollution remained constant. Above some price point (\$12.70 per mmBtu in the figure; EPA's analysis identified a range of \$12 to \$14 per mmBtu for this second threshold), the extra costs would outweigh the benefits, and a prohibition on conventional coal plants would begin to impose net social costs.

The final step in EPA's analysis was to consider the likelihood that the price would fall into each of those three ranges. Historical data showed that natural gas prices had never reached the \$9.60 per mmBtu threshold on an annual average basis and had exceeded it even temporarily in only 8 of the previous 120 months. And the advent of fracking meant that natural gas prices were expected to be much lower than historical levels (more on this in Chapter 6). Even the most pessimistic natural gas forecasts available projected prices remaining below \$9.60 per mmBtu through 2035. These facts suggested that the probability of the proposed standard yielding net benefits was positive but small, and the probability that it would impose net costs was essentially zero. As a result, *expected* net benefits were greater than zero (although small in magnitude).

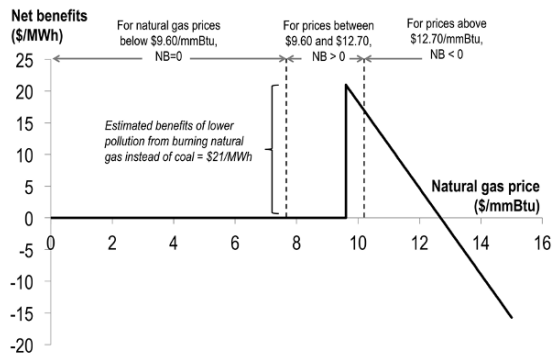


Figure 3.3 Illustrative estimated net benefits of EPA's proposed carbon pollution standards for new power plants.

is a *distribution* of benefits and costs rather than a single estimate. Rather than simply comparing the outcomes under a discrete set of parameter values, as in a sensitivity analysis, the Monte Carlo approach allows analysts to explore a much broader range of possible benefits and costs and to see which outcomes are more or less likely (given the specified distributions of the parameter values). For example, two policies with similar *average* net benefits may have different probabilities that net benefits will be much higher or lower than expected, a finding that could be of interest to policymakers interested in minimizing the risk of bad outcomes (such as a policy imposing net costs).

Conclusion

This chapter has fleshed out the concept of economic efficiency that we developed in Chapter 2. In order to put the efficiency criterion into practice, we need to be able to measure costs and benefits. Costs are more straightforward to define and to measure. Even so, economics raises two key insights: The costs of doing one thing depend on the opportunities forgone, and the burden is often shared by many groups in society. The benefits from environmental amenities—as from any other good—ultimately derive from what people are willing to pay to secure them. Economists have devised a range of clever approaches to infer such willingness to pay from how people behave when they go on vacation, buy a house, or accept a job. Nonetheless, in some cases—for example, the existence of an endangered species—there is no alternative but to ask people how much they value a good.

Even if we can measure costs and benefits, it is important to have a firm grasp on why economic efficiency is a desirable goal for society. As a criterion for policymaking, efficiency has two important strengths: It is grounded in a strong welfare justification (recall the Pareto criterion it is based on), but at the same time it is readily implemented in practice. Nonetheless, there are powerful critiques of efficiency and of the related practice of benefit–cost analysis. In particular, a single-minded focus on efficiency can mask deeply unfair outcomes, if the positive net benefits reflect large gains to one group of people at the expense of another. Keeping this blind spot in mind, however, we shall use efficiency throughout the remainder of the book as our basic benchmark for evaluating environmental policy and the management of natural resources.

4

The Efficiency of Markets

How well do market economies deal with environmental problems? One school of thought holds that government regulation is costly and intrusive, hindering innovation and economic growth. Conservative talk radio host Rush Limbaugh captured this sentiment when he wrote, “The key to cleaning up the environment is unfettered free enterprise. . . . Capitalism is good for people AND for other living things.”¹ On the opposite side of the political spectrum, observers such as Amory Lovins and former vice president Al Gore agree that firms can cut costs and boost profits by “being green”—for example, by reducing pollution or increasing energy efficiency.² If such a connection between corporate profits and environmental protection were to hold in general, it would suggest that government regulation is unnecessary, because firms will find it in their own interest to protect the environment.

So how well would the free market perform on its own? In this chapter, we discuss the advantages of well-functioning markets. Under certain conditions, market outcomes are Pareto efficient. In other words, markets allocate the production and consumption of goods in a way that maximizes the net benefits to society. This result is evoked by Adam Smith’s famous image of the “invisible hand” (one inspiration for this book’s cover). Without any explicit coordination, the interactions of individual consumers and producers, each motivated by self-interest, nonetheless combine to advance the common good. In general, therefore, free markets are a socially desirable means of allocating goods and services. But not always: In some important cases, markets can fail. In the next chapter, we shall discuss several closely related notions of market failure that are ubiquitous in the environmental realm. As we shall see, competitive markets—although typically effective institutions for allocating resources—are unlikely to

provide adequate levels of environmental quality without some government intervention.

Before diving into our discussion of markets and market failure, one point must be emphasized. Although market economies may give rise to environmental problems, the capitalist system should not be construed as the sole *cause* of those problems. The scale of environmental problems in capitalist economies such as the United States often pales next to the devastation wrought in socialist economies such as the former Soviet Union. We focus here on how environmental problems arise in market economies, not because such problems arise only in market economies but because markets are the dominant form of economic organization in the world today.

Competitive Market Equilibrium

Let's start by defining what markets are and describing how they work.

Defining Markets

In everyday activity we take part in a range of different kinds of markets: retail stores such as supermarkets or clothing shops, farmers' markets where local farms sell produce from individual stands, open-air flea markets, or capital markets such as the New York Stock Exchange. More broadly, the United States and other Western countries have decentralized market economies in which the price and quantity of goods are determined by the interaction of supply and demand rather than by a central state government.

Markets are not the only way to allocate goods or services. An obvious counterexample is a planned economy, such as that of the Soviet Union during most of the twentieth century. China's economy combines central planning and state-owned enterprises (which are typically insulated from market forces by subsidies and regulatory control) with some free enterprise. State-owned enterprises are also common in Europe, though to a smaller degree. But much more broadly, even in a market economy a range of other allocation mechanisms exist: You might purchase antique furniture at an auction, bid for collectibles on eBay, win a prize in a raffle or lottery, receive food stamps from government agencies, gain admission to a competitive college or university through the admissions process, or secure a prestigious medical residency through the medical matching system.

These are all ways to exchange goods and services. What, then, are the key characteristics of a market?

A market is a decentralized collection of buyers and sellers whose interactions determine the allocation of a good or set of goods through exchange.

Note several points from the definition. First, a market is an institution for allocating goods from those who produce or own them to those who want to buy them. Of course, markets are not the only means of allocating goods, as the examples in the preceding paragraph demonstrate. Second, markets are based on the *exchange* of payment for goods or services rather than one-way allocation of scarce resources (as in the cases of food stamps, a lottery, or the medical match system). Finally, and perhaps most importantly, markets are *decentralized*. This crucial characteristic distinguishes them from auctions (where buyers and sellers are brought together but exchange is arranged by the auctioneer) and from centralized economies (where a government planner orders producers to make specified quantities for sale or distribution to consumers).

Demand and Supply

Even if you have never studied economics, you probably know the bedrock of microeconomics: supply and demand. To understand market outcomes, we need to describe the behavior of sellers and buyers. First, consider buyers, who make up the demand side of a market. As the price of a good falls, some people who were already buying the good decide to consume more, while some people who chose not to buy the good at higher prices start to purchase it. In other words, as the price falls, the quantity demanded by consumers rises. This relationship between price and quantity, sometimes called the “law of demand,” holds for almost all goods.

We can describe consumers’ behavior by a *demand curve*. A demand curve summarizes how much buyers in the aggregate will buy at a given market price, with all other factors (such as the prices of other goods) held constant. The underlying relationship between price and quantity can also be stated another way. At each quantity, the demand curve summarizes what buyers are willing to pay for one more unit of a good, given how much they have consumed already—their marginal willingness to pay.

Figure 4.1 depicts a hypothetical demand curve for coffee. Note that quantity is on the horizontal axis and price on the vertical axis. The demand curve slopes downward, because lower prices lead to larger quantities demanded. For the sake of discussion, let’s imagine that this represents the demand for gourmet coffee in a college town.

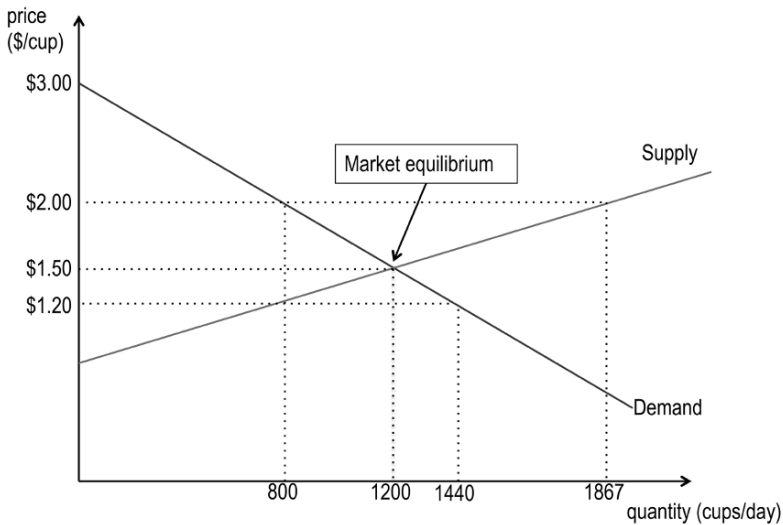


Figure 4.1 Demand and supply curves, and market equilibrium.

Now consider the sellers of a good—the supply side of the gourmet coffee market. A *supply curve* summarizes the relationship between price and quantity on the supply side (see figure 4.1). The supply curve represents how much the coffee shops in your town (collectively) are willing to sell or produce at a given price. Equivalently, for any given quantity, the supply curve represents the amount the coffee shops are willing to accept in order to produce one more unit of a good—one more cup of coffee. Note that the supply curve slopes upward. As the price of a good increases, so does the quantity firms are willing to supply. At higher prices, existing producers can expand their output in a range of ways: hiring more workers (perhaps increasing wages to do so); paying higher prices to secure larger quantities of needed inputs, say coffee beans; running their shops more intensively, in-

A demand curve summarizes how much buyers in the aggregate will buy at a given market price, with all other factors held constant. A supply curve summarizes the relationship between price and quantity produced by the sellers of a good.

curring higher costs; or perhaps even opening new shops. Moreover, higher prices allow new firms with higher costs of production to enter the market, expanding output further.

Market Equilibrium

What happens when supply and demand interact? The answer is that the market will eventually settle where

demand equals supply. This combination of quantity and price is called the *market equilibrium*, where the forces of demand and the opposing forces of supply just counterbalance each other. In fact, this market equilibrium

The market equilibrium is the combination of quantity and price for which demand equals supply.

is almost always a stable equilibrium; the market will “automatically” tend to establish this equilibrium price and quantity and to maintain them as long as the underlying factors that drive demand and supply (such as income, people’s tastes, and production costs) remain unchanged.

Market equilibrium is illustrated in figure 4.1. At a price of \$1.50, the demand curve shows us the amount consumers will want to buy (1,200 cups per day), which is exactly equal to the amount sellers will want to produce at a price of \$1.50, as shown on the supply curve. This point—and only this point—is a market equilibrium. To see why, suppose that instead of selling coffee at \$1.50 per cup, producers raise the price to \$2.00 per cup. Why can’t this be an equilibrium? At a price of \$2.00 per cup, some consumers will cut back on their purchases (switching to tea or another substitute, or reducing consumption of hot beverages altogether); they will buy only 800 cups per day at that price. However, sellers will want to sell even more than before because of the higher price. At \$2.00 per cup, it may be profitable to hire more workers and perhaps stay open later; figure 4.1 shows that sellers are willing to produce more than 1,800 cups at that price. With supply greater than demand, this situation would create a coffee surplus and hence market pressure to move back toward the equilibrium we identified earlier. How would this happen? Given the surplus, sellers would have a hard time finding buyers for all the coffee they produced. A smart coffee vendor would soon lower the price to attract more buyers, and other sellers would follow in order to compete for buyers’ business. The market price would continue to fall, until it reached the equilibrium price of \$1.50 per cup (and the quantity sold would rise in response, until it reached the equilibrium quantity of 1,200 cups).

Now consider a price lower than \$1.50, such as \$1.20. Such a price cannot be an equilibrium price either. Once again, there is a gap between the amount producers would be willing to sell and what consumers would be willing to buy. At a price of \$1.20, coffee vendors will want to sell only about 800 cups, and at this bargain price, consumers will want to purchase more than 1,400 cups. This time, the gap creates a shortage of coffee (demand greater than supply) and attendant market pressure to return to equilibrium. Seeing that demand is strong enough to sustain

a higher price, coffee vendors would increase prices (decreasing buyers' demand) until the equilibrium was reached again.

Only when the amount producers want to sell equals the amount consumers want to buy—in this case, 1,200 cups per week, at a price of \$1.50—is there an equilibrium, with no tendency for sellers to lower their prices or buyers to offer to pay more. In other words, the market “clears”—all buyers are able to find sellers and all sellers are able to find buyers—leading us to refer to the equilibrium price (\$1.50) as the market-clearing price.

The Efficiency of Competitive Markets

So far, we have described the workings of the market mechanistically—that is, we have predicted what will happen in a particular market, given the behavior of buyers and sellers (as represented by the demand and supply curves). But we've also told you that under certain conditions, the market outcome maximizes net benefits to society. That is, as a general matter, markets are Pareto efficient. How does this happen?

Demand and the Benefits to Consumers

Suppose you would be willing to pay up to \$1,200 for a laptop computer—in other words, you would be indifferent between paying \$1,200 and getting the laptop or not having the laptop at all. If a laptop computer actually sells for \$1,000 and you buy it, you have in effect received a surplus of \$200—the difference between what you would have been willing to pay (the benefit you get from using it) and what you actually paid.

The same principle is at work in markets in general. Remember that a demand curve traces out consumers' marginal willingness to pay, which is the measure of their benefit from consuming a good. Because the market price is determined at the margin, and demand curves slope downward, it will generally be the case that almost all consumers who make purchases in a market pay less for the good than they would be willing to pay. For example, in figure 4.1 the price of coffee is \$1.50 per cup, but some consumers would be willing to pay as much as \$3 a cup. Economists call the resulting net benefits consumer surplus: the benefit consumers get from purchasing a good over and above what they actually pay for it. In figure 4.1 we can see this amount for each cup of coffee by reading from the demand curve to find the benefit from consuming that cup and subtracting the price paid to obtain that cup (the vertical distance between the demand curve and the price at each quantity). When we combine these vertical distances for all the cups of coffee purchased in a market, consumer

surplus covers the whole area below the demand curve and above the price. This area is shaded and labeled in figure 4.2, for a market in which the equilibrium quantity is Q^* and the equilibrium price is P^* .

The difference between what consumers would be willing to pay for a good and what they actually pay is called consumer surplus.

Note that we are in effect using the demand curve as a measure of the marginal benefits to consumers. That should make sense. After all, we defined benefits in Chapter 3 as measured by willingness to pay. Although the discussion of that chapter was focused on environmental amenities, the same principle applies to all economic goods. From an economic perspective, the marginal benefit an individual receives from consuming a particular good or service can be measured by her marginal willingness to pay for that good or service.³

Supply and the Costs to Producers

Next, consider the supply curve. Recall that we defined it as representing the relationship between the quantity of a good produced and the price producers are willing to accept to produce one more unit.

It turns out that the supply curve traces the *marginal cost curve* of the industry. Marginal cost is the cost of producing one more unit of a good.⁴

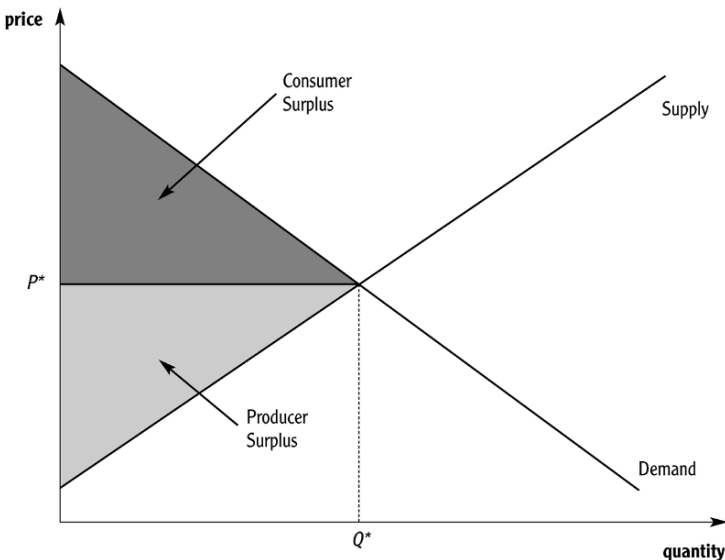


Figure 4.2 Demand, supply, market equilibrium, and welfare measures.

Suppose you are in charge of a firm that produces widgets and sells them for \$5 each. If the cost of producing an additional widget is less than \$5, you can make money by increasing production. For example, if the marginal cost of making a widget is \$2, you will earn a net return of \$3 (\$5 minus \$2) from making another widget. On the other hand, if the marginal cost is greater than the price, then you should not produce any more.

Now suppose that your marginal cost increases with the number of widgets you make. (That will usually be the case, at least in the short run when the size of the factory is fixed: As output rises, it becomes more and more costly to increase production as firms near the limit of their available capacity.) How many widgets do you want to make? The answer is: Produce widgets up to the point where the marginal cost equals the price. As long as marginal cost is less than the price, increasing output will raise revenues by more than cost; the opposite occurs once marginal cost exceeds the price. The same reasoning can be applied at the level of an industry made up of many firms. Firms in the aggregate will produce output at the level where price equals marginal cost. Therefore, the supply curve runs along the marginal cost curve.

In figure 4.2, we have shaded the area below the price and above the supply curve. This area is analogous to the measure of consumer surplus and—not surprisingly—is called producer surplus. What does producer surplus represent in economic terms? Recall from the discussion in Chapter 2 that the area under a marginal cost curve corresponds to total costs. The same principle applies here. As a result, the area under the supply curve corresponds to the costs of production. Meanwhile, the entire area under the price line equals price times quantity, or total revenues. Thus producer surplus equals the difference between the revenues to suppliers and their cost of production—that is, net revenue to producers.

Market Efficiency

The sum of producer and consumer surplus is a measure of the net benefits produced by market interactions. A careful look at figure 4.2 should convince you that this measure of net benefits is maximized at the point where supply equals demand. Thus the market equilibrium achieves the maximum possible net surplus from the market.

Indeed, we can see this efficiency result as a direct application of the equimarginal rule we discussed in Chapter 2. We have already seen that the market equilibrium is given by the intersection of the demand and supply curves. We just argued that those curves can be thought of as, respectively, marginal benefit and marginal cost curves. This suggests that

market outcomes equate marginal benefit with marginal cost and hence must be efficient.

This is an extremely powerful result, but the efficiency of markets is not guaranteed in all cases.⁵ In particular, it depends on three necessary conditions. First, markets must be competitive, in the sense that all firms and consumers must take prices as given. That is, individual firms and consumers must be unable to manipulate the market price in their favor. This is a reasonable assumption in settings where there are many firms or producers competing with each other to sell a good to a large number of potential buyers. When this assumption fails, however, market efficiency fails with it. For example, a single firm with a monopoly over a product will have the power to set its own prices rather than allow them to be dictated by the market. Rather than set a price that just covers the incremental cost of production, therefore, the monopolist will set a higher price that maximizes its own profit. This higher price means that less is sold than in a competitive market, with correspondingly lower surplus to consumers. For example, at the turn of the twentieth century, Standard Oil (owned by John D. Rockefeller) had a near monopoly on U.S. oil production and sales. More recently, governments in the United States, Europe, and elsewhere have sought to reduce monopoly power in the markets for personal computer operating systems (Microsoft), diamonds (De Beers), and eyeglasses (Luxottica).⁶

The second condition that must be met for markets to be efficient concerns the information available to firms and consumers about the quality of the good or service being traded. This information must be *symmetric*, understood equally by buyers and sellers. A classic example of asymmetric information is the used car market, in which the seller of a car typically knows much more about its quality than any prospective buyer. In such cases, sellers have an incentive to take advantage of their private information; knowing this, wary buyers are likely to be less willing to pay for goods or services whose quality is uncertain. Similarly, a buyer of health insurance knows more about her personal health conditions (and how likely she is to need that insurance) than does the seller of a health insurance policy. When information is asymmetric, markets may fail to allocate goods and services efficiently.

Finally, for markets to be efficient, they must be *complete*, that is, they must capture all the good and ill effects resulting from a market transaction. In particular, the costs of a good or service must be fully paid for by those who produce it, and the consumer who buys that good or service must enjoy the entire benefit from it. When this condition fails, demand

and supply no longer reflect marginal benefit and cost, because part of the benefits or costs accrues to other consumers or firms.

Our market efficiency result can be formally stated as follows:

- A market equilibrium is efficient if the following conditions are met:
 1. The market is competitive, meaning that firms and consumers take prices as given.
 2. Firms and consumers both have good information about the quality of the good or service being traded.
 3. The market is complete, in the sense that all relevant costs and benefits are borne by the market participants (the firms and consumers involved in transactions).

In many real-world markets these three conditions are likely to hold, at least approximately. Although market power certainly arises in the real world, many markets are reasonably competitive. For example, retailers such as supermarkets, bookstores, and clothing stores operate in fairly competitive industries. If your local supermarket raised the price of bananas or cereal too far above prevailing prices, consumers would simply buy those items elsewhere. If Barnes & Noble charges more for books at its online site, consumers will switch to Amazon. Similarly, consumers and firms typically have adequate (or at least symmetric) information about the quality of goods and services, and when information asymmetries do crop up, market solutions often arise as well, as when car dealerships certify used cars. In general, therefore, free markets are a socially desirable means of allocating goods and services.

When any of these three conditions are not met, however, markets are no longer efficient. In the language of economics, they are said to fail. Failures of the first two conditions are the focus of extensive literatures in economics and much debate in the real world. When it comes to the environmental realm, the problem most often is with the third condition, that

is, the complete market assumption.

The resulting market failures are the subject of the next chapter.

In general, free markets are a socially desirable means of allocating goods and services. When the key conditions for market efficiency are not met, however, a market failure is said to occur.

Conclusion

This chapter has explained how supply and demand interact in a market. The key result that emerges from our discussion is that competitive

markets are efficient, at least in many cases. When buyers and sellers take prices as given and have good information about the quality of goods and services, and when all costs and benefits are borne or enjoyed by consumers and producers, the decentralized interaction of self-interested individuals leads to socially desirable outcomes. This is a powerful result, and demonstrating it stands as one of the crowning achievements of economics. Nonetheless, although it is quite general, it does not apply in every case. In particular, as we shall see in the next chapter, market failure is prevalent in the environmental realm, because the costs and benefits of environmental protection and sound resource management are often left out of the calculations of individuals and firms.

5

Market Failures in the Environmental Realm

We've seen that competitive markets are often efficient. If that were the end of the story, this would be a much shorter book (and we, as environmental economists, would have much less to think about). Of course, when it comes to the environment, the assumptions underlying the efficiency of markets commonly fail to hold. Economic activity often gives rise to unwanted byproducts, such as water and air pollution, that impose indirect costs on consumers or firms downstream or downwind. And these costs are not usually captured in markets. For example, car exhaust is a major contributor to smog. But although drivers pay for gasoline, tires, maintenance, and so on, they do not pay for the pollutants they send into the atmosphere.

Economists call this sort of thing a *negative externality*. By driving their cars, people impose a cost on others in the form of poor air quality. That cost is invisible to the drivers themselves, however; in the parlance of economics it is *external* to their decision making. Note the asymmetry between benefits and costs. If you drive to work, you gain the entire benefit from driving—in terms of convenience, comfort, and so on—but you share the costs of greater pollution with everyone else around you.

The problem of incomplete markets can also arise in other ways. Some environmental amenities, such as biodiversity, are enjoyed by lots of people, whether or not those people help pay for them. Economists call such goods *public goods*. A market failure arises because some individuals will end up being free riders: Rather than helping to provide the public good themselves, they merely enjoy what others provide for them.

A third class of environmental problems is known as the *tragedy of*

the commons. When a natural resource—such as a fishery or an underground aquifer—is made available to all, individuals will tend to exploit the resource far beyond the optimal level. This problem arises because the incentives of individuals diverge from the common good.¹ We call it a tragedy because everyone would be better off if they could all commit themselves to act less selfishly. Thus, individually rational actions add up to a socially undesirable outcome.

In this chapter, we examine each of these three problems in turn. We demonstrate why markets fail to be efficient in each case. We also show how these seemingly distinct categories of environmental problems are linked at a fundamental level. You may already see the deep similarities between them. For example, air quality can be described in terms of a negative externality (your automobile exhaust makes my air worse), as a public good (clean air is enjoyed by all, so individuals have too little incentive to provide it), or as a commons problem (each driver overuses the shared atmospheric commons). Similarly, overfishing problems, though typically couched in the language of the tragedy of the commons, can also be described as a negative externality or as a free rider problem.

Externalities

We gave an intuitive definition of a negative externality earlier, in the example of automobile emissions. More generally, we can define an *externality* as follows:²

An externality results when the actions of one individual (or firm) have a direct, unintentional, and uncompensated effect on the well-being of other individuals or the profits of other firms.

Note three keywords in the definition: *direct*, *unintentional*, and *uncompensated*. For example, because your health and happiness depend in part on how clean the air is, automobile drivers have a direct effect on your well-being. *Unintentional* is included in the definition to rule out acts of spite or malice. (It is the *effect* rather than the action that is unintentional. I may decide deliberately to use a gasoline-powered lawnmower, without the intent of my action being to pollute the air or disturb the neighbors.) Finally, *uncompensated* implies that the responsible actor does not compensate the damaged parties (or is not fined) for his actions. This rules out market transactions or bargaining between individuals.

Second-hand cigarette smoke is a common example of an externality: A smoker in a bar ignores the effects of her smoking on nearby patrons

(except for her companions, which is why smokers turn away from their friends and blow smoke over their shoulders). At a larger scale, air pollution from factories or power plants represents an externality for downwind populations. Externalities can also impose costs on other firms as well as individuals. In the Pacific Northwest, logging in forested headwaters degrades spawning habitat for salmon, and hydroelectric dams hinder the fish on their way upstream. Both activities adversely affect commercial fishers.

Although environmental problems are typically framed as negative (harmful) externalities, positive (beneficial) externalities are also common. For example, a firm that carries out research and development often produces knowledge that its rivals can use—a positive externality, because some of the benefits from the research are captured by firms that do not contribute to its expense. If my neighbors keep their houses and flower gardens well maintained, the value of my house is likely to rise; thus I benefit from their actions, but they do not reap that additional gain. (We will return to the topic of such positive externalities when we consider public goods in the next section.)

How Do Externalities Cause Market Failure?

To see why the efficiency of markets breaks down in the presence of an externality, consider an oversimplified version of the steel industry. Steel furnaces typically burn coal, emitting sulfur dioxide, nitrous oxides, and particulate matter. Suppose for simplicity that there is a fixed relationship between the amount of steel produced and the amount of pollution emitted: For example, for every thousand tons of steel, 1 ton of sulfur dioxide is emitted. (This would be true if steel mills were unable to install pollution control devices or switch to cleaner fuels and thus could reduce pollution only by reducing output. We will relax this assumption when we consider policy instruments in Chapter 8, but for now it proves useful.)

Now consider the marginal damages of pollution as a function of the amount of steel produced. (We will continue to assume, just as we did in Chapter 2, that the marginal damages of a ton of pollution rise with the amount of pollution; as a result, marginal damages also increase with the amount of steel produced.) In figure 5.1, we have drawn such a function (labeled *MD* for *marginal damages*), along with the familiar supply and demand curves. In the absence of regulation, each steel producer will (quite rationally) ignore the damages caused by its pollution when deciding how much to produce and consider only its own costs. Indeed, this explains the term *externality*: The damages from pollution are external to the firm. The supply curve therefore corresponds to the private marginal costs of

steel production: the costs of the labor, fuel, materials, and so on needed to make one more unit of steel. (We have labeled the curve *PMC* on the graph, for *private marginal costs*.) This is exactly the same kind of supply curve we saw in Chapter 4. But now there is another cost of production, which the steel producers do not pay: the damages from pollution. What is the efficient level of steel production, given this externality? As always, efficiency means maximizing net benefits to society as whole, and it involves equating marginal benefit and marginal cost. Just as in the previous chapter, marginal benefit corresponds to the demand curve. However, the externality means that the *social* marginal cost is no longer equal to the supply curve, which reflects only the private marginal cost. Instead, social marginal cost (labeled *SMC* in the figure) equals the private marginal cost (paid by the steel industry) plus the marginal damage from pollution. The efficient quantity, where $D = SMC$, is labeled Q^* on the graph.

Now, let's compare that efficient outcome with the one that would result in a free market. Because the supply curve is unaffected by the pollution damages, so too is the market outcome. The unregulated market equilibrium occurs where supply equals demand: the quantity labeled Q_M on the figure. Note that $Q^* < Q_M$: The unregulated market equilibrium results in too much output (and thus too much pollution). In the absence of regulation, the market yields too much of a bad thing.

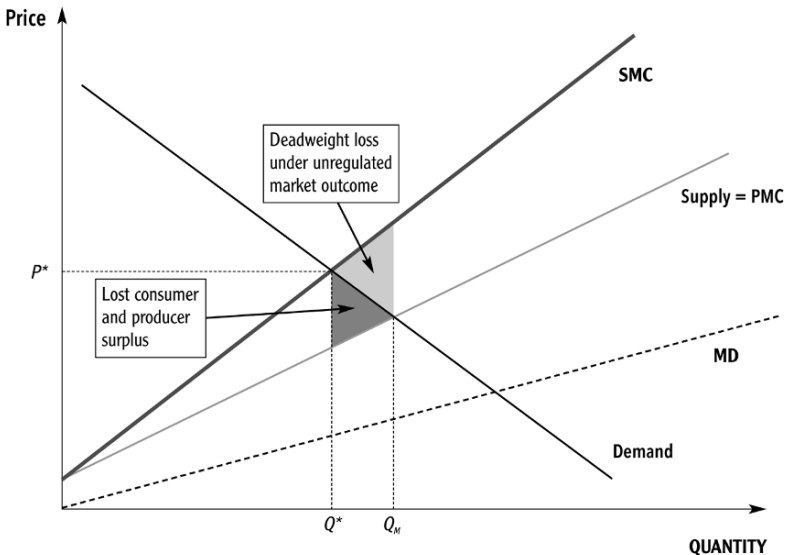


Figure 5.1 A market with a negative externality.

In the previous chapter, we saw that the competitive market outcome is efficient when there are no externalities. In the presence of an externality, that result no longer holds. The divergence between the market equilibrium and the efficient outcome arises precisely because the steel producers do not bear the full costs of production. If they paid the marginal damages from each unit of pollution, then their private costs would coincide with social costs, and the supply curve would trace out the *SMC* curve rather than the *PMC* curve in figure 5.1. In that case, the pollution damages would be *internalized*, and the market outcome would once again be efficient.

To understand the consequences of the market failure, consider what happens when we correct it. Suppose we move from the unregulated market outcome (Q_M) to the efficient outcome (Q^*). The amount of steel produced falls, while the price paid by consumers rises. Therefore, reducing output lowers the sum of consumer and producer surplus. (This must be the case, because it remains true that the market equilibrium maximizes their sum.) The value of this lost surplus is illustrated by the lower shaded triangle in figure 5.1.

How, then, can moving to Q^* raise welfare, as measured by the total surplus to society? The answer is that social surplus now has an additional component: It equals the sum of consumer and producer surplus, minus the damages from pollution. In a sense, there is another segment of society that must now be accounted for: people who suffer from pollution caused by steel mills. Although curtailing output hurts consumers and producers of steel, it benefits people who are harmed by pollution. Starting at Q_M , the gain from reducing pollution damages outweighs the lost consumer and producer surplus, on the margin. That remains true until output falls to the efficient point, Q^* , where marginal social cost equals marginal benefit.

The upper shaded triangle in figure 5.1 equals the difference between the avoided pollution damages and the lost consumer and producer surplus. Thus it represents the net gain from reducing output from the unregulated level, Q_M to the efficient level, Q^* . Equivalently, we can think of the same triangle as the net loss to society that results from the unregulated market, relative to the efficient outcome. Economists call this kind of vanished social welfare *deadweight loss*; the name underscores the notion that such losses are not transfers from one group to another but rather losses to society as a whole. The deadweight loss, in this case, represents the lost social surplus due to the overproduction of steel (and pollution) in the unregulated market. Its size depends on the slopes of demand and supply and on the magnitude of pollution damages. Like consumer and producer surplus, the deadweight loss can be expressed in monetary

terms: With knowledge of the slopes of the curves in figure 5.1, we could attach a dollar value to the size of the inefficiency represented by the dead-weight loss triangle.

The deadweight loss from a policy refers to lost social surplus, not transfers from one group to another but rather a loss to society as a whole.

Public Goods

A second type of market failure in the environmental realm arises with public goods—goods that are shared by all and owned by no one. National defense is a classic example of a public good. A country's armed forces offer protection from invasion to the citizens living within the country's borders. Importantly, all citizens within the same country are afforded the same level of protection; furthermore, the security enjoyed or "consumed" by one citizen does not diminish that of her neighbor. Biodiversity is a leading example of an environmental public good. Greater genetic diversity makes the food supply more resistant to threats from parasites and disease and offers the potential for new medicines or industrially useful chemicals. Many people value the existence of rare or exotic species of animals and plants or of uninhabited wild places, even if they will never walk in those wild places or glimpse those animals and plants. As with national defense, everyone enjoys these benefits of biodiversity, and no one person's enjoyment reduces the amount available to others.

These public goods have two fundamental characteristics in common. First, public goods are *nonrival*: The amount of any individual's consumption does not diminish the amount available for others. Second, they are *nonexcludable*: Individuals cannot be prevented from enjoying a public good. In particular, even individuals who did not help to provide the public good can still benefit from it. (You breathe the same air whether you drive a Ford Expedition to work and a John Deere lawnmower on the weekends or ride a bicycle and cut the grass with a push-mower.) Any good that is nonrival and nonexcludable is by definition a public good. Of course, a particular good might be relevant only to a given region: Cleaner air in Denver is not a public good in Dallas. Similarly, some goods are "public" only for a limited and well-defined population: For example, a city park may be open to all city residents free of charge but not to outsiders.

We can represent these two characteristics in a box diagram (figure

A type of market failure in the environmental realm arises with public goods, goods that are shared by all and owned by no one.

5.2). The vertical axis measures nonexcludability; the horizontal axis, nonrivalry. Pure public goods (like the examples mentioned earlier) are located in the upper right corner of the box: They are both nonexcludable and nonrival. On the other hand, pure *private* goods—those that are fully rival and excludable—are in the opposite corner. A candy bar is a simple but illustrative example: If I have a candy bar, I can completely exclude you from enjoying it, and whatever I eat leaves less for everyone else. Most goods commonly traded in markets—shoes, clothes, furniture, and so on—are purely (or nearly so) private goods.

In between these extremes of pure public or private goods, we can think of goods as having varying degrees of “public-goodness.” Some goods exhibit one characteristic but not the other. For example, cable TV is highly nonrival: Under most circumstances, the quality of the signal does not diminish appreciably with the number of users. However, cable TV is perfectly excludable. In contrast to broadcast TV, the cable company can shut off the signal to a consumer who fails to pay. Thus cable TV occupies the lower right-hand corner of figure 5.2. (Because congestion is possible, even if infrequent, cable TV is not all the way at the far right-hand side of the box.) Goods such as this one that are nonrival but excludable are known as *club goods*.

At the other corner of our diagram we might put an open-access resource, such as a fishery or forest that is open to all. In this case, the good is fully rival: If I harvest a tree, it reduces the timber left for you. But by definition an open-access resource is nonexcludable. Note that in this case (as in others) the nonexcludability is not an inherent characteristic of the good but rather is a product of institutions. For example, contrast a freeway (such as an interstate highway) with a toll road: By institutional design, the freeway—but not the toll road—is nonexcludable. Or compare a major artery leading into a city at 8 a.m. to the same highway a few hours later. At rush hour, the highway is a rival good: My presence on the road lengthens your commute. But in late morning, when traffic is light, the highway is effectively nonrival: Additional drivers do not affect those already on the road.

Why Do Markets Fail to Provide Public Goods?

These characteristics imply that private individuals and firms, left to their own devices, will undersupply public goods. In other words, the market outcome will fail to be efficient. To see how this happens, let’s start with a simple example. Adam and Beth live on either side of a flower garden that they own in common. As far as they are concerned, the garden is

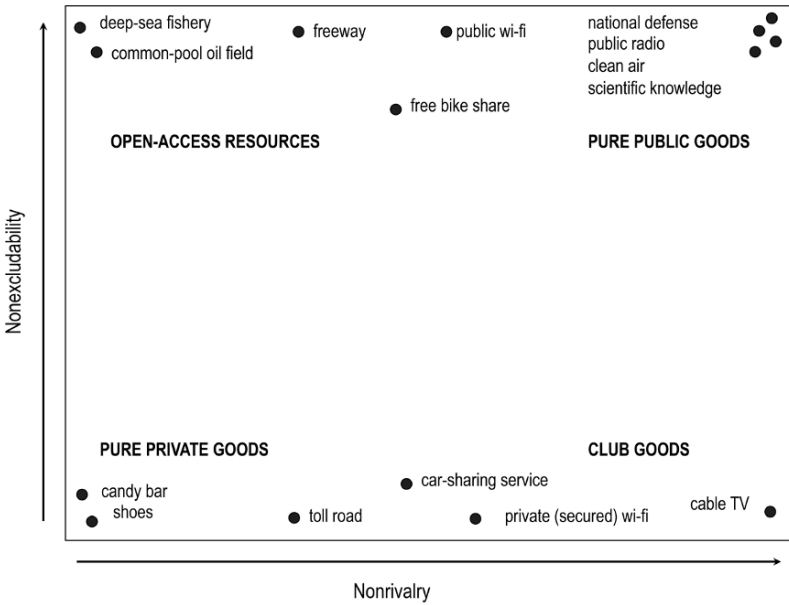


Figure 5.2 Characteristics of public goods, and real-world examples of various categories: pure public goods, pure private goods, open-access resources, and club goods.

essentially a public good, because they both have access to the garden, and one neighbor’s enjoyment does not diminish the other’s. (Most public goods involve many more than two people, of course; we’re keeping the example simple for the sake of illustration.) Figure 5.3 depicts the marginal cost of tending the garden, along with the two neighbors’ marginal benefit curves. The horizontal axis measures the quantity of this public good, which is to say the aesthetic appearance of the garden (e.g., its lushness, lack of weeds, plant health). Note that although both neighbors enjoy the garden, Adam values it more highly than does Beth.

As we did in the cases of the market and of externalities, we start by asking what private provision would yield on its own. In the real world, we might think of this as the “free market” outcome for public goods, in the sense of an unregulated market without government intervention. In our simple example, this corresponds to a lack of cooperation between the neighbors. Left to her own devices, Beth would tend the garden up to the level at which her private marginal benefits equal the marginal cost of provision (denoted Q_B on the graph). At that point, however, Adam’s marginal benefit exceeds the marginal cost: He prefers more of the public




Figure 5.3 Private provision of a public good.

good, even if he has to bear the entire cost himself. Indeed, Adam will willingly provide the higher amount denoted Q_A . In this case, Beth provides Q_B and Adam supplies the difference $Q_A - Q_B$. Of course, once Beth recognizes that Adam enjoys greater benefits from the garden, Beth will have an incentive not to supply any of the good. Instead, she may choose to free ride on Adam, who would willingly supply the entire amount Q_A on his own. (Note that even when Beth contributes, Adam will not provide more than Q_A , because beyond that point the incremental benefits to him are less than the incremental costs.) As a result, under the free market outcome Q_A units of the good are produced.

But this presents us with a seeming paradox: Both Adam and Beth would be better off if more of the good were provided! To see why, recall that at Q_A , the marginal cost of increasing the public good, just equals the marginal benefit to Adam alone. Thus the *combined* marginal benefit from tending the garden a bit more, to the two neighbors together, must be greater than the marginal cost. Yet as long as Adam and Beth act only in their own selfish interests, without cooperating, neither will make the extra effort to go beyond Q_A . The benefit to either individual is too small to make the extra cost worthwhile.

This parable illustrates an important result: Private provision of public goods is inefficiently low. We have already seen that efficiency requires that marginal benefit equals marginal cost. In the case of the flower garden (or any other public good) the relevant measure of marginal benefit is the social marginal benefit (SMB)—in this case, the sum of Adam’s and Beth’s private marginal benefits. Figure 5.4 compares the efficient outcome with the private provision we found earlier. In the figure, the curve marked *SMB* equals $MB_A + MB_B$. Note that it intersects marginal cost at Q^* , a level greater than what Adam provides on his own. Indeed, if Adam and Beth could find an equitable way to share the costs (perhaps if each one knew exactly how much the other valued the garden) and managed the garden jointly, they would tend it up to the efficient level Q^* .

The crux of the problem is that the quality of the flower garden is exactly the same for Adam as it is for Beth, regardless of how they divide up the total time spent tending it. The same principle holds for public goods in general. For example, New Haven, Connecticut, has a city park with a new playground and a renovated carousel. The most frequent users, of course, are families with young children, but all city residents enjoy the same free access and hence the same potential consumption. Or consider




Figure 5.4 Efficient provision of a public good, and the underprovision by the market.

the example of clean air. The clean air you consume is not measured by the volume of air that passes through your lungs but rather by the amounts of particulate matter, sulfur dioxide, and so on that are present in the air—concentrations that are the same for all individuals in a given city or region.

Because all individuals experience the same level of the public good, we must sum their marginal benefits when we consider the efficient amount of the good to provide. This stands in sharp contrast to a market for private goods. Because the benefit from consuming one more unit of a private good will be enjoyed solely by one person, we need only consider that individual's marginal benefit from the good. Moreover, each person consumes the good until her marginal willingness to pay just equals the price of the good. As a result, every individual in a market for a private good ends up with the same marginal benefit for that good (because all face the same price). In turn, this means that the marginal social benefit of a private good must equal its price, ensuring an efficient outcome.³

As a final point, note that public goods problems in the real world are much more complex than our simple two-person example. Neighbors like Adam and Beth might well cooperate; even if not, the difference between the efficient level of the public good and the free market outcome (what Adam provides on his own) may not be all that large. As the number of individuals grows, however, cooperation becomes more and more difficult. Moreover, the marginal benefits of individuals shrink relative to the marginal benefits to society. As a result, the free-riding problem becomes more acute as the size of the relevant public increases, with the gap between the efficient level of a public good and what is privately supplied growing apace.

Public Goods Provision as a Positive Externality

At first glance, public goods may seem to be a fundamentally different problem than externalities. After all, public goods are defined by a pair of specific characteristics, and our discussion to this point has said little about production or cost, which was central to the pollution externality of the previous subsection.

The free-riding problem—in which some individuals don't contribute at all to a public good, instead relying only on the contributions of others—becomes more acute as the size of the relevant public increases.

This apparent difference masks a deep underlying connection. In the two-person example, Beth enjoys the flower garden that Adam alone

tends. Adam is not compensated for those benefits, however; that's why we call Beth a free rider, after all. By caring for the garden, therefore, Adam provides a *positive externality* for Beth.

The externality involved in public goods provision plays out on the demand side, whereas the pollution externality considered before involves the supply side. The market failure from a negative externality such as pollution results from a divergence between private and social marginal costs: Private marginal costs exclude external damages and thus lie below true social marginal costs. The market failure that arises in public goods provision, on the other hand, comes from the divergence between private and social marginal *benefits*. When the public goods provider (Adam, in our example) considers how much of the public good to supply, he does not take into account the benefits enjoyed by the free rider (Beth) and thus understates the true marginal benefits to society. As a result, too little of the public good is provided. Just as the unregulated market tends to produce too much of a bad thing, so private provision of a public good yields too little of a good thing.

The Tragedy of the Commons

Our third category of environmental market failures is known as the tragedy of the commons, made famous by a well-known article of that name written in 1968 by Garrett Hardin.⁴ A number of people sharing common access to a natural resource will tend to overexploit it, unless they can develop effective government institutions (or social norms) to regulate its use. Hardin used the metaphor of an English pasture, or commons. The more sheep graze the commons, the less food is available for each of them. Each shepherd bears only a portion of the cost to the commons from the grazing of an additional animal (because that cost, in the form of less fodder, is spread over the herd as a whole), but he receives the entire gain from increasing his private flock. The result is that each shepherd puts too many sheep to pasture, from the point of view of the commons as a whole.

The tragedy is that the resulting overgrazing reduces the pasture's productivity. As a result, every shepherd would be better off if all could agree to restrict their flocks, but none has an incentive to do so on her own. If one shepherd pares down her flock, another may respond by adding a sheep to his own.

We can delve more deeply into the tragedy of the commons in the context of a very different resource (but one more familiar to modern minds): commuter roads. Consider the commute from a suburb to a hypothetical city. We'll suppose that drivers can reach this city by any one of

a number of smaller back roads; this route always takes 40 minutes (we'll assume that there are enough of these back roads that they never fill up with traffic).

On the other hand, it takes 30 minutes to reach the city on the highway—at least if there's no traffic. As more and more drivers use the highway, traffic slows the commuting time. For the sake of this simple example, let's say that each additional driver after the first one slows the commuting time (for every driver) by 1 second. If 121 drivers use the highway, for example, the commuting time becomes 32 minutes (thirty minutes plus 1 second for each of the 120 additional drivers). We'll use this extra commuting time as our measure of the marginal cost of additional drivers. Meanwhile, the marginal benefit from an additional driver is her own time savings from taking the highway rather than the back roads. The net benefit from the highway is just the total commuting time saved by all drivers, relative to what would happen if they all drove the slower route. (To focus on the problems of congestion and open access, let's set aside the well-known negative externalities associated with automobile emissions.)

What is the efficient number of cars on the highway? Suppose we start with no drivers taking the highway at all. The first driver saves 10 minutes in commuting time; that is the marginal benefit. Because there are no other drivers on the road, the marginal cost of the first driver is zero. Now suppose there are already N commuters taking the highway. What happens if we add one more? The marginal cost of the additional driver is 1 second for everyone else, or N seconds. The highway commute now takes $30 + N/60$ minutes, so the marginal benefit to the $N + 1$ driver (the time saved relative to the back roads) is $10 - N/60$.

A little algebra will show you that marginal cost equals marginal benefit when there are 301 drivers. That is, $N^* = 301$ (the asterisk denotes the efficient outcome). At that point, the commute takes 35 minutes by highway. The time savings to the 301st driver is 5 minutes, which is precisely the slowdown she imposes on all the other drivers.

Now what happens if access to the highway is unrestricted? Will the number of cars be efficient? The answer, as you have probably guessed, is “no.” To see why, consider what happens in the efficient scenario. When there are already 301 drivers on the highway, every driver taking the back roads thinks to herself: “If I take the highway instead, I will save almost five minutes (actually, four minutes and 59 seconds). Even though I will slow down everyone else, therefore, from *my* perspective I will be better off by taking the highway.”

What happens if everyone reasons this way? The number of drivers on

the highway will continue to increase until there is enough traffic that the highway commute takes exactly as long as the back roads—that is, 40 minutes! When each driver rationally makes her decision based only on her own costs and benefits, the aggregate outcome is far from rational for the group as a whole: The net benefit from the highway (the total time saved) ends up being reduced to zero.

The Tragedy of the Commons as a General Model

Hardin's metaphor of the commons applies to many natural resources but only under two important conditions. First, access to the resource must be unrestricted. Economists describe such resources as *open-access resources*. In terms of our discussion of public goods, an open-access resource can be thought of as a resource that is *nonexcludable* (like pure public goods) but not nonrival. The lack of exclusion usually stems from a combination of institutional and physical factors. A classic example is a deep-sea fishery, such as the cod fishery in Georges Bank in the northern Atlantic, where the distance from shore and the lack of national jurisdiction—along with political obstacles—make restricting access difficult. Similarly, forest reserves are typically open to harvesting by surrounding populations—especially in the developing world, where the funds necessary to patrol boundaries and prevent poachers are often lacking. Large underground aquifers such as the Ogallala, which underlies the Great Plains from north Texas to Nebraska, provide a common water source for the farmers and ranchers who live over them. The same can be said of wireless Internet routers. It is easy enough to install security measures, such as requiring a password, but as you no doubt know from experience, open-access networks are common. This is usually not a problem for the owner of the router, unless the residents in an apartment building across the street discover the free Internet access.

A second important condition is *diminishing marginal returns*. In plain English, as the number of people using the resource grows, the benefits from the resource must increase at a slower rate. In our commuting example, the total benefits from the highway increase as more and more drivers use the resource, because they each save time. However, the incremental time savings for each additional driver diminish, because traffic slows down as more and more drivers use the highway.

As in the case of public goods provision, there is a deep connection between the tragedy of the commons and the notion of externalities. In particular, diminishing marginal returns imply that each user of a resource imposes a negative externality on the other users. In the highway example,

the negative externality comes in the form of extra commuting time for the other drivers. This “open access externality” implies that unrestricted use will result in the overexploitation of a resource, because individuals will ignore the negative externality their effort imposes on others.

The Tragedy as a Collective Action Problem

The tragedy of the commons represents a particularly stark example of an externality, because when access to a resource is open, overuse tends to drive net benefits to zero. In contrast, of course, net benefits would be positive in the efficient scenario.

In the commuting example, the net benefits at the efficient number of drivers is $301 \cdot 5$ minutes = 1,505 minutes of commuting time. In the terms of our discussion in Chapter 3, a move from open access to the efficient outcome would be a Pareto improvement: Drivers permitted to use the highway would enjoy net benefits from doing so, whereas those excluded from it would be no worse off than they would be in the open-access case (because the time savings from using the highway are zero when access is unrestricted).

If people vary in their costs and benefits of using a resource, of course, restricting access to it will produce losers as well as winners. This helps explain why governments in the real world have found it so hard to impose restrictions on previously unregulated fisheries, such as the Atlantic cod fishery that once supported thriving New England towns. Even so, with appropriate transfers the increase in net benefits under the efficient scenario means that everyone could be made better off by imposing restrictions on access to the resource. (Recall our discussion of the modified Pareto criterion in Chapter 3.)

For this reason, economists (and political scientists) often describe the tragedy of the commons as a *collective action problem*. A collection of individuals—people, or firms, or even nation-states—may find itself in a situation where the group as a whole is better off if all contribute to the common good, but each individual member of the group has incentives to free ride.

The logic of collective action can also shed light on the difficulties of securing international cooperation on global environmental issues such as climate change. Every country would benefit from limiting the concentration of greenhouse gases in the atmosphere in order to reduce

Economists (and political scientists) often describe the tragedy of the commons as a collective action problem.

the probability of dangerous climate change, but because one country's efforts to reduce emissions benefit the rest of the world, each individual country faces an incentive to free ride on the efforts of others.

To make this more concrete, consider the following simplified example of countries considering whether to cooperate in reducing greenhouse gas emissions. To make things tractable, we'll imagine that there are only two countries in the world and only two options (reducing emissions or not), but the logic, and the lessons, apply more generally. To think about the incentives facing the two countries, we need to specify the costs and benefits to each country from cooperation. For the purposes of this example, we will draw on the most recent report by the Intergovernmental Panel on Climate Change (IPCC), an international body charged with producing comprehensive reports summarizing the scientific literature on climate change.

To keep things simple, let's assume that the two countries are considering whether to reduce emissions enough to limit greenhouse gas concentrations in the atmosphere to around 550 parts per million at the end of this century.⁵ As a starting point, drawing on estimates by the IPCC, we assume that if neither country acts, abatement costs are zero, but high concentrations of greenhouse gases raise average global temperatures by 4.5°C (8°F) above preindustrial levels. Drawing on economic models, we'll assume that this rise in temperature would result in damages equal to 6 percent of gross domestic product (GDP).

What if both countries take action sufficient to achieve the 550 ppm target? Again according to estimates reported in the IPCC, that concentration of greenhouse gases would raise the average global temperature at the end of the century by a little over 2°C (3.6°F) above preindustrial levels. Such a temperature change would entail estimated damages of 1 percent of GDP (relative to no warming at all). As a result, the *benefits* of limiting concentrations to 550 ppm—corresponding to the avoided damages from climate change—amount to 5 percent of global GDP. Modeling estimates reviewed by the IPCC suggest that achieving a 550 ppm target could cost 3.8 percent of GDP. Finally, for the sake of this example we assume that if only one of the two countries takes action, it still incurs abatement costs of 3.8 percent of its GDP, but it achieves only half the benefits that joint action would yield (i.e., resulting damages are 3.5 percent of global GDP). Importantly, these benefits are enjoyed by both countries equally, regardless of which one (or both) takes action.

Because each country chooses either to act or not, there are four possible outcomes: both countries cooperate, neither does, or one or the other acts alone. To help think through the incentives facing the two countries,

figure 5.5 summarizes the costs and benefits under all four scenarios in a single two-by-two matrix.

The rows correspond to actions that country A can take: “contribute” (i.e., take action to reduce emissions) or “shirk” (do nothing). Similarly, the columns represent the actions taken by country B. The numbers in the cells represent the net payoffs—damages plus abatement costs—to the two countries, with country A’s payoff in the upper left-hand corner of each cell. Note that all the payoffs in the matrix are negative, reflecting the fact that some amount of climate change is unavoidable and will leave the world worse off even if we take action to mitigate it. For example, if both countries agree to contribute, each incurs climate damages of 1 percent and abatement costs of 3.8 percent, for a payoff of -4.8 percent (depicted in the upper left-hand corner of the figure). If only country A contributes (upper right-hand corner), then it incurs costs of 3.8 percent while country B pays nothing. Because the damages from climate change in that case are 3.5 percent for both countries, the payoffs are -7.3 percent and -3.5 percent, respectively.

Notice that cooperation is economically efficient, in the sense that it yields the highest net benefits (lowest combined damages). If both countries contribute, the sum of damages and abatement costs is 4.8 percent of GDP; if neither does, the total payoff is -6 percent. (You can also see this by noting that joint action yields positive net benefits relative to doing nothing: A cost of 3.8 percent of GDP yields benefits of 5 percent.) But even though joint action is optimal, if each country considers only its own payoffs it has a strong incentive *not* to contribute.

To see why, we need to compare the payoffs to each country under each of the four possible outcomes. Suppose you are country A, choosing between contributing and shirking. Suppose that you expect country B to contribute. In that case, you would get a payoff of -4.8 percent from contributing and would receive -3.5 percent from shirking; thus shirking yields a higher payoff and is individually rational. What if you expect B to shirk? In that case, contributing yields -7.3 percent, whereas shirking yields -6 percent, so shirking is again preferable to contributing. The identical logic holds for country B. Thus, regardless of what action the other country takes, each country does better by shirking. In the language of game theory, the field of economics that studies strategic interaction, shirking is a dominant strategy.

The “game” we have just described is an example of a Prisoner’s Dilemma, commonly used to describe collective action problems—not just in the environmental arena, but in a broad range of social and economic

Player B
(payoffs in lower RH corner of each cell)

		Player B	
		Contribute	Shirk
Player A (payoffs in upper LH corner of each cell)	Contribute	- 4.8% - 4.8%	- 7.3% - 3.5%
	Shirk	- 3.5% - 7.3%	- 6.0% - 6.0%

Figure 5.5 Payoff matrix for the climate change collective action problem.

situations. (The name evokes a district attorney who induces two criminals in separate cells to rat out each other by promising each that he will get off easy if he confesses to the crime while the other stays silent.) Individually rational decisions produce a collectively suboptimal outcome. In the case of climate change, each country has an incentive not to limit its own emissions, instead free riding on the efforts of others. But if all countries follow this logic, all end up worse off than if they cooperated. Similar logic applies in the case of an open-access resource. In such cases, each actor has an incentive to exploit a resource—whether a fishery, a freeway, a pasture, or the global atmosphere—as long as her private gains from doing so outweigh the costs. But when all actors follow this reasoning, they all end up worse off than if they had cooperated to limit their use of the resource. This is the dilemma at the heart of the collective action problem. Solving it requires institutions of some sort—whether they take the form of government policies, treaties or other international agreements that promote cooperation, or more informal systems of social norms.⁶

Conclusion

If Chapter 4 extolled the potential virtues of markets, our extensive discussion of market failures in this chapter is an important reminder that

laissez-faire markets are a prescription for environmental problems. When markets are incomplete, individuals face the wrong incentives to change their behavior—to reduce the pollution they produce, contribute to the provision of public goods, or refrain from exploiting a common resource.

Although we have discussed negative externalities, public goods, and the tragedy of the commons in turn, these are not distinct problems but rather different ways of framing the same underlying market failure. The key to all three is the notion of *nonexcludability*. An externality arises when the good or ill effects of one person's actions, or a firm's operations, are not borne exclusively by that person or firm. People who cannot be excluded from enjoying a public good may prefer to free ride on what others provide rather than contributing to the good themselves. And those with open access to a common resource have little incentive to moderate their use of the resource, knowing that others will take what they do not.

Viewed in this way, economics provides a morally neutral explanation for environmental problems. Pollution and overfishing do not arise because polluters or fishers are bad people. Rather, the managers of polluting firms, like fishers, are simply trying to maximize their profits. (Indeed, the managers of private firms have legal responsibilities to their shareholders to maximize the firm's profits.) We might see this as selfish behavior, as indeed it is, but we have seen that selfish behavior by itself is not the problem, because in well-functioning markets it is the engine of efficiency. To an economist, the root cause of environmental problems concerns the *incentives* people face. The driving factor is not that individuals pursue their own interests but rather that in an unregulated market nothing aligns self-interest with the broader interests of society.

Framing the problem as a problem of incentives—rather than as a problem of morality or of markets per se—also points the way toward possible solutions. As we will see in Chapter 7, economic theory suggests that the way to deal with market failures in the environmental realm is not to avoid the use of markets but rather to fill in the incomplete nature of the market, whether by providing artificial price signals, assigning property rights, or even creating a market in environmental goods such as clean air.

6

Managing Stocks: Natural Resources as Capital Assets

Many people assume that natural resources have infinite value. But economics does not assume this, as we discussed in Chapter 2. As a general matter, economists treat natural resources as a subset of society's capital assets, no more or no less important than other types of capital. This human-centered approach is fundamentally different from other perspectives, such as deep ecology and intrinsic rights. You may or may not find it difficult to reconcile this approach with your own values with respect to natural capital. Nonetheless, we will approach natural resource management problems from the perspective of economic efficiency. This approach highlights the tradeoffs that must be made between competing uses of scarce natural resources, such as recreation, species habitat, and resource extraction.

In this chapter, we discuss the economics of nonrenewable natural resources such as oil and minerals. We begin with a discussion of scarcity as an economic concept, which incorporates more than simply the limited availability of physical resource stocks. We then present a simple two-period model of nonrenewable resource extraction, using it to understand the economic notion of scarcity and how resource owners will take it into account as they decide how much to extract. We develop the concept of marginal user cost, an extra cost of extracting nonrenewable resources that represents the opportunity cost of forgone future consumption. The discussion then considers the role of market power in natural resource markets. Finally, we stress the critical role of property rights in determining whether nonrenewable resource extraction in real-world markets will be efficient.

Economic Scarcity

Natural resource scarcity has economic and geologic dimensions. The critical point in the economic analysis of resource management is this: The stock of a natural resource, such as oil, depends not only on the physical availability of that resource within the earth's crust but also on its marginal extraction cost and the prices people are willing to pay to purchase it. For example, some parts of the oil sands of Alberta, Canada, where oil is produced by sucking sticky tar out of sandy soil, are viable commercial sources of oil only when the market price of crude oil is above about \$35 per barrel. Most North American tight oil resources (those that must be extracted from deep, impermeable rock formations such as shales) have a higher breakeven price—a price at which the net return to extraction is positive—of \$50 to \$70 per barrel. Below these prices, wells tapping these resources may lie idle (and new wells are unlikely to be drilled) because they cannot be operated profitably.¹ Effective stocks of natural resources continually expand and contract in response to technological change and resource prices, with high prices increasing the quantity of resources that are worth extracting and low prices reducing it.

A useful way of representing both the economic and the physical dimensions of resource stocks, known as a McKelvey diagram, is presented in figure 6.1. The original McKelvey diagram was developed by the U.S. Geological Survey to classify the stock of nonrenewable resources along its physical (horizontal) and economic (vertical) dimensions. Earth's total resource endowment is, of course, both unknown and fixed; it has only physical dimensions. But the portion of this endowment that is potentially useful to humans depends on both geological availability and economic value. In addition, technological progress continually affects both the costs of resource extraction (generally driving down these costs) and the value of specific resources to society. For example, mechanization and large-scale surface mining have lowered the cost of coal extraction over time. But the advent of gasoline-powered engines made coal obsolete in certain uses.

In thinking about natural resources from an economic perspective, we must be careful in interpreting some common indicators of physical scarcity, such as the reserve-to-use ratio or static reserve index. These ratios divide current known reserves of a resource by current annual consumption, thus measuring the number of years until the resource is exhausted. The problem is that such measures ignore the economic dimensions of scarcity and so offer an inaccurate picture of resource limits. For example,

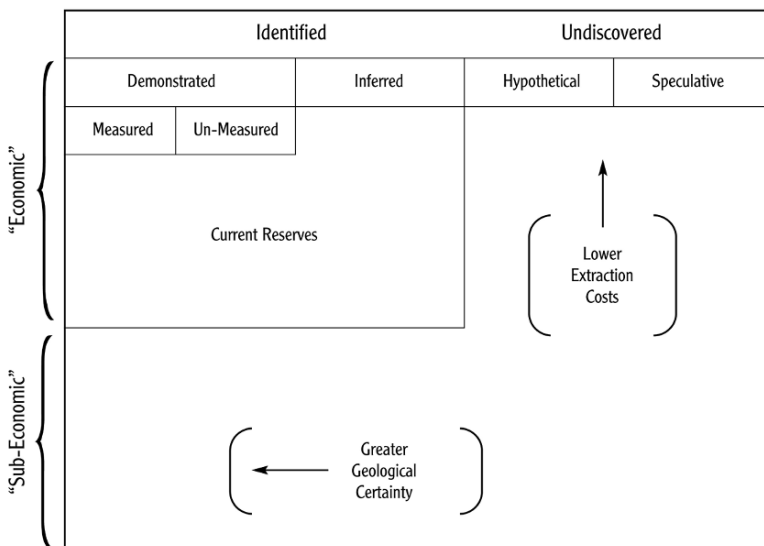


Figure 6.1 McKelvey diagram.

despite substantial increases in annual consumption, static reserve indices for iron, copper, aluminum, and zinc all increased over the period between 1950 and 2000. These increases may reflect exploration for and discovery of new reserves; technological change, which can lower extraction costs, making known reserves exploitable where previously they were not; and increasing commodity prices. More recently, global oil and natural gas reserve-to-use ratios have increased because of the boom in production from deep shale formations, made possible in part by hydraulic fracturing.

Failing to take the economic dimensions of scarcity into account is a common mistake—so common that even some prominent economists throughout history have made it. Stanley Jevons, a renowned nineteenth-century British economist, predicted in a book called *The Coal Question* that as Britain depleted its coal reserves, its economic power would decline precipitously. Jevons failed to account for the fact that an increase in the price of coal would spur the development of alternative sources, new extraction technologies, and greater efficiency in coal use, and he failed to anticipate the rise of oil and natural gas as alternative fuels.²

Although physical scarcity is only one dimension of economic scarcity, the limited physical availability of nonrenewable natural resources does affect their optimal rate of use. In particular, faced with limited stocks, maximizing the net benefits of a resource to society will gradually require

Nonrenewable Resources Can Become Less (Economically) Scarce

In 2000, natural gas and oil extracted from deep shale rock formations represented a very small fraction of total U.S. production. Since then, production from shales has boomed, and U.S. “proved reserves” of oil and gas (as well as the reserves-to-use ratios for these resources) have grown in recent years (figure 6.2).




Figure 6.2 Trends in U.S. proved reserves of oil and natural gas, 1973–2013, showing recent increases in both due to exploitation of newly economic resources.

Source: U.S. Energy Information Administration. “U.S. Crude Oil and Natural Gas Proved Reserves” (December 2014), Figure 1. Data from Form EIA-23L, “Annual Survey of Domestic Oil and Gas Reserves, 1977–2013,” American Petroleum Institute, 1973–1976.

Shale formations have existed since well before society began exploiting oil and gas resources. What has caused this sudden surge in exploitation? These changes have been made possible by innovations in technologies including hydraulic fracturing (“fracking”), directional drilling, and seismic imaging. Entrepreneurs in the early 2000s combined these technologies to unlock vast reserves of oil and gas from shale formations previously considered “sub-economic,” to borrow a term from the McKelvey diagram in figure 6.1. Fossil fuel prices and growing global energy demand were key driving forces in this process; they made it worthwhile for firms to take risks and invest in the research and development necessary to make these technologies profitable.³ The so-called shale revolution is but one recent example of the importance of incorporating economic factors, and not just current physical availability, when defining resource scarcity.

using less of it today and keeping more in the ground to use in the future than if the resource were in infinite supply. In addition, the market prices of nonrenewable resources will be higher than they would be if stocks were not limited, reflecting the impact of scarcity. Let's use a simple example of an oil well to look at the use of a resource over time and to demonstrate these dual effects of scarcity on the efficiency problem.

Efficient Extraction in Two Periods

Suppose we own an oil well, and we plan to pump oil from the well in two time periods, "today" and "tomorrow."⁴ (The labels "today" and "tomorrow" are just for convenience; the goal of the model is to explore the balance between current and future use of a scarce resource.) The demand for oil in each period is $MB = 10 - .5q$, where q is the quantity extracted; the marginal cost of extracting a barrel of oil (which might include labor and electricity, for example) is constant at $MC = \$3$.

First, let us assume that our oil supply is not limited but infinite. What would be the efficient quantity of oil to extract today? In order to figure this out, we would set the marginal benefits of extracting oil today equal to the marginal costs.

$$\begin{aligned} MB &= MC \\ 10 - 0.5q &= \$3 \\ q^* &= 14 \text{ barrels} \end{aligned}$$

Using this static efficiency rule as developed in Chapter 2, we would extract 14 barrels of oil today.

Now, let's introduce a limited stock; only 20 barrels are available. If we extract 14 barrels today, as we would like to, what would that leave for tomorrow? We would be left with only 6 barrels of oil in the ground. Given no change in demand and marginal cost between today and tomorrow, if we apply the static efficiency rule again tomorrow, we will want to pump another 14 barrels. But our remaining 6 barrels will fall well short of this goal (see figure 6.3).

Has the efficiency rule failed us? The problem we have just solved twice sequentially is *myopic*. We have intentionally ignored the limited oil supply and acted as though extraction of oil today is independent of the quantity left to extract tomorrow. In doing so, we have not identified a loophole in the efficiency rule. We have simply left out a very important cost from our scenario.

$$\text{Demand: } MB = 10 - 0.5q$$




Figure 6.3 The problem with (myopic) static efficiency in the case of scarce resources. Static efficiency would imply extracting fourteen barrels in each period—more than the total stock of twenty barrels.

When a resource is finite, like our oil well with a mere 20 barrels, one cost of extracting a unit of that resource is the lost opportunity to extract that unit in the future. In addition to the marginal cost of extracting a barrel of oil, we must account for the marginal cost of using up a barrel of oil, which leaves one fewer to use in the future. This extra cost is called *marginal user cost*, or *scarcity rent*. Accounting for the marginal user cost associated with oil extraction, like any cost increase, will reduce the amount of oil that we can efficiently extract today, leaving more in the ground for tomorrow.

Let us solve this problem again, this time taking the limited stock directly into account. The dynamic two-period problem we now solve differs from the static efficiency problem presented earlier in three important respects. First, because we are interested, today, in the value of extracting oil both today and tomorrow, we will need to discount the returns to oil extraction tomorrow to reflect the time value of money. This will help us to account for the fact that any oil left in the ground until tomorrow cannot be sold on the market today, and the proceeds from its sale cannot be invested to increase in value between the two periods. Thus the marginal benefits and marginal costs of oil extraction will be expressed in terms of *present value*—their value in today's dollars.

Second, we will introduce the stock constraint directly into our efficiency problem. To do this, we will define the quantity of oil available to extract tomorrow, q_2 , as the difference between the total stock (20 barrels) and the amount extracted today, q_1 .

Third, rather than setting the marginal benefits and marginal costs of

extraction in a single period equal to each other, we will equate the net marginal benefits (benefits, less costs) of oil extraction in each period. That is, we will start from the presumption that, in order to maximize the net benefits of this oil well, we must ensure that the net benefit of the last barrel pumped today is equal to the net benefit of the last barrel pumped tomorrow. If this were not the case, we could increase the overall net benefit of the oil well by redistributing our pumping plans over time. We will explore this assumption in detail later in the discussion. Here, we solve for the efficient quantities of oil to extract today and tomorrow, assuming a discount rate of 10 percent.

$$PV(MB_1 - MC_1) = PV(MB_2 - MC_2)$$

$$10 - 0.5q_1 - 3 = \frac{10 - 0.5q_2 - 3}{1 + .10}$$

$$7 - 0.5q_1 = \frac{7 - 0.5(20 - q_1)}{1.10}$$

$$7 - 0.5q_1 = \frac{0.5q_1 - 3}{1.10}$$

$$q_1^* = 10.19 \text{ barrels}$$

$$q_2^* = 20 - q_1^* = 9.81 \text{ barrels}$$

We suggested earlier that the efficiency rule in the presence of resource scarcity would cause us to use less of a resource today than we would if the resource were infinite. That is certainly the case in our oil well example. When we solved this problem myopically, thinking only about the net benefits of extracting this oil today, heedless of dwindling stocks, we planned to extract 14 barrels of oil. Now that we have incorporated the limited stock into our problem, the rules of efficiency tell us to extract just over 10 barrels of oil today, leaving the rest in the ground for tomorrow.⁵

A Closer Look at the Efficient Extraction Path

Why not split the well's contents exactly in half, extracting 10 barrels today and 10 tomorrow? The time value of money is the reason we extract just over half today and just under half tomorrow. Because the value of the oil we extract today can earn interest in an alternative investment between today and tomorrow, it is efficient to extract a bit extra today. In fact, if you experiment with interest rates other than the 10 percent

assumed here, you will notice that the higher the interest rate, the greater the difference between the amount of oil we will extract today in an efficient scenario and the amount we will leave for tomorrow. In seeking to maximize the net benefits of this oil well to society, these two different types of capital—oil and money—are fungible.

In fact, there are an infinite number of extraction quantities, today and tomorrow, that sum to our 20-barrel limit. Why is the specific extraction path we arrived at—10.19 today and 9.81 tomorrow—the efficient one? A diagram may help illustrate the intuition behind these numbers. Figure 6.4 plots the marginal net benefits, in present value terms, of oil extraction in each period. Oil extraction today increases along the horizontal axis from left to right, and extraction tomorrow increases from right to left. The two marginal net benefit curves intersect at the efficient allocation of extraction over time, the pair of extraction quantities for which we have just solved algebraically. Because the curves on this graph represent marginal net benefits, the total net benefits of this resource to society are measured by the area under these curves. And we simply cannot generate greater total net benefits by choosing any extraction path other than the efficient path we have identified. If we move to the right of the efficient allocation, extracting more today and leaving less for tomorrow, the value of net benefits lost tomorrow would exceed today's gains. And if we move to the left of the efficient allocation, extracting less today and leaving more for tomorrow, the value of net benefits lost today would exceed those gained tomorrow.

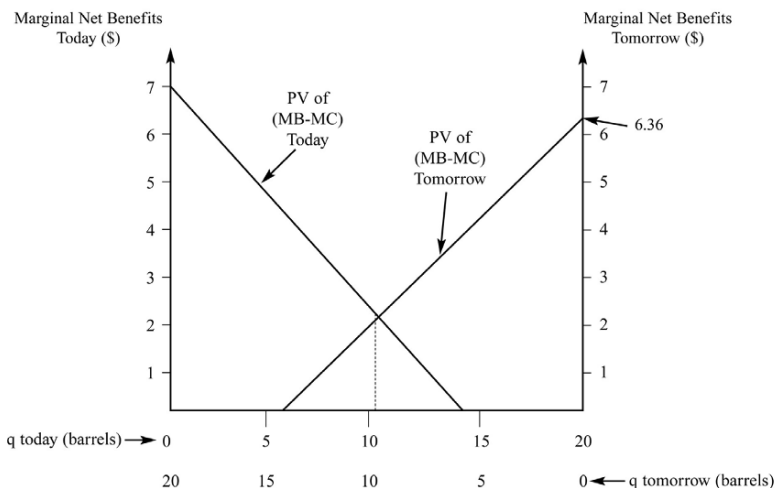


Figure 6.4 Nonrenewable resource extraction: the two-period model.

Marginal User Cost, a Special Externality

Earlier, we mentioned that extraction of scarce resources, such as oil in a finite well, imposes a cost above and beyond the marginal cost of extraction—a marginal user cost. Is it possible to identify this extra cost, either in our algebra problem or in the diagram? The marginal user cost of a barrel of oil from our well becomes apparent when we solve for the prices we can expect to collect for a barrel of oil in each period, today and tomorrow:

$$p_1^* = 10 - 0.5(q_1^*) \approx \$4.905$$

$$p_2^* = 10 - 0.5(q_2^*) \approx \$5.095$$

The market price of a barrel of oil is approximately \$4.91 today and \$5.10 tomorrow, yet the marginal extraction cost is only \$3. If you have taken an introductory microeconomics course (or carefully read Chapter 4 of this book), it may appear as though we have just violated one of the fundamental tenets of a competitive market: The price received for a good or service should be exactly equal to the cost of producing the last unit, or price equals marginal cost. This difference between price and marginal cost in the case of scarce resources like oil is, in fact, the user cost we discussed earlier.

When resources are limited, current consumption comes at the cost of forgone potential future consumption. The present value (at the margin) of these forgone future consumption opportunities is marginal user cost, or scarcity rent.

We can also think of marginal user cost as a negative externality to current oil consumption. Extracting today, we impose an extra cost on tomorrow: diminished supplies. This is not true only of oil, of course. If residents in Las Vegas, Nevada, an extremely arid city in the western United States, use large quantities of water to grow lush, green lawns, this may involve no scarcity rent if the water is from a large, quickly replenishable supply. However, if lawn watering draws down nonrenewable groundwater supplies, then this extra cost of diminished future supplies (lower aquifer levels) should be incorporated into water prices.

Water prices are rarely determined in a market. But oil prices usually are.

When resources are limited, current consumption comes at the cost of forgone potential future consumption. The present value (at the margin) of these forgone future consumption opportunities is marginal user cost, or scarcity rent.

In our example, the marginal extraction cost does not change between today and tomorrow, but the market price of a barrel of oil rises. Thus it appears that user cost (the difference between price and marginal cost) rises over time. This fact helps us to understand the alternative name for user cost: scarcity rent. As our stock of oil dwindles, oil becomes scarcer, thus scarcity rent—the extra cost of using up a barrel of oil today, which the owner collects in the form of a higher price—increases.

Earlier, we used a McKelvey diagram to demonstrate the physical and economic dimensions of scarcity. Marginal user cost is an economic indicator of scarcity that takes into account both the known physical limits of a resource and what we are willing to pay for that resource. If marginal user cost is an economic indicator of scarcity, how can we be sure that it will really be incorporated into market prices? After all, we have discussed many examples of externalities in this text, and in most of those examples, markets fail to account for environmental and resource damages, leading to inefficient outcomes and often requiring government intervention.

The answer to this question depends critically on the structure of property rights with respect to a scarce resource. Note that if we own the oil well we have been discussing, by extracting oil we impose a marginal user cost on *ourselves*, diminishing our own future supplies. Thus we have a strong incentive to account for that cost as we decide how much oil to extract. If we do not, we will not maximize the profits from our oil resource over time, and in a competitive market we will soon be out of business. This is a strong contrast to the examples we discussed in Chapter 5, in which environmental costs were borne by parties other than the externality-generating firms.

So when nonrenewable resources are privately owned and extracted in a competitive market, resource owners will account for scarcity in determining the optimal timing and quantity of extraction (the extraction path). They will treat oil resources, and other nonrenewable resources, like any other capital asset in their portfolio—as stocks that generate returns by the very nature of their scarcity.

The Hotelling Rule

In fact, when we consider nonrenewable resources as capital assets, it is clear that they must generate these returns at a very specific rate. Nonrenewable resource stocks should increase in value at a rate equal to that of other types of assets in the market. A useful benchmark here is the prevailing rate of interest, which represents the risk-free return an investor can earn in the market. If oil stocks in the ground were gaining in value at a

rate faster than the rate of interest, resource owners would extract nothing in the near term, leaving stocks in the ground to increase in value relative to money in the bank. If oil stocks in the ground were gaining in value at a rate slower than the rate of interest, resource owners would do better to extract all of the oil, sell it, and invest the proceeds.

Let us test this theory with our oil well example. We can calculate the rate of change in marginal user cost between the two periods as follows:

$$\begin{aligned} \frac{MUC_2 - MUC_1}{MUC_1} &= \frac{(P_2 - MC_2) - (P_1 - MC_1)}{(P_1 - MC_1)} \\ &= \frac{(5.095 - 3) - (4.905 - 3)}{(4.905 - 3)} \\ &\approx 0.10 \end{aligned}$$

This interesting theoretical result, that marginal user cost rises at the rate of interest, is called the Hotelling Rule, named for statistician and economist Harold Hotelling. The key to understanding the Hotelling Rule lies in realizing that in a competitive market the limited availability of nonrenewable resources such as oil strongly affects resource prices and extraction paths. Oil in the ground generates returns for its owner over time. It is a capital asset that can be spent today (through extraction) or saved for tomorrow. The price of spending it today is the lost scarcity rent it would generate by remaining in the ground. Likewise, the return to saving it for tomorrow is the rate of increase in scarcity rent.

If a market is in dynamic equilibrium, private owners of capital cannot increase their profits by reallocating their portfolios; if they could make more money by holding less capital in oil and more in some other asset, or vice versa, private owners would take advantage of that opportunity. The intuition here is like a *no-arbitrage condition*. The competitive pressures of the market for privately owned nonrenewable natural resources are extremely powerful, and thus one particular negative externality to depleting these resources—the fact that they will not be around to consume tomorrow—will be reflected in market prices and extraction rates.⁶ In this sense, at least, the markets for nonrenewable resources are complete.

This interesting theoretical result, that marginal user cost rises at the rate of interest, is called the Hotelling Rule.

Of course, there are other externalities generated through the extraction and consumption of petroleum; this should be clear from our earlier discussion of the economics of pollution control. And we have not reached the bottom line in our discussion of the impact of scarcity on markets; we will return to the potential impact of dwindling resource stocks on human welfare when we discuss the economics of sustainability in Chapter 11. But the Hotelling Rule helps us to understand why economists tend not to worry about the extraction decisions of private owners of nonrenewable resources in competitive markets.

What about Market Power?

One thing that may have contributed to global worries about running out of nonrenewable resources is the role of market power in markets for fossil fuels and other mineral resources. For example, one of the key players in the global oil market is the Organization of the Petroleum Exporting Countries (OPEC), a cartel that, similar to the monopolies discussed in Chapter 4, restricts the output of petroleum to raise market prices. OPEC does this by assigning individual production quotas to member countries. When this strategy succeeds, oil extraction is slower than the dynamically efficient extraction rate identified by the Hotelling Rule, and prices are higher. Thus, Nobel laureate Robert Solow has dubbed the monopolist the “conservationist’s friend,” if a conservationist is one who supports the use of natural resources at slower-than-efficient rates. As is typically the case when one of the major assumptions required for market efficiency fails to hold, the net benefits to society of an oil resource in the presence of market power are not maximized.

Perhaps more importantly from the perspective of non-OPEC countries, the periods in which the cartel succeeds in manipulating world oil prices, as it did during the “oil shocks” of the 1970s, leave us with memories of lining up for gasoline and a sense that our supply is scarce and vulnerable. But in fact OPEC has rarely succeeded in manipulating market prices in this way.⁷ High prices encourage production outside of the cartel (by non-OPEC members). And within the cartel, the incentive to “cheat,” producing more than a country’s production quota, often wins out, increasing production and driving down prices.

Recent concerns about China’s exercise of market power in the market for rare earths, a group of seventeen rare elements that play important roles in the production of high-tech devices such as phones and hybrid vehicles, and in gasoline refining, had a similar outcome. China produced almost all of the world’s rare earths by the early 2000s, and

political conflict between China and Japan in 2010 led to a Chinese export embargo, raising global concerns about access to these nonrenewable resources. Prices of these resources rose dramatically, drawing attention to their limited availability. High prices, in turn, encouraged the users of these elements (e.g., high-tech firms) to change their production processes, decreasing demand, and spurred new production in the United States, Japan, Australia, and other countries, increasing supply. These shifts in demand and supply stabilized rare earth prices.⁸ Thus, even when some firms or countries are able to exercise market power over nonrenewables, when prices communicate information about resulting resource scarcity, history suggests that markets themselves can often mitigate the impacts.

The Critical Role of Property Rights

We have expressed confidence that the individual decisions of private, competitive owners of nonrenewable resources will maximize the present value of nonrenewable resource stocks to society. However, part of the reason for this lies in the structure of property rights with respect to this class of resources. With few exceptions, nonrenewable resources such as oil and other minerals tend to be privately owned and traded in reasonably competitive markets (with exceptions noted earlier). Recall that marginal user cost is basically an externality that is really not external to the transactions between buyers and sellers of oil—sellers incur this consumption externality (reduced future supplies) themselves and thus will take it into account in determining how much oil to extract.

Were we to have used groundwater aquifers, rather than oil wells, as our example in the preceding discussion of nonrenewable resources, we would not have arrived so cleanly at the Hotelling result. We can make this general statement because, in contrast to mineral resources, groundwater supplies tend not to be privately owned and traded in competitive markets. When a farmer pumps water from a nonrenewable groundwater aquifer, the marginal user cost associated with pumping a unit of water is not incurred by that farmer; instead, the cost of diminished future supplies is spread among all of those who benefit from the aquifer. Thus, individual users of the resource have no incentive to take scarcity into account when deciding how much water to pump.

What would happen if a forward-thinking farmer did try to save a cubic meter of water in the aquifer for tomorrow rather than pump it today? Unlike the owner of the oil well, the farmer cannot assume that this unit of water will remain in the ground for his own benefit tomorrow; it is much more likely that another user of the resource (perhaps one

of his competitors down the road) will pump that unit of water, spoiling our farmer's good intentions. It is easy to see how we might arrive at a race to pump, quickly draining the resource, rather than extracting it in a dynamically efficient manner. Economists refer to this type of good—one from which potential consumers cannot be excluded and one that consumers compete to capture—as an *open-access* resource, as discussed in Chapter 5.

Although a nonrenewable groundwater aquifer would seem to be directly comparable to a nonrenewable oil well, the difference in property rights regimes between the two cases leads to very different outcomes. In the case of oil, we expect markets to do a reasonable job of ensuring

An Open-Access Resource: The Ogallala Aquifer

A good example of an open-access resource is the Ogallala Aquifer, which underlies approximately 174,000 square miles in the U.S. Great Plains, including portions of the states of Texas, New Mexico, Kansas, Nebraska, Colorado, Oklahoma, South Dakota, and Wyoming. The aquifer contains approximately 3.8 billion acre-feet of water and provides about 30 percent of all groundwater used for irrigated agriculture in the United States. One economic analysis, based on the variation in land values between irrigated and dryland farms in the region, estimated that the water's value ranges from 30 to 60 percent of the irrigated farmland sale price in the region.⁹

In most states that sit on top of the aquifer, groundwater is an open access resource. Unlimited quantities of water can be extracted by individual farmers, who incur only the costs of extraction (the water itself is free). Use of the aquifer for irrigation has proceeded at a rate conservatively estimated to be about ten times the rate of recharge (which in many parts of the aquifer is negligible). Since the resource was first exploited on a large scale in the 1940s, groundwater levels have dropped precipitously in some areas, particularly northern Texas, Oklahoma, and southwestern Kansas.

Because the Ogallala is an open-access resource, the marginal user cost associated with pumping water from the aquifer is truly an externality—no party incurs the full cost of pumping, thus the resource is depleted at a rate faster than the dynamically efficient rate. Collectively, society would be better off if irrigators did take into account marginal user cost, the marginal cost of aquifer depletion. But individual irrigators have no incentive to internalize this cost when making their water use decisions. Thus, without a change in property rights structures in the overlying states, the Ogallala aquifer will be depleted inefficiently soon.

dynamically efficient extraction; in the case of groundwater, we expect markets in their current structure to fail at this task.

Conclusion

In this chapter, we have introduced the first application of dynamic efficiency in the determination of optimal extraction rates for nonrenewable natural resources such as oil and coal. Although such resources are available in limited quantities on the earth, the economic concept of scarcity that we developed in this chapter takes many things other than physical limits into account, most importantly the effects of rising prices on demand for scarce resources.

We showed how private, competitive owners treat nonrenewable resource stocks as capital assets. In doing so, they extract resources at a rate that takes into account the limited physical stocks. The Hotelling Rule told us that this optimal extraction rate of nonrenewable resources maintains an asset market equilibrium, in which the rate of return to stocks in the ground equals the rate of return to alternative investments.

In this first natural resource management application, we encountered a situation in which real-world market outcomes do a pretty good job at approximating efficient outcomes, in large part because nonrenewable resources tend to be privately owned and traded in reasonably competitive markets. Even in situations of market power, market pressures tend to mitigate the influence of noncompetitive actors on prices and scarcity. We ended with a reminder that, where property rights are not well defined, extraction rates of nonrenewable resources can be expected to exceed efficient extraction rates. The situation of poorly defined property rights will be even more relevant to the discussion of renewable resources in Chapter 7.

7

Stocks That Grow: The Economics of Renewable Resource Management

Many renewable resources are more like an aquifer than an oil field. That is, they often are open-access or common property resources, and we do not expect them to be used in a dynamically efficient manner when markets are left to their own devices. This absence of property rights complicates economic analysis. A second complication is that the stock of a renewable resource, though limited, is not fixed. For example, the stock of fish in a fishery might be a certain number at any particular moment, but over time it depends on factors such as reproduction rates and predation, including human fishing effort.

In the case of nonrenewable resources such as petroleum, we were concerned with calculating the dynamically efficient rate of depletion of the resource. In contrast, with renewable resources we hope to calculate the size or timing of the efficient harvest, in many cases maintaining a sustainable flow of the resource in perpetuity. Because renewable resource stocks are functions of both natural systems and human behavior, the models we use to analyze them combine biology and economics—they are *bioeconomic* models. The first such model we will discuss is a bioeconomic model of a forest.

Economics of Forest Resources

When we were considering how fast to pump oil in the preceding chapter, we concluded that efficient management of an oil well (or any nonrenewable resource) requires extraction at a rate that maximizes net benefits.

Similar logic applies to forests. In economic terms, standing trees are capital assets that increase in value as they increase in volume over time. But allowing the trees to stand is also costly; we must consider the opportunity cost of alternative investments. Thus, we seek to identify the length of time to wait between timber harvests that maximizes the difference between total benefits and total costs (in present value).

The economic analysis of forest management raises two issues that were largely absent from our discussion of nonrenewable resources. First, oil and coal have value primarily as inputs to the production and consumption of other goods, such as energy. In contrast, the value of a forest is more complex. In addition to their value as timber for potential harvest, standing trees offer other benefits, providing species habitat and carbon sinks. To keep things simple, we will start with the problem of commercial timber extraction, and we will add other types of values later.

Second, forested lands exhibit a wide variety of property rights regimes, ranging from private ownership to open access. We start by considering a private landowner who makes rent-maximizing decisions about harvesting her trees. Later in the chapter, we will consider other property rights arrangements.

Forest Growth and the Biological Rotation

We begin with a simple model of forest growth. We can model the volume of timber in a stand of homogeneous trees as a function of time. Here we use the following volume function, pictured in figure 7.1, to describe this process:

$$V(t) = 10t + t^2 - 0.01t^3$$

At first, the rate of growth is very fast. Over time the trees continue to grow, but the rate of growth begins to decline (in our model, after about 33 years). At some point, depending on the species, climate, and a variety of other factors, the trees stop growing and begin to decay, resulting in declining volume (in our model, after about 71 years).

One candidate for the best interval at which to cut and replant these trees is the age that maximizes the *mean annual increment* (MAI), the average volume of the stand, $V(t)/t$. If we divide volume by time, we obtain the MAI curve depicted in figure 7.2, which reaches its maximum after 50 years of stand growth. This decision criterion makes some intuitive sense, because no other rotation yields a greater average volume of wood. For this reason, the maximum MAI is often called the *biological rotation*.

$$V(t) = 10t + t^2 - .01t^3$$




Figure 7.1 Timber volume in a forest as a function of time.

The biological rotation maximizes timber volume, but does it maximize the net benefits of the stand to society? We haven't yet introduced any economic information into our discussion, so you should guess that the answer to that question is "probably not." Figuring out the efficient rotation requires that we think about tradeoffs, as we did in the case of petroleum. For example, if we were to cut the trees after 40 years instead of 50, we would obtain fewer board feet of timber, but we would obtain the smaller cut 10 years sooner. Given the time value of money, this might make sense. Now we introduce some economic information in the simplest possible case: a single harvest.

Optimal Aging Problem: The Wicksell Rotation

Say that we are interested in the returns to harvesting this stand of trees once, with no concern for what will happen to this currently forested land after we extract our timber. This single rotation problem is essentially an "optimal aging" problem. The question "How long should I age a stand of trees?" is quite similar to the same question involving a bottle of wine or a fine cheese. To answer it, we must think about the returns to alternative investments, again represented by the rate of interest.

To solve this problem, think of the situation a private landowner would face each year. She would compare the net returns to cutting her trees this year to the net returns to waiting for 1 more year. As long as the net




Figure 7.2 Mean annual increment ($V(t)/t$) in a forest, as a function of time.

returns to cutting now are less than the net returns to waiting, she would prefer to keep her assets in standing trees. The net benefit-maximizing year in which to cut the trees would occur just as the net returns to waiting equaled the net returns to cutting. We can represent this point in a simple equality, in which the net returns to cutting now are on the left-hand side, and the net returns to waiting (in present value) are on the right:

$$(p - c) V(T_0) = \frac{(p - c) V(T_1)}{(1+r)}$$

where:

- p = timber price
- c = unit harvesting cost
- $V(T_0)$ = stand volume this year
- $V(T_1)$ = stand volume next year
- r = discount rate

If we rearrange some terms, we obtain the following:

$$r = \frac{V(T_1) - V(T_0)}{V(T_0)}, \text{ or } r = \frac{\Delta V}{V(T_0)}$$

Thus, it is efficient to harvest the stand when the rate of growth in timber volume, the rate of return to our capital asset (standing trees), is equal to the interest rate.

This is called the Wicksell Rule, and it can be applied to any optimal aging problem. If we harvest the stand before this point, the lost value of the incremental growth we would expect between this year and next would exceed the value of the incremental gains we would earn by depositing our net harvest proceeds in the bank to earn interest for 1 year. If we wait to harvest the stand beyond this point, the opposite would be true. The Wicksell Rule, like the Hotelling Rule, is a no-arbitrage condition. Just as in the case of oil extraction, if forest owners could make more money by holding less capital in trees and more in some other asset, or vice versa, they would take advantage of that opportunity.

There is an inverse relationship between the Wicksell rotation and the rate of interest. If the expected returns to alternative investments are very low, the Wicksell rotation is very long; a high interest rate implies a shorter rotation. For interest rates above 2 percent, the Wicksell rotation for a stand of trees described by the volume function in figure 7.1 is shorter than the biological (MAI-maximizing) rotation. Note that incorporating the time value of money into our model had the predicted result. As we guessed at the end of the previous section, under reasonable assumptions about the interest rate, we would prefer to accept a smaller total volume than that afforded by the biological rotation in exchange for cashing in our trees at an earlier date.

Efficient Forest Management over Time: The Faustmann Rotation

The time value of money is not the final wrinkle in the problem of optimal forest rotation. We have one more concern we did not worry about when discussing oil extraction: the value of the land on which our trees are growing. The problem of optimal rotation is really one of optimal land use. A landowner deciding when to harvest a stand of trees is concerned not only with the growth rate in the value of alternative

The Wicksell Rule for the optimal single rotation tells us that it is efficient to harvest a stand of trees when the rate of growth in timber volume, the rate of return to our capital asset (standing trees), is equal to the interest rate.

assets—that is, how much she might earn by cashing in her trees once and putting the money in the bank—but also with the value of her property as a whole. The problem requires an understanding of ongoing returns to forestry on a tract of land over time and a comparison of these returns to those from other potential land uses.

The landowner who is mindful of

the value of her resource as a whole faces a variety of choices each year. She could cut her timber and replant; she could wait 1 more year, then cut and replant; she could cut this year and convert the land to a new use, such as planting watermelons or building suburban tract housing; or she could cut this year and sell the land to a new owner, who might choose one of these land uses or something else entirely.

To solve the optimal rotation problem, which takes all these options into account, we introduce the concept of *site value*. Site value is the value of a forested piece of land, assuming that the landowner will implement efficient forest rotation in perpetuity; or—if forestry is not the most profitable use of that land at any point in the future—convert the land to its most profitable use. Site value allows us to compare the present value of expected future rents (benefits less costs) from forestry to those from other potential land uses, such as farming or residential development. In economic analyses of land use and land use change, this is how land prices are determined. Land prices are equal to the present value of expected future rents from land in its most profitable use. Thus, site value, which we will represent here using the letter S , captures all the competing land use options we mentioned earlier.

When the landowner considers site value, as well as the annual return from cutting and selling her timber, her yearly problem looks a bit different than the one we defined earlier. As before, she will seek to cut her timber and replant in the year in which the marginal net benefits of cutting are equal to the marginal net benefits of waiting 1 more year. We can represent this point in a simple equality, in which the net returns to cutting now are on the left-hand side, and the net returns to waiting one more period (in present value) are on the right. Everything is as before, but we have added two additional terms, site value (S) and the cost of replanting trees after the timber harvest (D).¹

$$(p - c) V(T_0) - D + S = \frac{(p - c)[V(T_0) + \Delta V] - D + S}{(1 + r)}$$

The easiest way to understand the intuition behind including site value on both sides of this equation is to think of S simply as the sale price of the land. The net returns to cutting in each period include not only the per-unit returns from timber less the cost of replanting but also the amount of money the landowner would make if she sold her land immediately after replanting. Even if the landowner has no plans to sell this land, S still must be included in each year's potential returns, because it represents the

opportunity cost, to her, of holding this land in forest rather than doing something else with it.

To simplify things, let us now consider $V(T)$ to represent not simply the volume of timber in the forest at time T but the net value of that volume, or $(p - c)V(T)$. In this case, we can reduce the equality above to

$$r[V(T_0) - D] + rS = \Delta V$$

The left side of the equation is the marginal benefit of harvesting now. The right side is the marginal cost of harvesting now (the extra timber volume that would accrue between this year and next, which the landowner forgoes through her impatience). Figure 7.3 illustrates efficient timber rotation. In early years, timber volume in the forest is growing quickly, so the increase in value is large; as growth slows, the benefits of cutting the forest approach the costs. At T^* the marginal benefits and costs of cutting are exactly equal. If the landowner rotates this stand at T^* in perpetuity, she will maximize the net benefits of the forest resource.

To continue with our discussion of efficient natural resource management guidelines as no-arbitrage conditions, we can rearrange terms to obtain

$$r = \frac{\Delta V}{V(T_0) - D + S}$$

The landowner should time her harvests so that the rate of return to her forest assets is equal to the prevailing rate of interest. Notice that, in contrast to the Wicksell Rule, here the landowner is interested in the rate of return to the value of her forested land, not just the timber volume. Thus we include S in the denominator. This rule for efficient forest rotation, taking into account both the time value of money and the opportunity cost of land, is called the Faustmann Rule.

We noted earlier that the biological rotation will, in general, be longer than the Wicksell rotation. How do these compare with the Faustmann rotation? In general, the Faustmann rotation will be the shortest of the three. Why is this the case? This should be obvious from a comparison of the Wicksell and Faustmann rules, given that the Faustmann Rule has a larger denominator on the right-hand side. Adding site value to the problem shrinks the rate of growth in the value of standing trees as capital assets in comparison to the Wicksell case.

But there is an intuitive explanation also. If the landowner solves the

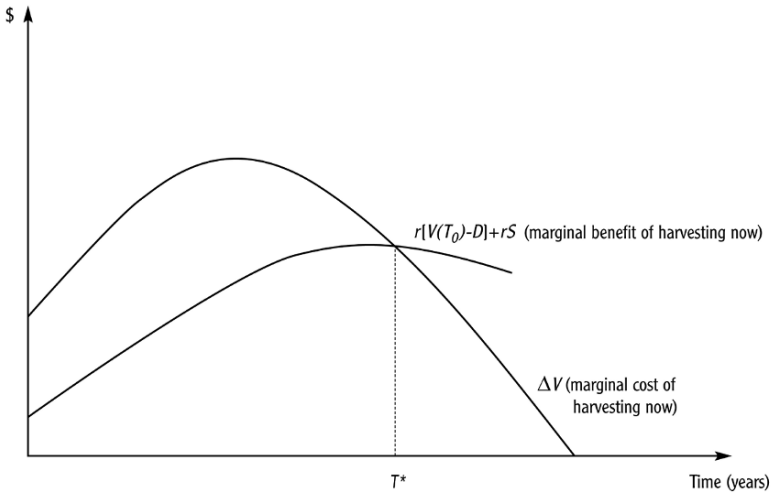


Figure 7.3 Efficient (Faustmann) forest rotation. The efficient rotation length, T^* , equates the marginal benefit of harvesting and the marginal cost.

optimal rotation problem once, she has solved it forever, given no change in the basic parameters. That is, T^* will remain the same in year 100 as it is in year 1. So in addition to delaying the net returns from the current harvest by 1 year, if she waits 1 extra year before cutting the trees she also delays the net returns from each future harvest by 1 year. And each of those delayed harvests has an associated cost: the forgone growth in the value of money in the bank from the harvest over the period of 1 year. The Faustmann model captures the present value of this change in the timing of all future harvests in S . The perpetually delayed harvests decrease the present value of expected future returns to the land (which before T^* is outweighed by the marginal increase in the present value of expected future returns that results from greater timber volume). And taking into account this future loss from delaying the harvest shortens the optimal infinite rotation, in comparison to the single rotation. Although this characteristic of the efficient rotation is extremely important, it is also quite a subtle point. It is so subtle, in fact,

The Faustmann Rule identifies the dynamically efficient forest rotation, maximizing the present value of future net benefits. It takes into account the time value of money and also site value, the opportunity cost of keeping the land in forest rather than converting it to another use, such as farming or residential development.

that even though German forester Martin Faustmann described the rule in an article published in 1849, the substance of his point went unnoticed by forest economists for more than a century.

Efficient Rotation with Nontimber Forest Benefits

Forests offer many benefits aside from the commercial value of their timber. Nontimber forest products and benefits include species habitat, watershed protection, carbon sequestration, and recreation. How do these nontimber forest values affect the efficient rotation?

When a forest provides multiple nontimber benefits, it can be hard to sort out their effects on the optimal rotation. Imagine, for example, that our landowner's forest stand provides habitat for both the red-cockaded woodpecker (a threatened U.S. species that prefers mature southern forests where trees reach 60 to 180 years of age) and the white-tailed deer (which prefers tender new growth and lots of understory on which to graze). These two uses pull in opposite directions: Benefits accrue to one species with a longer rotation, but this imposes costs on another species. To understand the general impact of nontimber forest benefits on efficient forest management, we will consider an example with benefits from old-growth standing trees. Keep in mind, however, that the problem is often multidimensional and therefore more complicated.

In figure 7.4 we introduce species habitat value—assumed to accrue only after the stand is sufficiently old—into our forest management problem.² We do this by incorporating the forgone benefits of woodpecker habitat as an additional cost of harvesting timber (because the benefits of harvesting are captured in the commercial value of the timber). Thus, the *social* marginal cost of harvesting timber includes both the forgone potential growth in timber volume from waiting 1 more year (V) and the forgone benefits of bird habitat (H).

The effect is to increase the optimal rotation period. If the landowner is managing her forest efficiently, she will now let the trees stand longer in each rotation than she would have in the absence of bird habitat. In figure 7.4, we represent this as an increase in the optimal rotation from T^* to T^*_{bird} . In fact, if the value of woodpecker habitat in old-growth forests is large enough, the social marginal cost of harvest curve may never intersect the curve that describes the marginal benefits of harvest. In other words, the optimal rotation may be infinite, meaning that it would be efficient to never harvest certain stands.

Later in this chapter, we describe an economic benefit valuation study of northern spotted owl preservation that used contingent valuation, one

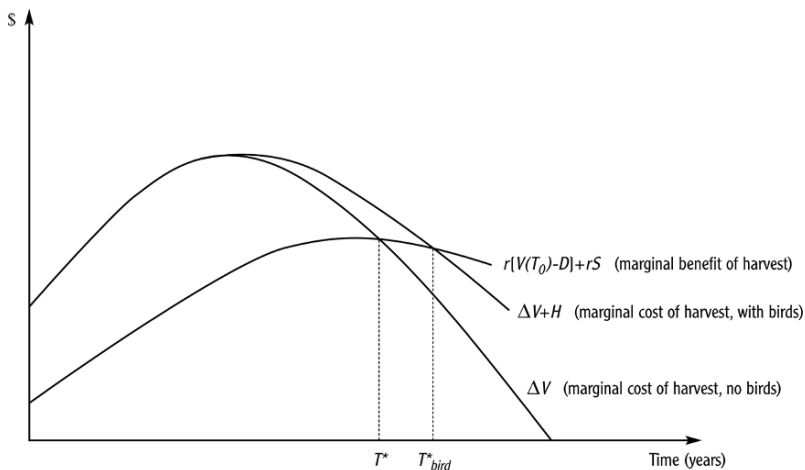


Figure 7.4 The effect of non-timber benefits on the Faustmann rotation. The value of bird habitat provided by old-growth forest represents an additional cost of harvesting timber. As a result, the efficient rotation length increases in this case, from T^* to T^*_{bird} .

of the benefit valuation techniques described in Chapter 3. Benefit valuation techniques can also be used to determine the effect of nontimber forest benefits on optimal forest rotation. For example, in the coastal forests of British Columbia and boreal forest of northern Alberta, Canada, the Faustmann rotation increases by approximately 20 percent, on average, when carbon sequestration benefits of standing trees are considered.³ Another analysis of a forested watershed in Victoria, Australia, makes a similar type of calculation.⁴ They account for the value of both watershed protection and carbon sequestration, finding that under some reasonable assumptions with respect to the value of water, carbon, and other factors, the particular forest studied by these researchers should never be harvested.

Public Goods, Property Rights, and Deforestation

The Food and Agriculture Organization (FAO) of the United Nations estimates that 13 million hectares of forest was lost each year between 2000 and 2010, a rate of loss offset only partially by new plantings, natural forest expansion, and establishment of forest plantations. The greatest net losses occurred (and continue to occur) in tropical regions. Given the richness in biodiversity of tropical forests, deforestation, particularly in the global South, has been an ongoing issue of concern to environmental advocates and policymakers. As northern countries have increasingly regulated the

Nontimber Forest Benefits: The Spotted Owl Controversy

In response to new biological information on the habitat needs of the northern spotted owl, the U.S. Forest Service was directed in the mid-1980s to issue a revised management plan for the Pacific Northwest Region. The “preferred management alternative” presented in this plan generated a large amount of media coverage and controversy. There were lengthy arguments for and against measures to set aside old-growth forests for owl habitat. By the late 1980s, there was open disagreement between federal agencies about the management plan, and environmental advocates were attempting to block timber sales from old-growth areas. Court rulings alternately halted and permitted timber sales, and most decisions were appealed by either the logging industry, environmental groups, or both. By March 1989, twenty-five timber mills in the region had shut down, causing many workers to lose their jobs. In 1990, the owl was listed as a threatened species under the Endangered Species Act.

These are difficult tradeoffs: jobs, threatened species dependent on old-growth forest, recreation, and cheap timber for construction and other purposes. A number of researchers strove to make these tradeoffs transparent in the form of benefit–cost analysis. One group of economists estimated the benefits to U.S. residents of preserving old-growth forests in the Pacific Northwest as habitat for the spotted owl.⁵ Their benefit estimate, determined through contingent valuation to be in the neighborhood of \$1.5 billion, exceeded the U.S. Forest Service’s estimate of owl preservation costs (including forgone timber harvests), which ranged from \$500 million to \$1.3 billion. Another study also showed total benefits of owl habitat protection greatly exceeding total costs (their most conservative estimated ratio of total benefits to total costs was 3:1, and the estimated ratio most favorable to preservation was 43:1).⁶

extraction of timber, especially from old-growth forests, partially because of greater recognition of nontimber forest values, the resulting restrictions in global timber supply and increases in global timber prices have increased the incentives for deforestation in tropical regions (which are less heavily regulated).

To the extent that the world’s forested lands are in private hands and markets are complete, some of this observed deforestation is actually efficient use of a scarce resource: land. In previous chapters, we argued that private owners of nonrenewable natural resources in competitive markets faced powerful incentives to maximize the present value of net benefits from their resources, such as oil wells and coal mines. In the same way, private landowners face powerful incentives to maximize the present value

of the net benefits they receive from their land. So, for example, farmers think carefully about the profit-maximizing mix of crops to grow in a given year, within certain constraints. In the long term, they also think carefully about whether farming itself is the most profitable use of their land or whether they should instead consider selling their land for other purposes. Farmers in developing countries, likewise, make similar decisions about converting forested land to agriculture, although the context is different in important ways.

Whereas nonrenewable resources lie, in large part, in private hands, forests exhibit a wide variety of ownership regimes. For example, in 2008 federal, state, and local governments owned approximately 44 percent of U.S. forested land, and 56 percent was privately owned (21 percent by the forest industry, conservation organizations, and Native American tribes and 35 percent by families and individuals).⁷ In developing countries, public ownership is more prevalent; about 83 percent, on average, in Africa, and 92 percent in South and Southeast Asia, for example.⁸ In addition, many nontimber forest benefits are public goods, so we would expect private landowners to rotate their trees at a rate that provides less than the efficient quantity of some of these services, such as species habitat, watershed protection, and carbon sequestration. These two factors—the nature of property rights and the prevalence of public goods among forest services—frame the discussion of deforestation from an economic perspective.

Well-defined property rights are particularly important for the efficient management of forest resources, because optimal rotation periods for some species can be very long. Imagine trying to ensure that a valuable hardwood left standing today will be there for the efficient forester in 50 years, in a country where the revenue from cutting the hardwood can feed a hungry family for months. Even where well-intentioned governments intervene in markets to establish property rights to forested lands, the incentive to poach trees in such countries is very high. Efficient forest management hinges critically on the ability to regulate the capture of forest resources.

Economists have provided empirical evidence about the influence of property rights on deforestation rates, and these studies provide mixed evidence about the influence of secure property rights on deforestation. Studies of the Amazon basin in Brazil have shown that possession of land title leads to longer rotation periods and increased efforts at reforestation and conservation by small landholders.⁹ Land title holders in Brazil are also less likely to participate in timber markets altogether; they are less

likely to sell trees for a living than to use forested lands for other purposes. In developing countries, there is a strong relationship between deforestation rates and factors (such as political instability) that indicate uncertainty over property rights.¹⁰

Economists have also shown the importance of protection and enforcement of property rights, in addition to their establishment.¹¹ In Brazil, although possession of land title reduces the incentive to deforest one's land, the positive effect of land title is eroded where title holders face a continued risk of fire contagion from neighbors clearing land for agriculture. Social institutions can mitigate this effect. For example, if landowners are aware of the timing of neighbors' fires, they can take preventive measures. Where such coordination occurs, title holders maintain longer rotations and better conservation and reforestation practices.¹²

Note, however, that increasing the security of property rights may also increase returns to land uses that compete with forests. For example, land titling reforms in Nicaragua increased deforestation, because the reforms provided an incentive to increase investments in agriculture, increasing returns to deforestation.¹³ Thus, policymakers and activists concerned with deforestation cannot rely on property rights approaches alone. Internalizing negative externalities from deforestation (or positive externalities from preservation or afforestation) through payments-for-ecosystem-services policies can be another efficient approach.

One common mechanism for preserving forest is public ownership and management of protected areas, which impose strong limits on extractive activities. (It may surprise you to learn that resource extraction, including timber harvest, is common in most publicly owned forestland worldwide, including in the United States.) Economists have raised four main issues regarding the effectiveness of protected areas.

First, when a protected area is established, private incentives to exploit the resources within its boundaries are not eliminated. Government protection may raise the cost of activities such as tree poaching, but effective enforcement is essential if protected areas are to do their job. Second, in order for protected areas to truly be "additional," they must set aside land that would otherwise be at risk for deforestation. Unfortunately, these at-risk lands are often most costly—politically and economically—for societies to set aside. Therefore, it is not surprising that assessments at the global and country levels suggest that many protected areas have been chosen specifically because they were unlikely to be suitable for exploitation, at least in the near term.¹⁴ These low-cost protected areas may also provide little benefit in terms of avoided deforestation. Third, economists

Paying Landowners to Preserve Tropical Forests

Starting in 2003, Mexico's national government began paying landowners to maintain forest cover, enrolling participants through 5-year contracts in a program originally called Federal Payments for Hydrological Services (and eventually called PROARBOL, then PRONAFOR). The primary goals are preserving forest for watershed protection and aquifer recharge, improving water quality and quantity for downstream communities, and reducing flood risk.¹⁵ The program has important rural poverty alleviation goals also and is one component of Mexico's national strategy for reducing carbon emissions (given the carbon sequestration benefits of avoided deforestation).

Water users pay federal fees that fund the program, which enrolled more than 2.6 million hectares between 2003 and 2011. The program uses a point system for selecting participants, based on risk of deforestation, local surface water scarcity, and location in an area of either high poverty or high indigenous population. Annual per-hectare payments are scaled according to the ecological value of the particular parcel; cloud forest (the preservation of which has a stronger impact on water quality and availability) receives a higher payment than other forest types. Satellite imagery and field inspections are used to monitor compliance. The Payments for Hydrological Services program reduced Mexico's forest cover loss by 40–51 percent between 2003 and 2012, compared with what would have taken place without the program, and also provided small poverty alleviation benefits during this period.¹⁶

have pointed out the potential for spillover effects of protected areas. If a government protects one patch of forest and through this action pushes land conversion into neighboring areas, with a net result that is the same as, or worse than, the counterfactual—what would have happened without the protected area—then these forest preservation policies may not have net benefits.¹⁷

Fourth, and finally, there are questions as to the economic impacts of protected areas on local populations, who may lose access to land, fuelwood, hunting or foraging areas, and other resources. On the other hand, locals may benefit from markets for ecosystem services, protection of water, fish, and other resources used locally outside a preserve, or even ecotourism revenues. The net local effect can be expected to vary with circumstances. The hypothesis that protected areas either increase or decrease poverty can be difficult to test empirically. Protected areas tend to be distant from cities, on difficult terrain, or in other ways negatively correlated with economic development. So it is necessary to control carefully

for these confounding factors. Recent work in Thailand and Costa Rica suggests that forest protected areas have decreased local poverty rates.¹⁸

Fisheries

Forests exhibit a range of property rights regimes, from private ownership to open access. Unless we include aquaculture in our analysis, the same cannot be said of fisheries. In fact, most fisheries are not owned by any party, and many are open access; thus, completeness of markets will be an integral part of our fish story. As in the case of forests, we will first take a look at a simplified version of fishery biology, and then add the economic dimension to determine the efficient quantity of effort to exert in harvesting fish.

Logistic Growth

To represent the biological side of our bioeconomic model of a fishery, we use a common model—the Schaefer logistic model for growth of a species population—that describes incremental growth in a fish stock as a function of the size of the stock (usually measured in tons of biomass).¹⁹ The general form of the logistic function is as follows, where X is the fish stock, r is the fish species' intrinsic growth rate, and K is environmental carrying capacity:

$$F(X) = rX \left(1 - \frac{X}{K} \right)$$

The logistic curve is symmetric and bell shaped (figure 7.5). To the left of the curve's peak, the annual growth in stock is increasing in the size of the stock (bigger stock leads to faster growth), although this is happening at a decreasing rate. When we reach the peak, the annual rate of growth is maximized (at X_M). On the right side of X_M , the rate of annual growth in stock is decreasing in stock size: More fish mean less growth, as the larger population begins to result in crowding and competition for food, for example. At the far right-hand side of the logistic curve, we reach the carrying capacity of this fishery, K . This is the fish population that would persist in the absence of any outside perturbation; mortality is exactly offset by new births.

Bioeconomic Model

Now we introduce fishing activity. Harvesting fish, like extracting oil, is fundamentally a dynamic problem. The population left in the fishery

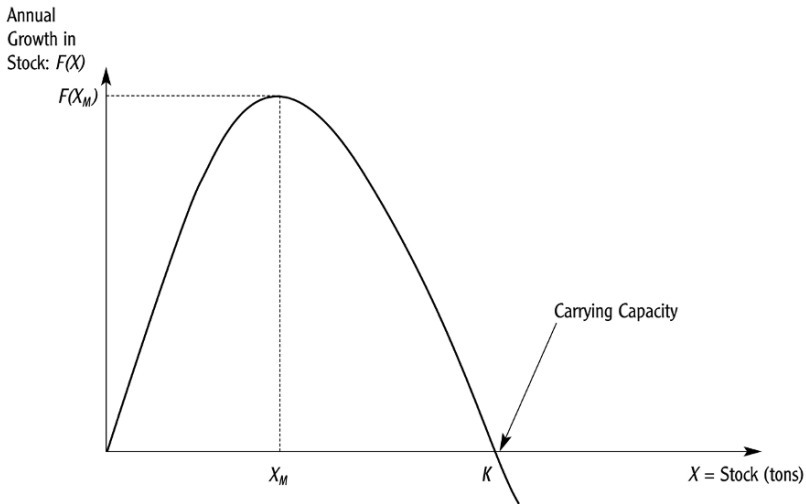


Figure 7.5 The logistic growth curve in the model of a fishery.

tomorrow will depend on what is happening today, such as the size of the fish stock, environmental conditions, and human fishing effort. Solving the full dynamic problem turns out to involve some complicated mathematics. To keep things simple, we'll focus instead on a *steady-state model*—that is, on what happens in a fishery in long-run equilibrium. This approach will still give us a great deal of insight into the economics of fishing and the inefficiency of open access.²⁰

In the steady state, the fish population remains the same from period to period, thus the fishing harvest is equal to net growth of the stock. Notice that any harvest level would be sustainable if the harvest rate were equal to the growth rate of the fish stock. Such a harvest could be sustained forever, and the population size would remain constant. For any level of the stock (X), $F(X)$ is equal to the rate of annual growth in the stock, and it is also equal to the maximum sustainable long-run level of harvest for that stock. Thus in the steady state, the logistic curve is also a sustainable yield curve.

Looking at figure 7.5, we can expect to maintain any population size along the horizontal axis between zero and K simply by harvesting the number of fish each year equal to the natural change in population from the last year. For example, X_M is known in biological fishery models as the stock that results in *maximum sustainable yield*. A stock of this size maximizes the average level of growth; therefore, it also maximizes the sustainable (non-stock-reducing) yield, that is, the largest catch that can

be perpetually maintained, $F(X_M)$. If the fishing harvest were to reduce the stock beyond this point, the fishery would be biologically overfished.

To understand the relationship between fishing effort and the returns to fishing in the steady state, we need to make an important assumption. We assume that the yield per unit of fishing effort is proportional to the size of the fish stock: More fish mean a greater return per unit of fishing effort.²¹ Given this assumption, the relationship between the logistic growth curve we have already drawn and a curve describing the steady-state returns to fishing as a function of the level of effort is quite straightforward; they look very much alike.

In figure 7.6, we graph the yield–effort function $Y(E) = 10E - E^2$. We represent fishing effort (any constant unit will do; here we use the number of boats) on the horizontal axis. With the level of fishing effort increasing from left to right, the fish stock is increasing from right to left, simply because fishing effort reduces the stock. Thus, at the origin, effort is equal to zero and the stock is equal to K , carrying capacity or natural equilibrium, the level of stock that will prevail in the absence of fishing. The vertical axis in figure 7.6 measures the returns to fishing effort. We could continue to measure returns in tons of fish, but we find it more useful to proceed in dollars. To convert fish to dollars, we need only to know the price of fish; here, to keep things simple, we assume it to be \$1 per ton. The hill-shaped curve in figure 7.6 gives us the total revenues we can expect from fishing at varying levels of fishing effort in the long run. Notice that after the curve peaks, effort continues to increase while returns decline. We have also drawn in the total cost curve, assuming that the marginal cost of fishing effort is \$3 per boat.

Efficient Fishery Management versus Open Access

In biological fishery models, the stock that maximizes the average level of growth also maximizes the sustainable (non-stock-reducing) yield, that is, the largest catch that can be perpetually maintained. If the fishing harvest were to reduce the stock beyond this point, the fishery would be biologically overfished.

To decide how much effort we should put into this fishery, we compare benefits and costs. To maximize the value of this resource to society, the fishing effort should maximize the rents from fishing (the difference between total benefits and total costs). Rents in the fishing model are like scarcity rent in the nonrenewables model; in both cases, rents are what we obtain from exploiting a scarce resource above and beyond




Figure 7.6 Efficiency versus open access. The efficient level of fishing effort, E^* , sets marginal cost equal to marginal benefit. In the open-access equilibrium, total costs and benefits are equal, resulting in a much higher level of effort, E^{OA} .

the cost of extraction or harvest. We can locate the rent-maximizing level of effort using figure 7.6; it is the point along the horizontal axis at which the distance between the total benefits curve and total cost curve is maximized (E^* in figure 7.6).²² In our example, this occurs when four boats are employed in the fishery. Fishing effort beyond this point is *economic* overfishing. Notice that the fishery is economically overfished before it is biologically overfished. This will always be true because fishing is costly. If the marginal cost of fishing effort were zero, the efficient level of effort would be equal to the level of effort that would preserve a stock that ensures maximum sustainable yield.

Is the level of fishing effort that we observe in real-world fisheries economically efficient? By now, you have probably guessed that it is not. We have discussed this problem as if fisheries had walls and gates and were either privately owned or staffed by a social planner who had the power to limit the number of boats admitted. Although there are privately owned fish farms in many countries, deep

Fishing effort beyond the point at which marginal benefits are equal to marginal costs is economic overfishing.

ocean fisheries (those beyond individual nations' 200-mile exclusive economic zones) are characterized by open access.

Under open access, boats will continue to enter a fishery as long as there are profits to be made (as long as rent is positive). At the efficient level of fishing effort identified in figure 7.6, profits still remain for other boats. This will be true as long as the total benefits curve lies above the total cost curve—that is, past the efficient number of boats, past the number of boats that result in the fishery's maximum sustainable yield and total returns begin to decline, to the point at which total benefits are exactly equal to total costs. In our model, in the absence of private ownership or regulation, fishers will continue to enter until there are eight boats on the fishery, twice as many as the efficient number (E^{OA} in figure 7.6).

At the open-access equilibrium level of fishing, the rents to fishing are completely dissipated. At first glance, the absence of such rents might seem to be the result of healthy competition in a well-functioning market. But that overlooks a crucial aspect of this problem. When eight boats enter this model fishery, they generate sufficient returns to cover the costs of their harvest, including things such as fuel, gear, and depreciation of fishing vessels. But there is nothing left to cover the depletion of an important resource—the fish!

Here again, actors in a market generate costs (fish depletion) that are external to the transactions between buyers and sellers. No individual fisher has an incentive to account for the externality of depleting the fish stock. Like our conscientious farmer who leaves a gallon of water in a nonrenewable groundwater aquifer, hoping to save it for tomorrow, the conscientious fisher who leaves a succulent tuna in the sea for tomorrow will quickly be put out of business by her competitors, who indulge in no such collectively minded behavior. The result is an inefficient race to fish, with too many boats pursuing too few fish. This depletion of the resource resulting from a conflict between collectively beneficial and self-interested behavior is a classic example of the tragedy of the commons that we discussed in Chapter 5.

Common Property Is Not Open Access

In light of the preceding discussion, it may seem as though we are implying that all shared resources, from fisheries to groundwater aquifers, are doomed to inefficient overexploitation. This is not the case, however. Although the terms *open access* and *common property* are often used interchangeably to describe shared resources, they are really very different property rights regimes. Open-access resources, by definition, lack

any restriction on who can use the resource or how much they can extract. In contrast, common property resources, though shared among a group (for example, a village), may be governed by formal or informal institutions, ranging from explicit rules to informal social norms.²³ Importantly, many common property arrangements steer clear of the extreme overexploitation typical of open-access resources.

Open-access resources lack any restriction on who can use the resource or how much they can extract. In contrast, common property resources, though shared among a group (for example, a village), may be governed by formal or informal institutions.

To get a feel for the difference between open access and common property, let us consider two very different social situations. In the first, you go out to dinner at a local restaurant, where rather than charging each diner the precise cost of her meal, the manager has decided to split the evening's total billings evenly among all of the guests in the dining room. You will pay a bill that amounts to the average of the entire restaurant's tab. How will your ordering behavior differ from the standard situation in which you pay for your own food and drink? Most of us would be more likely to order a second glass of wine, or coffee and dessert, when splitting the bill with a crowd of strangers.

In the second scenario, you go out to dinner at the same restaurant with a small group of close friends. You agree to split the bill evenly. Do you order a second glass of wine, as you did when you were splitting the tab with a large crowd of strangers? You might be somewhat less constrained in your ordering than you would be if you were going to pay the full cost of your order yourself. But you are likely to be somewhat more constrained than you were in the first scenario. After all, your friends might think very poorly of you if you stuffed yourself at their expense.

The first scenario resembles open access, the second common property. Successful common property arrangements rely on the self-regulating (and self-enforcing) capacity of resource users, who internalize incentives not to act opportunistically. Individual interest can be constrained by formal or informal institutions, like the social norms observed by a group of friends in a restaurant.

There is an intuitive link here to the Coase Theorem, which we will address in Chapter 8. By limiting harvests, fishers can increase the total returns from a shared fishery even as they exert less effort. The simple fact that potential rents exist and will be dissipated in the absence of some

kind of disciplined use may provide a powerful incentive for imposing such discipline. Of course, even successful common property arrangements, though they may avoid the degree of depletion observed under open access, may not result in efficient use. For example, a common property fishery may be economically but not biologically overfished.

Caveats Regarding the Steady-State Approach

Before we use the intuition from this simple model to analyze what we see happening in real-world fisheries, we pause to discuss three important details that are lost in the translation from a full dynamic bioeconomic fishing model to the steady-state model we have just described.²⁴ First, the steady-state analysis does not account for the time value of money. In a fully dynamic model, we would discount future costs and benefits of fishing effort, thereby arriving at an efficient level of fishing effort that is somewhat higher than what we identified using the steady state but still lower than the open-access equilibrium. Second, in the dynamic analysis it becomes clear that even the open-access equilibrium may not be sustained in some fisheries. In these cases, fisheries may be so intensely

Open Access versus Common Property: Beaver Hunting in James Bay

An example of the difference between open-access and common property regimes in the realm of natural resource management is that of James Bay, Quebec, where hunters traditionally have used resources communally.²⁵ The local Native American peoples have a rich heritage of customary laws regulating beaver hunting. Beaver in the region are an important food species and, since the start of the fur trade in the region in the late 1600s, a commercial species, as well. Beaver are vulnerable to depletion because colonies are easily spotted. Historically, a common property arrangement with senior hunters and their families serving as stewards of specific territories ensured sustainable use.

In the 1920s, a large influx of outsiders arrived in response to high fur prices. Native communities lost control over traditional territories, and a race to hunt ensued. All trappers (native and nonnative) contributed to the resulting tragedy. Beaver populations reached an all-time low in 1930, and conservation laws were enacted, banning outsiders from trapping. The traditional native family territories were recognized, and customary laws became enforceable. The return to a common property arrangement generated productive harvests again after about 1950.

exploited that the stock may collapse entirely. Third, without going through the dynamic model, it is impossible to demonstrate the strong parallels between the choice of a dynamically efficient harvesting policy and the choices of the optimal extraction rate of a nonrenewable resource, or the optimal rotation length for a forest. But most of the intuition we developed in those cases does carry over to the fishery problem. In particular, efficient fishery management has a strong Hotelling-like feel, in which at the dynamically efficient level of effort, the rate of return to the resource (fish in the sea) must equal the interest rate.

Economics and Real-World Fisheries

According to the FAO, of the major marine stocks or species groups for which information is available, about 29 percent are biologically overexploited, meaning that they are fished beyond the point at which the stock is able to maintain maximum sustainable yield.²⁶ The collapse of the cod, flounder, and haddock fisheries on the Grand Banks off the coasts of New England and Canada are one commonly cited example. In the late 1400s cod were so abundant that they could be pulled from these waters in weighted baskets.²⁷ Groundfish landings peaked in the northwestern Atlantic in 1965 at 2.6 million metric tons; despite significant improvements in fishing technology and substantially increased effort, annual landings from 1994 to 1997 averaged about 121,000 metric tons.²⁸

Can we draw any direct links between the inefficiency of open access in theory and the numerous examples of fishery collapse in the real world? A number of studies have attempted to estimate the efficient level of fishing effort in various fisheries and to compare actual rates of effort with efficient rates. A regional study of a fishery in the Bering Sea and Aleutian Islands concluded that the optimal level of effort would allow 24 factory trawlers and 44–50 catcher vessels in the fishery; 140 vessels were operating in the fishery at the time.²⁹ In the U.S. Gulf of Mexico shrimp fishery, the number of fishing vessels more than doubled between 1965 and 1988, while average landings and revenues per vessel declined; in 1988, the fishery's shrimp harvest could have been caught by one third of that year's shrimp trawling fleet.³⁰ Other U.S. fisheries that have experienced near-total depletion as a result of open access and resulting overcapitalization include those for Pacific halibut, Gulf of Mexico red snapper, northern California sardine, Atlantic Ocean perch, and West Florida sponge. In Chapters 9 and 10, we will discuss an important market-based approach to regulating fisheries experiencing these problems: individual tradable fishing quotas.

Other Important Considerations: Subsidies and Externalities

We have two remaining issues to explore in our discussion of fishery economics: subsidies and bycatch. Fishing industry subsidies can worsen the open-access problem and accelerate fishery collapse. Subsidies are a widespread attempt of governments to support those employed in fisheries that have become economically marginal through open access or other forces. They can take many forms, including

- Direct income support
- Price supports
- Reductions in the marginal cost of fishing effort, such as fuel tax exemptions
- Subsidies for capital equipment, including low-interest loans and loan guarantees
- Inefficiently low (or even zero) charges for fishing in public waters
- Subsidies to shipbuilding, ports, and fish processing facilities

Where subsidies are in place, even in the absence of open access, the equilibrium level of effort in a fishery will be higher (and stocks will be smaller) than the efficient level. In the early 1990s, the FAO estimated that the world's 3 million fishing vessels had \$92.2 billion in annual operating costs (excluding vessel depreciation, debt service, and return on investment) but brought in only \$70 billion in gross revenues.³¹ The shortfall was made up by government subsidies. A more recent study estimates that global governments spend \$30 to \$34 billion per year subsidizing their fishing industries.³²

To understand the impact of subsidies on the equilibrium level of fishing effort, let us take, for example, a specific type of subsidy: a fuel tax exemption, which drives down the marginal cost of fishing effort. In figure 7.6, a decrease in the marginal cost of fishing effort would pivot the total cost curve downward from its stationary point at the origin toward the horizontal axis; its slope would be less steep. As a result, the open-access equilibrium level of fishing effort (where total costs intersect total benefits in figure 7.6) would be pushed even further to the right. Because we have chosen to measure effort in terms of the number of boats, in our example additional boats will enter the fishery in the presence of a subsidy, resulting in further depletion of the stock. Keep in mind, however, that subsidies reduce fish stocks in two ways: They lower the cost of fishing for boats currently operating in a fishery (potentially increasing current boats'

number of fishing days), and they make fishing profitable for previously marginal boats (the effect described earlier).

The second remaining issue in our discussion of fish has to do with externalities. We have already discussed an important externality: the “depletion externality” that is present in open-access fisheries. Another externality associated with fishing is that of bycatch, the unintended capture of nontarget fish and marine mammals. Bycatch imposes a cost on society—the loss of these innocent bystanders—that in the absence of regulation is external to the decisions of individual fishing boats. Thus, where bycatch occurs, the equilibrium level of fishing effort (even in the absence of open access) will be higher than the efficient level. One example is the dolphin–tuna controversy that erupted in the late 1980s and resulted in the labeling of “dolphin-safe” tuna. Between 1960 and 1972, an average of one hundred thousand dolphins were killed each year in bycatch incidents by the U.S. tuna fleet alone.³³ Regulation in the form of the Marine Mammal Protection Act required U.S. fishers to take measures to decrease dolphin mortality beginning in 1975; mortality levels in 2000 were approximately five thousand per year.³⁴

Conclusion

The models we have discussed in this chapter and the previous one are members of a class of problems regarding the economics of natural resource management. As we have seen, the management of a natural resource, whether nonrenewable (like oil) or renewable (like forests or fish), is inherently a dynamic problem. Economic analysis provides a framework for decisions about the allocation of scarce natural resources over time, just as it provides a framework to analyze the allocation of other scarce resources among firms and consumers. We have discussed natural resource stocks as capital assets, goods that provide returns over time, either in their natural state (growing trees, oil in the ground) or upon extraction or harvest.

The bottom line is that efficient management requires that extractors and consumers of natural resources face the right prices to enjoy these activities. In some cases, markets acting alone will get prices right. Private owners of oil reserves in competitive markets, for example, can be expected to take scarcity into account in deciding how to allocate their extraction over time. In other cases, markets require some type of intervention to ensure efficient prices. If forests generate public goods, such as species habitat, private owners of forested lands may underestimate the cost of the timber harvest and therefore harvest trees too quickly. If the

cost of harvesting a fish from a deep ocean fishery includes the labor and capital used to capture it but not the value of the diminished fish stock, fisheries will be depleted too rapidly.

Later in the book, we will come back to the question of getting prices right in markets for natural resources, ensuring that all the externalities associated with resource extraction and consumption are reflected in market prices. In Chapter 8, we will discuss government policies that can mitigate the open-access problem and other externalities. And in Chapters 9 and 10, we will see how some of these policies have been implemented in the real world.

8

Principles of Market-Based Environmental Policy

Is government intervention in the environmental arena needed at all? After reading the earlier chapters, you may think that the answer is self-evident. After all, in the absence of government policies we have seen that private firms and individuals may impose negative externalities on other members of society and will fail to provide efficient amounts of public goods.

Nonetheless, a strong argument can be made that individuals or firms—at least in some cases—can solve externalities on their own, through private bargaining. This argument, due to Ronald Coase, is the starting point for this chapter; it offers a presumption against government intervention that we must consider carefully. After we have satisfied ourselves that government still has a vital role to play in addressing environmental problems, we go on to discuss the various types of regulatory policies (or policy instruments) that governments can use. These include prescriptive regulations that mandate certain actions at the level of individual firms, as well as more flexible market-driven approaches. As we shall see, these market-based policies—which include taxing emissions and creating markets in pollution—are appealing from an economic perspective, because they directly address the market failure at the root of environmental problems. Although a main theme running through our discussion is the common logic underlying emissions taxes and cap-and-trade programs, we will close the chapter by considering when it makes sense to use one rather than another on the grounds of economic efficiency.

The Coase Theorem

We opened this chapter by asking whether government intervention was necessary to correct negative externalities such as air pollution. This topic is the basis of a famous debate in economics, one whose central protagonists faced off not in person but in print. On one side was Arthur Pigou, author of a classic treatise on public goods, *Economics of Welfare*, first published in 1920. Pigou used economic theory to argue for government intervention in the economy—a sharp departure from Adam Smith’s invisible hand but one that soon became conventional wisdom among economists. Consider (wrote Pigou) a railroad running through a woodland. Sparks from the railroad threaten to start a fire along the tracks, destroying the woods. However, the railroad will ignore this adverse effect and run as many trains as will maximize its profit. Pigou argued that this sort of external effect called for government action. To solve it, he suggested making the railroad liable for damages—either by requiring the railroad to compensate the landowner or simply by levying a tax.

Four decades later, Ronald Coase challenged this view—first in passing in an analysis of radio and television broadcasting, and then head-on in a classic article titled “The Problem of Social Cost,” which helped him win a Nobel Prize in Economics many years later.¹ Coase argued that under certain conditions private bargaining between the railroad and the landowner would result in the same outcome—the same number and speed of trains, say—regardless of whether the railroad is liable for damages or not.

How might private bargaining overcome negative externalities? Coase also used the example of a rancher raising cattle next to a farmer growing crops to demonstrate how this might work. Without fences to protect the crops or to contain the cattle, the livestock would stray and eat the crops. The Pigouvian remedy would be to tax the rancher for the damage. In that case, the rancher would pare back his herd or build a fence in whatever way maximized the gains from cattle raising, net of the damages caused. This would evidently be the efficient outcome. Coase’s key insight was that the identical outcome would arise without government intervention. If the rancher is permitted let his cattle roam with impunity, then the farmer will build a fence (or pay for a reduction in the herd) if and only if the cost of doing so is less than the avoided damage to the crops. Left to their own devices, therefore, the farmer and the rancher would still reach the efficient outcome.

In the jargon of economics, the Coase Theorem states that the allocation of property rights (i.e., whether the rancher has a property right to let his cows roam as they please, or the farmer has a right to a cow-free cropfield) has no bearing on how economic resources are used (i.e., whether the fence is built). Of course, how those property rights are allocated matters to the farmer

and the rancher, because it affects the distribution of income and wealth, but it does not matter from the point of view of society as a whole.

So striking was Coase's argument that when he originally proposed it, even the economists at the University of Chicago—the citadel of free-market economics—thought that he must have erred. They invited him to Chicago to make his case at a dinner party hosted by Aaron Director, chair of the formidable Chicago economics faculty. As recounted later by George Stigler (one of three eventual Nobel laureates in attendance), "In the course of two hours of argument the vote went from twenty against and one for Coase to twenty-one for Coase." By dint of persuasion, Coase had convinced his colleagues that the case for government intervention to solve negative externalities was weaker than Pigou had argued. If government actions do not promote efficiency, then why should government get involved at all? After all, government policy is costly and may have unintended consequences that make matters worse rather than better.

Hidden in our discussion so far, however, has been a crucial assumption: that bargaining is easy and inexpensive and that deals are easy to enforce. In the case of neighboring landowners, that is plausible enough. But what of the case in which soot from a factory settles over an entire town? Finding all affected individuals will be difficult. Determining the true damages will be nearly impossible (because if the factory is to pay compensation, individuals will have strong incentives to overstate their damages). And in the case where the polluter is not liable (so that efficiency might require the individuals to pool their funds and pay the factory to install pollution control equipment), the individuals themselves will face a classic collective action problem: Each person, realizing that their contributions benefit everyone else, will seek to free ride on the efforts of others. In such a case, the costs to the town and factory of reaching

The Coase Theorem states that private bargaining will result in the efficient resolution of negative externalities, without the need for government intervention, as long as property rights are fully allocated (but regardless of the distribution of property rights among affected parties).

and enforcing a bargain—what economists call *transaction costs*—are likely to be insurmountable. In this case, the assignment of property rights will matter for efficiency after all. If the factory is liable for its pollution (effectively giving the townspeople the “property rights” to clean air), it may be required to reduce its emissions by curtailing its operations or installing abatement equipment, but if not (so that the factory has an effective “right to pollute”), the townspeople on their own are unlikely to be able to pay the factory to do those things. As a result, the assignment of property rights will affect how much pollution control is done, not just who pays for it. For the Coase Theorem to hold, therefore, transaction costs must be negligible.

But of course transaction costs are not negligible in the real world.² As in the preceding example, transaction costs will generally be large when a large number of people suffer the damages from an externality. They will also arise when large numbers of firms or individuals contribute to the problem, when causation is difficult to establish, when information about damages is not widespread, or when firms or individuals act strategically in bargaining situations. In other words, transaction costs are ubiquitous in the environmental realm.

Why, then, should we bother with the Coase Theorem at all? One reason can be explained by way of analogy with physics. Much of classical mechanics takes place in a vacuum. “But everyone knows that we don’t live in a vacuum,” you might imagine someone saying. “How can this be relevant to the real world?” The answer is that knowing what would happen in a vacuum is a useful starting point for understanding what happens in the real world when you start to incorporate friction, air resistance, and so on. In the same way, the Coase Theorem can be thought of as a kind of “economics in a vacuum” result. By describing what would happen if private bargaining worked smoothly, it allows us to appreciate the importance of taking real-world frictions into account.

A second reason for considering the Coase Theorem concerns what it can teach us about the design of environmental policies. The ubiquitous presence of transaction costs provides a strong justification for government regulation, running counter to Coase’s faith in private bargaining. Nonetheless, as we shall see, Coase’s key insight—that the clear assignment of property rights may help resolve negative externalities—helped inspire the cap-and-trade policies that have been very successful in reducing pollution.

Finally, in some cases of real-world interest, the number of parties is small enough and the private incentives are large enough that Coasean

bargains can be struck, after all. The accompanying text box discusses just such a case that took place in the French countryside.³

The Array of Policy Instruments

Despite Coase's concerns, many economists do see a role for government in helping to solve environmental problems. The question is, How should it do so?

Let's start by considering the primary ways government might get involved. We use pollution regulation as a motivating example in this chapter. In Chapters 9 and 10, we will broaden our perspective to include

Perrier and Vittel: Paying Farmers to Change Their Agricultural Practices

Perrier and Vittel are two leading mineral water companies based in France, now part of the Nestlé Waters Group (the largest bottled water company in the world by revenue). In the late 1980s, Vittel initiated a program to reduce water pollution in the source area feeding the springs that are the source of its bottled water.⁴ As part of this program, the company signed long-term contracts of up to 30 years with dairy farmers in the watershed. The farmers agreed to adopt less intensive farming methods in order to reduce agricultural runoff of herbicides and other pollutants. In return, Vittel paid each farmer roughly \$230 per hectare per year for 7 years, adding up to about \$155,000 for the average farm. The company also provided free technical assistance and paid for new farm equipment and construction. When Vittel purchased Perrier in 1992, it applied a similar model to the Perrier springs in southern France, where the program introduced organic farming methods on more than 500 hectares of vineyards and wheatfields.

Although sometimes cited as an example of an ecosystem service market, this case is more aptly described as a straightforward illustration of the Coase Theorem. Because Perrier and Vittel bottle the water and sell it, they capture sizable benefits from the improvement in water quality. It seems reasonable to assume that the costs of the program were much less than the costs of alternative methods of cleaning or filtering the water.

At the same time, there may well have been positive spillovers from the private transaction. To the extent that the change in agricultural practices improved water quality throughout the watershed (and not just in the water coming out of the privately owned springs), Perrier and Vittel provided a public good as a byproduct of their Coasean bargain.

a range of real-world examples of environmental policies in areas ranging from fishery management to water pricing to solid waste disposal.

Prescriptive Regulation: Technology and Performance Standards

A first set of policies, known as prescriptive regulation or command-and-control, focuses on regulating the behavior or performance of individual factories and power plants. This has been the conventional approach for most environmental regulation, at least until the late 1990s. A *technology standard* requires firms to use a particular pollution abatement technology. For example, the 1977 Clean Air Act Amendments in the United States required new electric power plants to install large scrubbers to remove sulfur dioxide from their flue gases. Alternatively, a government regulator might impose a ceiling on the air emissions (or water effluent) an individual firm can release. This approach, which allows polluters leeway in determining how to meet those emission ceilings, is known as a *performance standard* (or emission standard, in the case of air pollution). These may impose a ceiling on total emissions in a period (e.g., tons per year) or a maximum allowable *emission rate* (e.g., pounds of pollution per unit of fuel consumed or output produced). Although performance standards could vary between firms in theory, in practice regulators have typically established uniform standards. Finally, real-world regulations are often a hybrid of these two approaches, sometimes called technology-based performance standards. For example, the Clean Water Act in the United States requires individual sources of water pollution to meet effluent limitations that are based on the best practicable (or, variously, the “best available” or “best conventional”) technology.

Market-Based Policies: Emission Taxes and Allowance Trading

In contrast, another set of policies—called market-based (or sometimes incentive-based) instruments—incorporate market principles into government policies. Rather than focusing on the technology or performance of individual firms, these approaches are much more decentralized, focusing on aggregate or market-level outcomes such as total pollution.

Market-based instruments can be divided broadly into two categories, depending on whether they work by influencing prices or by limiting quantities. A prime example of the first approach, inspired by the work of Pigou, is an emission tax, set by the government and paid by polluters on each unit of emissions. (Other writers refer to the same thing as a fee or charge.) Such a tax puts a price on pollution, forcing emitters to recognize the social damages from production along with their private costs. In

the language of economics, the tax *internalizes* the costs of pollution and therefore eliminates the externality (or at least the economic inefficiency associated with it). Subsidies are another price-based approach. Indeed, a pollution tax and an abatement subsidy can be viewed as two sides of the same coin. Charging a firm \$10 for each ton of emissions creates the same incentive to cut pollution as paying the firm \$10 for each ton of abatement.⁵

The second main approach is known as allowance trading or cap and trade. The government first establishes a total allowable quantity of pollution (the cap) for a group of emitters—for example, the electric power sector, or all industrial sources of pollution. It then allocates allowances (also called tradable permits) to the regulated entities, with each allowance corresponding to one unit of pollution. (For example, in the real-world case of the sulfur dioxide tradable allowance program in the United States, each allowance corresponded to 1 ton of SO₂.) At the end of each compliance period (usually a year), emitters must submit one allowance for each unit of pollution they have emitted. Under a cap-and-trade system, firms that find it relatively expensive to reduce pollution will buy allowances from firms that can abate at lower cost. In this way, the total amount of pollution is fixed by regulation, but the *allocation* of that pollution among firms—and therefore the amount of abatement that any single firm must do—is left up to the market.

A number of variations on this basic theme are possible. The pollution allowances may be freely distributed to firms (on the basis of size, past output, or some other measure) or sold in regular auctions to raise government revenue. Firms that generate excess allowances in a given period can typically bank them for use in a later year, giving firms greater flexibility in how to comply with the regulation and encouraging firms to abate even more than required in the early years of the program. Cap-and-trade systems can allow firms to meet their compliance obligations by purchasing and submitting offset credits, generated by reductions achieved outside of the capped sector (whether in another sector of the economy that is not under the cap or in another country). And cap-and-trade programs can incorporate price floors and ceilings in order to provide greater certainty that the allowance price will not go above or below certain levels.

Information-Based Approaches

Emission taxes and cap-and-trade policies are the core of what economists call market-based instruments. On the periphery are two other

approaches to promoting environmental protection that share a similar ethos. The first such approach involves government regulations, often known as right-to-know laws, that require information provision by private firms. A well-known example is the Toxics Release Inventory. Since 1988, manufacturing facilities in the United States have been required by law to report their annual releases of each of more than 300 toxic chemicals. The Environmental Protection Agency (EPA) makes these data publicly available. (If you are curious, you can find out how much pollution is released in your neighborhood by going to www.scorecard.org.)

The second approach includes ecolabeling and certification programs, which provide consumers with information about how a product was manufactured. For example, you can buy produce from organic farms that do not use pesticides, coffee beans from farms that provide diverse bird habitat, paper made with high postconsumer recycled content or without fiber from old-growth trees, or “eco-friendly” household cleansers that substitute natural ingredients for toxics such as chlorine. Ecolabeling and certification programs aim to advertise and verify such eco-friendly claims. Unlike right-to-know laws, they are voluntary rather than mandatory; indeed, many ecolabeling programs are operated entirely by nongovernment organizations such as the Forest Stewardship Council and the Marine Stewardship Council, while certification is typically carried out by private firms such as Scientific Certification Systems.

Like emission taxes and allowance trading, but in contrast to conventional command-and-control regulation, these information-based approaches are decentralized. That is, they aim to influence the behavior of individual consumers and firms by changing the incentives they face and then letting them make their own decisions about how to respond rather than by requiring or proscribing certain activities. For this reason, they are sometimes lumped in with market-based instruments such as emission taxes and cap-and-trade policies.

However, there is an important difference. Rather than setting up new markets (allowance trading) or introducing price signals (emission taxes), these approaches work by providing information. As a result, the incentives faced by firms are created indirectly (mediated through the behavior of consumers and citizen groups) rather than being the direct result of government policy (as in the case of emission taxes and allowance trading). This distinction is critical, because as a result these information-based approaches address a different gap in the marketplace. Right-to-know laws and ecolabeling programs overcome the market failure from *asymmetric information*. That is, they tell citizens more about the activities of

factories near them and inform consumers about the characteristics of the products they buy. But we are still left with the underlying problem of externalities and public goods provision.

In sum, although information provision and ecolabeling can help promote environmental protection, their role is fundamentally different—and more limited in scope—than what can be accomplished through mandatory government policies (such as emission taxes or tradable fishing quotas) that directly shape the behavior of the firms and individuals with the greatest impact on the environment. Another class of policy instruments are policy “nudges” designed using the insights of the field of behavioral economics. For the rest of this chapter, and indeed the rest of the book, we will leave aside these decentralized and behavioral policies aimed at providing information and instead focus on market-based policies.⁶

How Market-Based Policies Can Overcome Market Failure

Economics tends to favor market mechanisms to restore market efficiency. To see why this is, it is useful to recall the discussion of market failure in Chapter 5. There, we presented three ways of framing market failure in the environmental arena: negative externalities, public goods, and the tragedy of the commons. Each of these ways of thinking about environmental problems points toward an approach to solving them. One natural solution is to get the prices right, by using government policies to make firms and individuals pay for the environmental damage they cause. Once the negative externalities are internalized in this way, they will be incorporated into the prices of goods and services, and market outcomes will again be efficient. We can also think of government policies as filling in the missing demand curve for environmental quality, surmounting the sorts of free-riding problems we talked about in the context of public goods. Or we can think of policies as establishing property rights over resources that had previously been open to all, thereby overcoming the tragedy of the commons.

Getting the Prices Right

In Chapter 4, we saw that in the absence of government intervention, firms will ignore the external costs of their actions (e.g., pollution resulting from producing steel) when making output decisions. The market outcome is not efficient because the true social marginal cost of production (which includes the marginal damage from pollution) is greater than the private marginal cost.

Environmental Conservation “Nudges”

In recent years, advances in the field known as behavioral economics have found their way into the popular consciousness; you may have read books such as *Nudge*, by Richard Thaler and Cass Sunstein, or at least noticed related newspaper articles and blogs. One important message of research in this field is that people can be influenced to act in environmentally friendly ways in response to simple behavioral signals. Research emphasizes the impact of a policy’s appeal to social norms, such as comparison with one’s peers, on household conservation behavior for both water and energy.

Let’s consider a couple of examples, starting with household energy efficiency. Home energy reports produced by the company Opower and used by more than eighty-five U.S. utilities for their residential customers include information on energy use relative to a sample of neighboring households (a social comparison), as well as information on how to reduce energy use. These messages are targeted specifically to individual households. The reports prompt small but significant reductions in household energy use, which persist over time with repeated messaging, and appear to result from both changes in household energy-consuming activities (perhaps resetting a thermostat or turning off lights when they leave a room) and, over the long run, replacement of lightbulbs, appliances, insulation, and other energy technologies with more efficient models.⁷

Similar effects of social comparisons have been demonstrated for household water use. In the suburbs of Atlanta, one experiment in partnership with a water utility sent letters to a random sample of households containing technical information about how to conserve water. A first group received only that technical message, and two other groups received additional information treatments: (1) a simple message about the importance of conserving water during a drought or (2) the drought message, plus a comparison of the household’s water use with their neighborhood’s average (the social comparison). All three groups used less water than households who received none of the messages. But the households receiving the social comparison letter reduced consumption the most.⁸ As in the Opower energy efficiency example, the effects of social comparison on water use in this experiment were small, but this work is representative of a growing body of research suggesting that behavioral nudges like these hold promise as cost-effective mechanisms for inducing environmentally beneficial behavior.

As a starting point for our discussion of market-based policies, we'll stick with the simple scenario we used in Chapter 4. In particular, let's assume that the amount of pollution is directly proportional to the amount of steel produced. Thus the only way to reduce pollution is to make less steel. Although this is not realistic, it makes the intuition behind the analysis easy to understand. (We'll consider a more general model of pollution control later.)

Figure 8.1 (which is almost identical to figure 5.1) illustrates this example. The dashed (lowest) line represents the marginal damages of pollution; the middle line is the supply curve, corresponding to the private marginal costs of production; and the upper line represents social marginal costs (the sum of marginal damages of pollution and private marginal costs). As you know already, the efficient point (Q^*) is where marginal benefit equals social marginal cost. The unregulated market outcome (Q_M) yields too much output and thus too much pollution, with the size of the inefficiency measured by the shaded deadweight loss triangle.

So far, so good. We now ask: Can a tax produce the efficient outcome? (Note that in this simple framework, a tax on steel and a tax on pollution are identical, because steel and pollution are assumed to be produced in the same proportions.) To answer this question, we need to think about

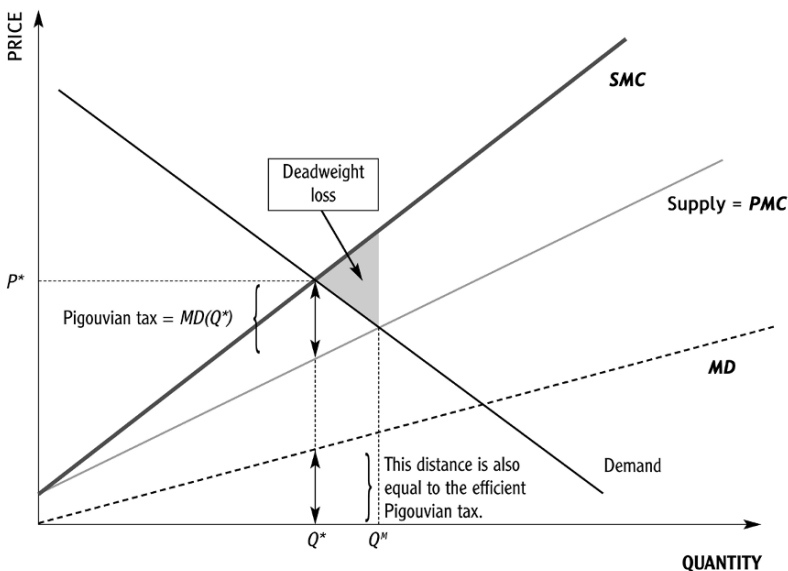


Figure 8.1 The efficient (Pigouvian) tax in the supply and demand framework. The tax equals the marginal damage from pollution at the efficient quantity, Q^* .

how a tax works. In market equilibrium without a tax, supply equals demand, implying that the price paid by consumers (along their demand curve) is the same price received by suppliers (along their supply curve). However, a tax introduces a gap between supply and demand. The price that the consumer pays is now higher than the price the supplier receives, by exactly the amount of the tax. To take a simple example, suppose that the government levies a 50-cent tax on every pack of cigarettes. If consumers pay \$4.50 for a pack (including the tax), the retailer receives only \$4; the government collects the difference.

On a graph, a tax can be represented by a vertical gap between the supply and demand curves. Market equilibrium, given a tax, is now the point at which the quantity supplied *given the price received by the supplier* equals the quantity demanded *given the price paid by the consumer*. Notice that as long as the supply curve slopes upward and the demand curve slopes downward, both sides of the market (producers and consumers) will share the burden of the tax. That is, the price to consumers will not rise by the full amount of the tax; instead, the price to consumers will rise by somewhat less, while the price to suppliers falls.

Now we can return to thinking about the case of a tax to correct a negative externality. Our goal is to set the tax in such a way that the resulting market equilibrium coincides with the efficient point, Q^* . How to do that? Well, we need to ensure that the tax is exactly equal to the gap (or vertical distance) between the demand and supply curves at Q^* . If we do so, the wedge created by the tax will be precisely large enough to drive the market to the efficient point.

But notice something important: The gap between the demand and supply curves at Q^* is exactly equal to the marginal damages from pollution at Q^* . This must be the case, because the demand curve intersects the social marginal cost curve at Q^* (that is the definition of the efficient point), and the difference between the supply curve and the social marginal cost curve is always equal to the marginal damages from pollution. In other words, the optimal tax is precisely equal to marginal damages at the efficient outcome.

Moreover, note that this is the only tax that will achieve the efficient outcome. In particular, a tax equal to the difference between the unregulated market price and the price paid by buyers in the efficient outcome (P^*) will not be large enough; that would result in a level of output somewhere between Q_M and Q^* .

This discussion leads to the following observation:

Getting the Prices Right by Removing Perverse Incentives: The Case of Fossil Fuel Subsidies

Our discussion in the text has focused on putting a price on pollution where none existed before—for example, through an emission tax or cap-and-trade program—in order to create incentives for firms and consumers to reduce pollution. In the real world, getting the prices right can also involve *removing* prices that create perverse incentives to pollute.

Consider fossil fuel subsidies, typically provided by governments (especially in oil- and gas-producing countries) to provide their citizens with cheap gasoline and diesel. The International Energy Agency (IEA) has estimated that subsidies on fossil fuel consumption amounted to \$523 billion annually worldwide in 2013. Countries with fossil fuel consumption subsidies spend an average of 5 percent of their gross domestic product on them; removing those subsidies would significantly improve those countries' fiscal outlooks and free up money for other government spending on education, health, and other priorities. By encouraging the consumption of fossil fuels, these subsidies also contribute to higher CO₂ emissions. The IEA found that even a partial phase-out of these subsidies in fossil fuel exporting countries could reduce emissions by 360 million tons annually, making fossil fuel subsidy elimination one of the most cost-effective near-term opportunities to address climate change.⁹

Partly as a result of similar analyses, in 2009 the G-20 group of large economies pledged to “phase out over the medium term inefficient fossil fuel subsidies that encourage wasteful consumption.” The results have been uneven, however, largely because of the political difficulty of eliminating subsidies that favor powerful industries (such as oil producers in the United States) or are intended to help consumers. Nonetheless, there have been a few bright spots. In 2014, building on years of attempts to eliminate fossil fuel subsidies in Indonesia, newly elected President Joko Widodo eliminated the subsidy on gasoline. And India, after steadily reducing diesel subsidies over several months, announced in late 2014 that it would deregulate diesel prices entirely.

- A tax on pollution equal to the marginal damage at the socially efficient level of pollution will achieve the socially efficient outcome.

The efficient tax is known as a Pigouvian tax, after the economist Pigou (whom we met in the first section). What such a tax does, in essence, is internalize the externality: It forces the producers and consumers of polluting goods to incorporate the full costs of their actions (including

The efficient (or Pigouvian) tax internalizes the externality, forcing producers and consumers to incorporate the full costs of their actions into their decisions.

the external costs from pollution and so on) into their output and consumption decisions. With the tax in place, the market outcome will be efficient: No other government intervention (such as telling the firms how much to produce) is necessary.

By correcting the market failure, the Pigouvian tax eliminates deadweight loss. You may have heard or read economists refer to taxes as distortionary. That is certainly true for most taxes (such as sales taxes or income taxes) that raise revenue but interfere with the smooth operation of the market by driving a wedge between supply and demand. In the case of the Pigouvian tax, however, the distortion arises from the *absence* of government intervention. The tax is *corrective*, not distortionary: It eliminates deadweight loss rather than introducing it. By incorporating the external cost into the price of the good, a tax gets the prices right and restores efficiency.

Filling in the Missing Demand Curve

We have just seen how a tax on output could restore the efficiency of markets, in the simple case where a unit of output always produces the same amount of pollution. But pollution and output are usually not so closely tied together. For example, an electric power plant can reduce its sulfur dioxide emissions by switching to lower-sulfur coal or installing a scrubber, without reducing its electricity production. In this more general case, the efficient tax would be on pollution directly, not on output, that is, on sulfur dioxide, not on electricity. Looking at the problem this way will also lead to another intuition behind market-based instruments.

Figure 8.2 plots abatement (rather than output of a good such as steel) on the horizontal axis. As in the graphs of marginal cost and benefit we drew in Chapter 2, the efficient level of abatement is marked X^* and coincides with the intersection of the marginal cost and marginal benefit curves. Looking at the figure now, however, you might notice something that may not have been apparent earlier: The marginal cost and benefit functions look an awful lot like supply and demand curves. Indeed, the resemblance is more than coincidental. Recall that the supply curve in a competitive industry is simply the marginal cost curve; therefore, we can think of the marginal cost of abatement as the supply of abatement. Similarly, the marginal benefits from abatement can be thought of as representing the demand for pollution control.

What Is the Optimal Gasoline Tax?

Although we have been discussing pollution taxes on firms (e.g., steel mills), we know from our earlier discussion (Chapter 5) that negative externalities can also result from individual actions, such as driving a car.¹⁰ You are probably familiar with the environmental consequences of automobiles: They are a leading source of emissions of particulate matter and nitrous oxides (which contribute to local air pollution) and carbon dioxide emissions (which help cause global warming). But driving also involves a host of other negative externalities, including traffic congestion and the costs of accidents (that is, the external costs that are not borne directly by the driver). What do these external costs add up to per gallon of gasoline?

Economists Ian Parry and Kenneth Small set out to answer this question and compute the optimal gasoline tax. They found that the Pigouvian tax reflecting marginal external costs would be 83 cents per gallon in the United States. They break that total down among the four externalities as follows: 6 cents for carbon dioxide emissions, 18 cents for local air pollution, 32 cents for congestion, and 27 cents for accidents. The small number for carbon dioxide emissions is surprising; it largely reflects the fact that there is less carbon in gasoline than one might think. Put another way, automobiles make a large contribution to global warming simply because of the sheer number of cars being driven and the amount of fuel consumed; on a per-gallon basis, the external costs due to global warming are fairly small. (The 6 cents per gallon estimate corresponds to marginal damages of \$25 per ton of carbon. Note that even doubling or quadrupling this number would still amount to a fairly small tax in cents per gallon terms.)

Parry and Small also point out that taxing gasoline is not necessarily the best possible policy. After all, the most expensive externalities—local air pollution, congestion, and accidents—depend on the number of miles driven rather than on the amount of gasoline consumed. Therefore, taxing gasoline is just an (imperfect) proxy for taxing miles driven. (This is a bit like taxing steel when we really care about pollution.) The authors estimate that replacing the gasoline tax with a mileage tax would increase social welfare substantially. Their argument illustrates the general principle that the choice of *what* to regulate may be just as important as (or more important than) the choice of *how* to regulate.




Figure 8.2 Market-based instruments to achieve efficient abatement. The efficient emissions tax equals the marginal benefit from abatement (as well as the marginal cost of abatement) at the efficient quantity X^* .

Our discussions of market efficiency can shed new light on the efficient level of pollution control—and on how to design policies that will achieve it. Suppose that individuals and firms harmed by pollution could be induced to pay for pollution control according to their true valuation. In that case, there would be a demand curve for pollution control, tracing out the social marginal benefits from abatement. In that imaginary scenario, a market for pollution control would arise naturally—just as markets arise for other goods. The interactions of buyers and sellers would determine the price and quantity of pollution control, corresponding to the intersection of supply and demand. Assuming perfect competition and complete information, this outcome would be efficient.

The central problem, of course, is that such a demand curve never arises, because pollution control is a classic case of a public good. Every individual, comparing the cost of paying for pollution control (borne entirely by herself) with the benefit (shared by others), finds that it is in her self-interest not to contribute. But when everyone free rides, the market demand for pollution control effectively falls to zero.¹¹ Note that the “supply curve” for abatement already exists; it simply corresponds to the marginal cost of controlling pollution. The hitch is that no firm will supply a good whose price is zero.

The role of government policy, then, can be understood as filling in the missing demand curve. Ideally, the government would reproduce the

whole downward-sloping marginal benefit curve (yielding a downward-sloping demand curve). To do so, however, the government would need to know the entire marginal benefit function and would have to be able to pay firms different amounts to reduce pollution depending on how much abatement was taking place.

Two simpler ways of filling in the missing demand curve are illustrated by the two dashed lines on figure 8.2. First, the government can require a fixed *quantity* of pollution control. (See the vertical dashed line in the figure.) This corresponds to a tradable allowance program with a cap on allowable pollution. In effect, a cap-and-trade policy is a commitment by the government to “buy” a certain amount of abatement from firms in the regulated industry.

Alternatively, the government might set a fixed *price* for pollution control. (This is represented by the horizontal dashed line in figure 8.2.) A tax on emissions is one way to set such a price: A firm saves the amount of the tax on every unit of pollution it abates. Recall from our discussion earlier that an emission tax creates the same incentives as an abatement subsidy, which is literally a commitment by the government to pay firms to reduce pollution. In effect, a tax on pollution amounts to charging firms for what they would pollute in the absence of regulation and then paying them back for every ton they abate.

A cap-and-trade system and a pollution tax, therefore, are complementary ways of filling in the missing demand curve—one by setting a quantity, the other by setting a price. Whether the government completes the market with a vertical “demand curve” (via a cap-and-trade policy) or a horizontal one (via a tax), it can achieve the efficient outcome. As you can see from the figure, the tax that achieves this level of abatement is the price that corresponds to the intersection of marginal benefit and marginal cost. This should not be surprising: It is the same thing we saw in the simple example of the steel market in a previous section. The efficient tax is the marginal damage of pollution at the efficient outcome.

Under a cap-and-trade system, the government can achieve a desired level of abatement directly by setting the appropriate cap. What about the price of pollution? In the cap-and-trade system, the price of pollution is simply the price of an allowance, which is determined by the market. It turns out (as figure 8.2 illustrates) that this price equals the tax that would achieve the same level of pollution.

Figure 8.3 explores this equivalence in a bit more detail. The marginal cost curve is the same as in the previous figure, as are the “demand curves” corresponding to an efficient tax and cap-and-trade program. Under a tax,

firms will abate up to the point where the marginal cost of abatement equals the tax T^* . By design, that level of abatement is the efficient quantity X^* . Abating more than this amount would cost more than paying the tax; abating less would mean paying the tax when it would be cheaper to reduce emissions. (This is the same logic that led us to conclude, back in Chapter 3, that competitive firms will produce up to the point where marginal cost equals the price.)

Meanwhile, under a cap-and-trade program, abatement is set at X^* by the cap. The resulting price of allowances must then be T^* : A higher price would lead to more abatement, and a lower price would lead to less. As a result, the price of an allowance in an efficient cap-and-trade system will be exactly equal to the efficient tax.

In theory, then, a tax and a cap-and-trade policy are essentially different ways of arriving at the same outcome. This equivalence also applies at the level of the individual firm. Because the two policies create the same incentives for reducing pollution, the allocation of pollution abatement among the firms in the industry will also be identical under the two approaches. That is, the amount of abatement due to any particular firm will be the same under the tax as under a cap-and-trade policy.

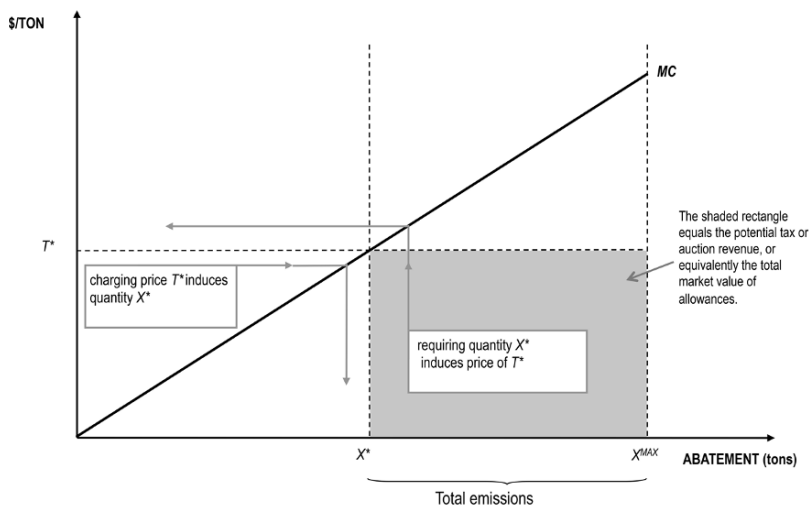


Figure 8.3 Equivalence between an emission tax (which sets the price on emissions T^*) and emission trading (which sets the quantity of abatement X^*), in the case of certain marginal costs of abatement.

A Tale of Two Trading Markets: Greenhouse Gas Emission Trading

As we will discuss in more detail in Chapter 10, emission trading has become a leading policy instrument to reduce greenhouse gas emissions (along with other approaches such as vehicle fuel economy standards and policies to promote electricity generation from renewable sources such as wind and solar power). Two of the most prominent greenhouse gas trading markets have been the European Union's Emissions Trading System (EU-ETS), which began a pilot phase in 2005 and entered full operation in 2008, and the Chicago Climate Exchange (CCX), launched in 2003.¹² The contrast between the two programs is striking, and it illustrates the important role played by government policy.

The EU-ETS is a government-run market, originally set up by the European Union (EU) to help meet its emission reduction obligations under the 1997 Kyoto Protocol. It is the largest emission trading program in the world, with an allowance cap of 2.3 billion metric tons in 2013, accounting for roughly 45 percent of total EU greenhouse gas emissions. The ETS covers more than eleven thousand facilities in thirty-one states (twenty-eight EU member states plus Norway, Liechtenstein, and Iceland), including electric power stations; manufacturing plants in sectors such as oil refining, iron and steel, ceramics, and pulp and paper; and civil aviation. Allowance prices were around €15 per metric ton for much of 2009 through 2011 before sliding below €5 per ton, largely as a result of the Great Recession; at the end of 2014 prices had recovered somewhat but were still below €10 per ton. In 2012, 7.9 billion metric tons were traded, with a gross value of €56 billion. In late 2014, the European Council agreed to tighten the EU's 2030 emission target and reaffirmed the central role of the ETS in achieving that goal.

CCX, on the other hand, was a purely voluntary exchange that ran from 2003 until 2010, with a membership of eighty-four companies (including DuPont, Ford, and Amtrak), municipalities, and universities, nearly all in the United States. Members committed to reducing emissions 4 percent by the end of 2006, relative to a baseline period of 1998–2001, and 6 percent by the end of 2010. A cumulative total of 150 million metric tons was traded over the 7-year life of the exchange. Prices rose from about \$1 per ton of CO₂ at the start of the program to about \$4 for much of 2006 and 2007, peaked at \$7.40 in June 2008, and then dropped to almost zero by October 2009. The exchange closed for good in November of the next year.

Why did the fortunes of the two markets diverge so much? The answer, of course, is that the EU-ETS is a mandatory government program, whereas CCX was purely voluntary. As we saw in Chapter 5, voluntary private provision of a public good is bound to be inefficiently low. Despite the great deal of hype

A Tale of Two Trading Markets: Greenhouse Gas Emission Trading *continued*

surrounding CCX, it had a miniscule impact relative to mandatory programs such as the EU-ETS. The contrast between the programs underscores the importance of the “cap” in “cap-and-trade.” Without the regulatory power of government to enforce a binding cap on all sources within a region or an industry sector, one cannot expect a voluntary program to accomplish much.

Creating Property Rights

As a final way of thinking about market-based policies, consider the problem of open-access resources that we discussed in Chapter 5 and again in Chapter 7. Unrestricted access to a resource typically results in overexploitation, as individuals act in their own self-interest rather than for the common good. How can market-based policies help in this case?

You have probably already guessed the answer. If we diagnose the problem as a lack of clear property rights, one promising remedy is to establish (and enforce) such property rights. Return for a moment to the shepherds in Hardin’s parable of the tragedy of the commons from Chapter 5. If the common pasture is divided between the shepherds, then each shepherd will have proper incentives to manage his own land wisely. Once private property rights are established, the market outcome will be efficient, because one shepherd’s stocking decision no longer affects the productivity of pastureland for everyone else.¹³

How would this be applied to natural resource management in the real world? In the case of a fishery, a property rights approach corresponds to what is known as *individual fishing quota* (IFQ) markets, also known as “catch shares.” Under such a policy, the total allowable harvest in a given year is divided up between a number of fishers. Each fisher receives some quota, which confers the right to take a certain percentage of the total catch. Fishers can then buy and sell their quota on an open market or lease them for a year. Such a policy is not simply a way of getting around open access, although it does achieve that goal. After all, traditional approaches to fishery regulation, such as setting a limit on the total allowable catch or restricting the length of the fishing season or the type of gear allowed, can also be seen as restricting entry or fishing effort, at least to some degree. The novel twist of the IFQ approach is that individual fishers receive *de facto* property rights in the resource, which they can trade

among themselves. This gives the fishers incentives not only to preserve the resource but also to harvest the resource in the most efficient manner possible.

This property rights approach might also remind you of the Coase Theorem that we discussed in the beginning of the chapter. Indeed, the economist who popularized the idea of pollution trading (J. H. Dales,

Property Rights for Halibut in the Gulf of Alaska

In the 1970s and 1980s, overfishing dramatically reduced the stock of Pacific halibut in the Gulf of Alaska. Regulators used command-and-control (CAC) approaches for two decades to try to stem the fishery's collapse, to no avail: The commercial halibut fishing season was restricted to 125 days in 1975, 25 days in 1980, and 2 days (24 hours in some areas) by 1994.¹⁴ What happened under this CAC approach? Fishers responded to shortened seasons through "effort substitution": They fished more intensively, for longer hours, and with more gear as the season shrank. During the 1994 season, crews fished for 48 hours straight, resulting in avoidable human injuries and deaths. Only frozen halibut was available in most of the United States for much of the year, while fresh fish decayed dockside in Alaska during extremely short seasons that overwhelmed local processing capacity and caused the price of halibut to plummet. This was an extreme case of the tragedy of the commons, but similar problems of overcapitalization and excessive effort exist to some extent in most fisheries managed using CAC policies. The "race to fish" in these areas diminishes rents, even driving them to zero.

Regulators put in place an IFQ system for Pacific halibut and other species in the Gulf of Alaska in 1995. In the first year of the policy, regulators were able to increase the fishing season from 2 days to 8 months, human mortality in the fishery was reduced to zero, and fish quality and availability increased. As we would expect from the discussion in Chapter 7, the number of fishing vessels dropped significantly, and the harvest increased. Fishery stakeholders who opposed IFQs when they were proposed soon supported the policy.

Recent research suggests that the phenomenon of IFQs—systems of property rights that establish ownership and encourage stewardship over marine resources—preventing fishery collapse is not limited to Pacific halibut. An analysis of catch statistics from more than eleven thousand global fisheries between 1950 and 2003 suggests that, had all non-IFQ global fisheries switched to management through tradable quotas in 1970, the percentage of collapsed fisheries by 2003 could have been reduced from more than 25 percent to about 9 percent.¹⁵

author of the 1968 book *Pollution, Property, and Prices*) made the connection explicit. Although pollution allowances are not private property in a legal sense, they do represent well-defined objects of trade whose value accrues to the holder. A market for tradable allowances effectively converts a nonrival, nonexcludable public good (clean air) into a collection of private goods (allowances). Because the allowances are excludable and rival (they cannot be shared or held in common), they can be bought and sold in a market just like any other private good.

The more general point is that the crux of commons problems is the lack of exclusion. Thus we can extend the intuition of property rights to cases in which we don't literally divide up a resource between individuals. Consider the example we gave of highway traffic in Chapter 5. Access to the highway can be restricted by erecting a tollbooth—transforming an open-access resource into what is effectively a privately operated one (where the owner, in this case, is the operator of the toll road). By raising the toll, the operator can reduce traffic and alleviate congestion. Indeed, the *efficient* toll is precisely equal to the external damages each additional driver imposes on everyone else—in other words, the Pigouvian tax. (This provides another example of the fundamental connections between our three ways of framing environmental market failures.) Tolls and other charges have been used successfully to reduce urban congestion in London (since 2003), Singapore (since 1975), and other areas.

In the case of managing a natural resource, the analog to charging a toll would be levying a “landing tax” on harvests. Just as a tollbooth excludes some people from using a highway, a landing tax on fish catches would in theory restrict the total harvest indirectly (through a price) rather than directly (through a system of property rights). However, although tradable quota markets have been used to manage fisheries (as we will see in detail in Chapter 10), landing taxes have not been implemented, presumably because of political opposition.

Raising Revenues

In focusing on how emission taxes and cap-and-trade programs can achieve efficient levels of abatement, we have ignored a major attraction of these policies for policymakers: their potential to raise government revenue. Under a tax, potential revenue is simply the tax times total emissions (see figure 8.3). Under a cap-and-trade program, the government can raise revenue by selling allowances at auction rather than giving them away for free; potential revenue is the market value of allowances, equal to the price of allowances times the quantity (equal to the cap). Because a

tax and cap-and-trade system can be designed to yield the same emissions and put the same price on pollution, they can also raise identical amounts of revenue. (By the same token, both types of policies can be designed to allocate some of the potential revenue to regulated firms: Just as allowances in a cap-and-trade program can be given away for free, an emission tax could be designed to exempt pollution levels up to some baseline level.) Far from being a fundamental difference between taxes and cap-and-trade, the potential to raise revenue is another example of the basic equivalence of the two types of policies.

Indeed, it is important to recognize that the amount of revenue raised by a tax or allowance auction is irrelevant to how well the policy deals with the negative externality. In correcting the market failure, what matters is that polluters have an incentive on the margin to reduce their emissions. Whether firms end up paying for every unit of pollution or instead receive some pollution for free (as with free allowances or a tax exemption) does not matter in terms of the incentives they face, as long as they face a price on the last unit of pollution they emit.

This might be surprising at first, but in fact it has been embedded in our argument all along. A firm's decision to abate 101 tons of pollution rather than 100 tons depends on the price of the 101st ton, not on the price of the first. You might think, "Well, if a firm is given 100 pollution allowances for free, then it doesn't have any incentive to cut pollution below 100 tons, since those emissions are all free." But now suppose that you are the manager of just such a firm, and suppose the market price of allowances is \$100. If you emit only 90 tons, rather than the 100 you are allowed, you can sell ten permits and receive revenues of \$1,000. Indeed, each ton of pollution you emit incurs an opportunity cost equal to the price of an allowance, whether or not you buy the allowance directly or receive it for free.

Nonetheless, we can still ask: From the point of view of society, should the government raise revenue from environmental policies? And if so, how should it spend the money? Until recently, economists thought that these questions didn't matter, at least in terms of efficiency. After all, we have just seen that the revenues raised by a policy have no bearing on the abatement incentives created by the policy.

The key is to widen our perspective and take into account other sectors of the economy. Quite apart from environmental regulation, governments raise needed funds by taxing income, capital investments, corporate profits, consumer purchases, and so on. These taxes are typically distortionary, undermining the smooth operation of the market: For example, an income

tax reduces the incentive to work, and a tax on capital affects investment decisions. Because income taxes hinder market efficiency, whereas environmental taxes (or cap-and-trade systems) increase it, a natural proposal is to use the revenue from emissions taxes (or allowance auctions) to fund reductions in distortionary taxes. That is, the government could cut the income tax and make up for the lost revenues with an emissions tax.

This is sometimes called a double dividend, because the pollution tax not only corrects the negative externality but also alleviates the distortions caused by taxes on income and capital.¹⁶ Revenues do matter for efficiency, after all. From an economic standpoint, environmental policies should raise revenue from polluters and use it to reduce distortionary taxes, such as those on labor income, sales, or capital gains.

Setting Prices versus Setting Quantities

As we pointed out earlier, there is a deep underlying equivalence between an emission tax and a cap-and-trade policy. One sets a price, the other a quantity, but as far as efficiency goes, the policies are just two ways of getting to the same point. The price of an allowance under an efficient cap-and-trade policy will be exactly equal to the efficient emission tax, and the abatement achieved by the tax will be the same as that imposed by the cap.

This theoretical equivalence is worth emphasizing because it underscores the common intuitions behind the two market-based policies. However, it relies on a key assumption: that the marginal abatement costs are known by the government, allowing the government to set the tax or cap to achieve the efficient outcome. What if the government lacks such precise information? It turns out that when marginal cost is uncertain, the choice of price versus quantity matters for efficiency. In particular, a tax (price) is preferable to a cap-and-trade policy (quantity) when the marginal benefit curve is flat relative to marginal cost, and vice versa. Surprisingly, however, uncertainty in the marginal benefit function does not matter for policy.

In theory, an emission tax and a cap-and-trade policy are just two ways of getting to the same point. In the real world, however, the choice does matter for efficiency.

First, let's consider what happens when marginal cost is uncertain. Suppose that the regulator knows what the marginal cost curve will be on average but not in any particular case at the time the policy is chosen. The actual cost could turn out to be above or below this average or

expected value. Figure 8.4 illustrates this case, depicting a high and low marginal cost curve, along with the average curve (labeled *EMC* for *expected marginal cost*) and the marginal benefit curve, which we assume is known.¹⁷

This uncertainty can be understood in two ways. One interpretation is that marginal costs are unknown at the time the regulation must be determined and that the regulation is “sticky”—that is, it cannot be changed—over some period of time during which the marginal costs will become known. Alternatively, you can imagine that the regulated firms know their own marginal abatement costs all along but for strategic reasons are unwilling to disclose them to the regulator.

What should the regulator do? Let’s consider a quantity instrument first—that is, a cap-and-trade policy. Because the regulator doesn’t know the true marginal cost curve, she can’t simply set the allowable pollution equal to its true efficient level. Instead, the best she can achieve is the




Figure 8.4 Comparison of price (emissions tax) and quantity (allowance trading) instruments under marginal cost uncertainty. The solid marginal cost line, denoted MC_H , represents the actual high marginal abatement cost curve, which is unknown to the regulator in advance. The bottom dashed line parallel to it, denoted MC_L , shows the alternative possibility (equally likely ahead of time) of a low marginal abatement cost curve. The middle line (EMC) is the expected marginal cost curve. The figure depicts a case in which marginal benefit is steeper than marginal cost, hence the cap-and-trade policy is preferable (smaller deadweight loss).

outcome that is expected to be efficient—that is, the level of abatement that equates marginal benefits with expected marginal costs. We have labeled this X^* on the figure. Despite her best efforts, this will result in some inefficiency. Because the true marginal cost curve will be different from the expected value, it will intersect marginal benefits at some other level of abatement. On the figure, we have shown what would happen if marginal cost were higher than expected. The efficient level of abatement after the fact would be X^{EFF} , which is less than X^* (because firms find it more costly to reduce their pollution than the regulator anticipated). But of course the industry abates all the way up to the required amount, which is X^* . The resulting deadweight loss is shown by the shaded triangle.

Now let's turn to a price instrument (i.e., a tax on emissions). The best the regulator can do is to set the tax (T^*) equal to marginal damages at the expected efficient level of abatement, X^* . Again, suppose that marginal abatement cost turns out to be higher than expected, so that the efficient level of abatement is X^{EFF} . How will the polluting firms respond? They will abate until their marginal costs of pollution control are just equal to the tax. This level of abatement is labeled X_{TAX} on the figure. Beyond that point, the tax savings from more abatement are outweighed by the costs of pollution control. (We will see this in more detail in the next chapter.)

Under the tax, therefore, firms will do too little abatement when reducing pollution is more costly than expected. (If the regulator had had better information on cost, she would have imposed a higher tax, which would have been sufficient to achieve the efficient level of pollution control.) Again, some deadweight loss will be realized.

The key point is that the deadweight loss from allowance trading is not the same as the deadweight loss from the tax when abatement costs are uncertain. The reason is that abatement differs in the two cases. Under a cap-and-trade policy, the amount of abatement is fixed at X^* by the cap. On the other hand, under a tax the amount of abatement varies with the true marginal cost.

Because the deadweight losses are different, it matters for efficiency which policy the regulator chooses. The preferred policy, naturally, is the one with the lower deadweight loss. In figure 8.4, you can see that the deadweight loss is smaller under the quantity instrument. But compare that with figure 8.5. In that case, the deadweight loss is smaller under the price instrument. What has changed?

If you look closely, you'll see that the difference in the two graphs is the *relative slopes* of the marginal benefit and marginal cost curves. In figure 8.4, the marginal benefit curve is steep relative to the marginal cost curve;

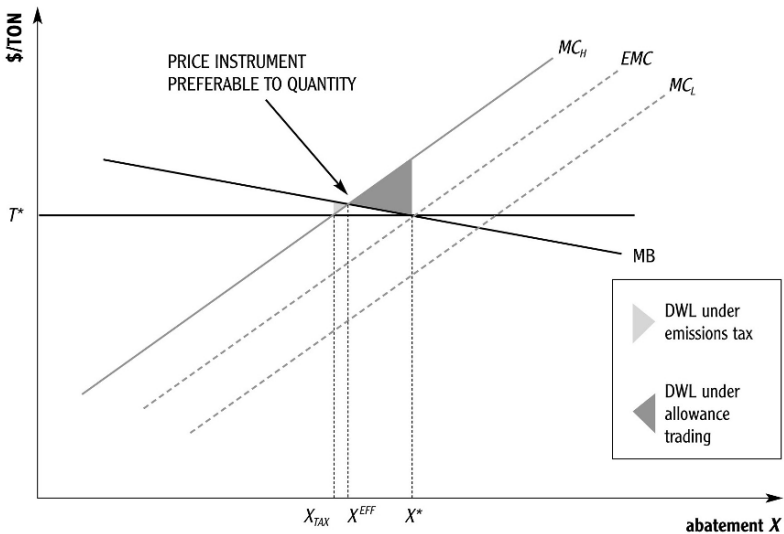


Figure 8.5 Comparison of price (emissions tax) and quantity (allowance trading) instruments under marginal cost uncertainty. In this case, marginal benefit is flat relative to marginal cost, and the emissions tax (the price instrument) is preferred. All else is as in figure 8.4.

in figure 8.5, the marginal benefit curve is relatively flat. In fact, those relative slopes are precisely what matters. This result is often referred to as the Weitzman Rule.

- When marginal costs are uncertain, the efficient choice of policy instrument depends on the relative slopes of the marginal benefit and marginal cost curves.

In particular, a price instrument is preferred when the marginal benefit curve is flatter than the marginal cost curve, and a quantity instrument is preferred when the reverse is true.

As we have seen, the two policy instruments differ in the amount of flexibility they give to the industry to respond to abatement costs. Uncertainty matters because firms can respond to the realization of marginal cost under the tax but not under a cap-and-trade system. Whether this flexibility is a good thing or not (from the point of view of efficiency) depends on the relative slopes of the marginal benefit and cost curves.

A simple way to understand the intuition behind this relative slopes rule is to focus separately on benefit and cost in turn. First, consider the

slope of the marginal benefit function. Recall that we can think of the regulator's goal as replicating the missing demand for abatement—that is, the underlying marginal benefit curve. A tax represents a horizontal demand curve, which is a better approximation when marginal benefits are flat. When marginal benefits are steep, on the other hand, a vertical demand curve—corresponding to the quantity policy—is preferable.

You can also gain intuition by thinking about what the shape of the marginal benefit curve implies about the real world. Steep marginal benefits correspond to a threshold effect around X^* , because the gains from further abatement drop off sharply at that point. (Conversely, the damages from pollution rise sharply as we increase pollution.) In this case, the regulator may want to ensure that the quantity target is met, because the damages from too little abatement will be very high. In contrast, a flat marginal benefit function implies no such urgency, because every ton of abatement brings roughly the same benefit.

Now consider the role of marginal costs. A flat marginal cost curve implies that the amount of abatement chosen by the regulated firms is highly sensitive to the tax, because a small change in marginal cost corresponds to a large change in abatement. Thus, when marginal cost is flat, a small error in the size of the tax (and recall that some error is inevitable, given uncertainty) induces a large error in abatement. Loosely speaking, the potential costs to society of the abatement flexibility afforded by the tax are much greater when the marginal cost curve is flat. When the marginal cost curve is steep, the tax has much less effect on the actual level of abatement.

So far, we have discussed uncertainty only in marginal costs rather than marginal benefits. Will the same arguments hold for the latter kind of uncertainty? The answer is “no.” In fact, uncertainty in marginal benefits makes no difference in the choice of instrument. This might surprise you at first. But note that—unlike marginal costs—marginal benefits are irrelevant to the polluting firms' decisions. After all, if the firms took marginal benefits into account, there would be no externality in the first place! Because firms ignore marginal benefits in their own calculations, the fact that the regulator is uncertain about marginal benefits has no impact on which instrument she should choose. Although deadweight loss will result (because the price or quantity chosen ahead of time will be different from the one that ends up being efficient), the actual abatement will be the same under both instruments, and thus the deadweight loss will be the same as well.

Conclusion

Now we've seen how economic theory can inform the design of environmental policy. We learned from Ronald Coase that there are conditions under which private bargaining will take care of negative externalities. In many cases of interest in the environmental realm, however, transaction costs will be large enough that government policies are needed.

In such cases, economics has much to say about what those policies should look like. Emission taxes and cap-and-trade policies use market principles to restore the efficiency of the market. Whether we think of these policies as getting the prices right, as filling in the missing demand for public goods, or as establishing property rights (and excludability) over common resources, the basic mechanism is the same: Market-based instruments align the incentives of private firms and individuals with the public interest. Finally, we saw how the choice between controlling price (through an emission tax) or quantity (through a cap-and-trade system) can have important implications for efficiency.

So far, our discussion has been focused on how market-based instruments can achieve *efficiency*. In the next chapter, we'll see how a strong case for such policies can still be made, regardless of how the ultimate goal of the policy—for example, the level of abatement—is determined.

9

The Case for Market-Based Instruments in the Real World

Now we know how market-based instruments can be used—at least in theory—to restore the efficiency of markets. However, efficiency may not be the relevant target in the real world, for several reasons. As we have seen, it is very difficult to ascertain just how much benefit we get from these policies; indeed, it may even be hard to estimate the actual costs. But without knowing costs and benefits, we cannot determine the efficient level of pollution. Even if the marginal benefits and costs of pollution control are known, moreover, there is no guarantee that the government will set the efficient target as the goal. Distributional equity and other worthy social goals may be at odds with efficiency. And of course the political process in the real world is driven by interest group competition and the desires of legislators to satisfy their constituents as much as (or more than) by an objective attempt to maximize social welfare.

Even if economic efficiency is elusive in practice, there are still advantages to market-based instruments. First, such policies are cost-effective, meaning that they can achieve any given abatement target at the minimum total cost—something that is not generally true of command-and-control approaches.¹ Second, over the long run, market-based instruments are likely to provide stronger incentives for the development of new pollution control technologies. This *dynamic incentive* will tend to lower abatement costs over time.

In this chapter, we first discuss each of these advantages in detail for the case of pollution control. Next, we consider how similar principles play out in the realm of natural resource management, focusing on the

case of a fishery. Finally, we briefly review some of the other arguments for (and against) market-based instruments. Although they are powerful tools, market-based instruments are not panaceas. We close the chapter by discussing some conditions under which command-and-control approaches might be preferred.

Reducing Costs

In the previous chapter, we discussed taxes and cap-and-trade systems as ways of achieving efficiency. But the question of how much environmental protection to achieve can be separated from the question of how to achieve it. In other words, we can distinguish between goals (ends) and instruments (means). This distinction is key to understanding the notion of cost-effectiveness. Taking the policy goal as given, we may still ask: How do the various policy instruments perform in terms of the total cost of achieving that goal?

Imagine you are a regulator who has been tasked with designing a policy to ensure that a specific pollution target is reached. For simplicity, suppose there are only two polluters in the industry. In the absence of any regulation, they emit 150 tons of pollution. You (the regulator) are supposed to cut that by two thirds—that is, you must achieve 100 tons of abatement, so that combined emissions are only 50 tons. The question is how to do that.

Figure 9.1 provides a useful way of depicting the problem. There are two firms; call them firm A and firm B. The horizontal axis measures abatement—but with a twist. Each point along the axis represents a different allocation of the hundred tons of total abatement between the two firms. As we move from left to right, the share of abatement done by firm A increases, while firm B's share decreases; their combined abatement stays constant. For example, at the left-hand corner of the figure, firm A's abatement is 0 and firm B's abatement is 100 tons. At the middle of the axis, each firm abates 50 tons. Finally, at the right-hand end, all 100 tons of abatement is achieved by firm A, while firm B does nothing. The key thing to realize is that at every point along the horizontal axis, the total policy target (100 tons) is met. What varies is how that target is divided between the two firms.

The vertical axis measures marginal cost, in (say) dollars per ton. Consistent with the way we have measured abatement on the horizontal axis, firm A's marginal cost curve (labeled MC_A) increases from left to right, while firm B's marginal cost curve increases from right to left. We have




Figure 9.1 Marginal costs in a two-firm polluting industry. Note that the horizontal axis measures the allocation of a constant amount of abatement between firms A and B. The share of abatement done by firm A increases as one moves to the right. At every point, however, total abatement is 100 tons. The shaded area labeled *a* represents firm A's total abatement cost under a uniform standard. Areas *b* and *c* combined correspond to firm B's total abatement cost under the same standard. The cost-effective allocation is where the two *MC* curves intersect. Area *c* represents the cost savings from the cost-effective allocation, relative to the uniform standard.

drawn the curves so that firm A is the “low-cost firm”; that is, firm A can achieve any given amount of abatement at lower marginal cost than firm B.

Cost-Effectiveness

Remember that your task, as the regulator, is to find the policy instrument that minimizes the total abatement cost. A natural approach would be to split the abatement equally between the two firms and require each firm to reduce pollution by 50 tons. The total abatement cost for firm A is the area under its marginal cost curve, labeled MC_A , up to 50 tons of

abatement. On the figure, this corresponds to the shaded area labeled a . Similarly, the total abatement cost for firm B is the area under MC_B up to 50 tons of abatement—the sum of areas b and c in the figure.

Is there a way to achieve the same total abatement at lower total cost? The figure suggests that there is. Starting from the uniform allocation, suppose we increase the share of abatement done by firm A. That lowers total abatement cost: the sum of the areas under the marginal cost curves, up to each firm's level of abatement. Total abatement cost continues to shrink as we assign more abatement to firm A, as long as firm A's marginal cost curve lies below firm B's curve. In fact, the combined area under the two marginal cost curves is minimized where the two curves cross. Beyond that point, MC_A lies above MC_B . Continuing to assign greater pollution control to firm A, at that point, will increase rather than decrease total cost.

In this simple scenario, therefore, the cost-effective allocation of abatement—that is, the allocation that achieves 100 tons of abatement at lowest total cost—is given by the point where the two marginal cost curves intersect. In figure 9.1, that corresponds to 60 tons of abatement by firm A and 40 tons of abatement by firm B. The cost savings from this allocation—relative to the uniform standard—is represented by area c .

It should not surprise you that the way to minimize total cost is to choose the allocation that equates the firms' marginal costs. Indeed, this is just another example of the equimarginal principle we saw in Chapter 2. To see the intuition on the margin, go back to the uniform standards case, when each firm must abate 50 tons. At that point, abatement is more costly on the margin for firm B than for firm A. (As we have drawn the figure, firm B's marginal cost at 50 tons of abatement is \$150, whereas firm A's marginal cost is only \$100.) Suppose we took one unit of abatement away from firm B and assigned it to firm A. The amount of abatement would be unchanged, but the total cost would go down—by exactly the difference in the marginal costs, or \$50.

As long as firm B's marginal cost is greater than firm A's marginal cost, we can continue to shift abatement from B to A, reducing cost without affecting abatement. This remains true until the two firms have equal marginal costs.

Note the important distinction between equal marginal abatement costs (given certain levels of abatement) and equal marginal abatement

The cost-effective allocation of abatement is the one that achieves a given level of pollution control at the lowest total cost.

cost *functions*. We do not assume that firms have the same marginal abatement cost functions; in figure 9.1, for example, any given amount of abatement is more costly for firm B than for firm A. Moreover, each firm is doing different amounts of pollution control, with the low-cost firm A doing more abatement. Cost-effectiveness simply requires that abatement costs are equal on the margin, given the amounts of abatement each firm is doing.

Although we can depict only two firms in a graph like figure 9.1, the same intuition applies when there are many polluting firms. For example, consider a policy to achieve 2,000 tons of abatement in an industry with 100 firms. If that were divided equally between the firms, each firm would have to abate 20 tons. Now pick any two firms in the industry. If they have different marginal costs, we can shift abatement from the higher-cost firm to the lower-cost one until their marginal costs are equal. We can continue to do this with every possible pair of firms until it is no longer possible to shift abatement from one firm to another without increasing total costs. At that point, the marginal abatement cost is equal across all firms. In fact, we have just identified a key condition for cost-effectiveness:²

- Cost-effective allocation of abatement occurs only when all firms that abate pollution have equal marginal abatement costs, given their abatement allocations.

In plain English, the last unit of pollution control done by every firm must cost the same amount. Otherwise, there would be a way to reallocate abatement and reduce total cost.

Command-and-Control Approaches

Now let's consider how the various policy instruments we introduced in Chapter 8 perform on this dimension of cost-effectiveness. First, consider technology standards, which require firms to install particular methods of controlling pollution. You may have already guessed from our discussion that such regulations are not cost-effective in general. In terms of our "necessary condition," technology standards typically fail because different firms will have different costs of installing the same sort of technology. For example, some power plants may have plenty of space to build a scrubber, whereas other plants have to build specially designed units to fit into a limited footprint on the ground. More fundamentally, technology standards are generally not cost-effective because they do not even minimize costs at the level of an individual polluter. Some firms are likely

to have other ways of limiting pollution that could achieve the same low levels of emissions as the required technology, at lower cost. Another way to reduce sulfur dioxide emissions from power plants is to burn very low-sulfur coal. Modifying a plant to burn low-sulfur coal, and even paying a premium for it (versus cheaper high-sulfur coal) is often much less expensive than installing and operating a scrubber—but nearly as effective in reducing pollution. We will see a vivid example of how costly a lack of flexibility can be in the real world, when we consider the case of sulfur dioxide control in Chapter 10.

Next, let's consider performance standards, which impose emission limits on individual firms. Return to our two-firm example, with a total emission target of 50 tons. A uniform performance standard would set a ceiling of 25 tons for each firm. Because we have assumed that firms A and B would emit the same amount of pollution (75 tons each) in the absence of regulation, this uniform standard on emissions amounts to a requirement that each firm does the same amount of abatement. We have already seen that such uniform regulation is not cost-effective in figure 9.1. Indeed, from our discussion of cost-effectiveness we can now see that uniform standards will never be cost-effective, as long as firms have different marginal abatement cost functions—that is, as long as firms face different opportunities to reduce their emissions.

“But wait a minute,” you might say, “that’s a problem with *uniform* standards, not with performance standards in general!” Indeed, that is correct. If you (as the regulator) knew the marginal cost curves of the two firms, you could establish firm-specific performance standards corresponding to the cost-effective allocation. In the scenario of figure 9.1, this would mean setting emission standards of 15 tons for firm A (requiring 60 tons of abatement) and 35 tons for firm B (40 tons of abatement).

Although cost-effective performance standards are theoretically possible, they impose an unrealistically high informational burden on the regulator. Firms have obvious incentives to misrepresent their true marginal cost curves: In our simple two-firm example, each firm would like the regulator to think that it was the high-cost firm, so as to minimize the abatement burden the regulator might assign. Meanwhile, deriving firm-level marginal cost functions without the firms’ cooperation would require data far more detailed than what real-world regulators have. As we shall now see, market-based instruments can achieve cost-effective allocations even when the regulator has much less information about abatement costs than what would be needed to set a cost-effective performance standard.

Emission Taxes

To see why an emission tax is cost-effective, let's consider what an individual firm will do when faced with a tax. Figure 9.2 depicts the marginal abatement cost curve for firm A. As usual, the horizontal axis measures pollution abatement. We have denoted firm A's maximum abatement (unregulated emissions) on the horizontal axis as well.

The dashed line on the figure represents the emission tax in dollars per ton. At any given level of abatement, firm A's total compliance cost is the sum of its abatement cost (the area under the marginal cost curve to the left of its abatement level) and its tax bill (the area of the rectangle under the tax and to the right of the abatement level—that is, the tax times emissions, where emissions are just the difference between abatement and unregulated emissions, X^{MAX}). This compliance cost is minimized at the point where marginal abatement cost equals the tax (denoted X' on the figure). (This is just like a cost-minimizing firm that produces up to the point that its marginal cost of production equals the price it receives for its output.) If the firm were to abate less than X' , its compliance costs would be higher, because it would pay more in emission taxes than it saved in abatement costs. On the other hand, costs of more abatement than X' would exceed the tax on the margin.

Nothing is special about firm A in this regard. Indeed, we could repeat the same analysis with firm B, or indeed with any firm. In every case, a cost-minimizing firm would choose the abatement level to equate its marginal abatement cost with the emission tax.

What does this imply about cost-effectiveness? Because every firm is setting its marginal abatement cost equal to the tax, and every firm faces the same emission tax, it follows that marginal abatement cost is equal across all the abating firms. Therefore, the condition for cost-effectiveness is met. Return to figure 9.1, which depicts firms A and B on the same set of axes. We showed earlier that the cost-effective allocation (60 tons by firm A, 40 tons by firm B) coincides with the intersection of the marginal abatement cost curves. Notice also that the cost-effective allocation is the only point at which the marginal abatement costs of the two firms are equal. Because the emission tax ensures that the two firms have the same marginal abatement cost, it follows that the emission tax achieves the cost-effective allocation. Again, the cost savings relative to a uniform performance standard are represented in the figure by the shaded triangle labeled c .

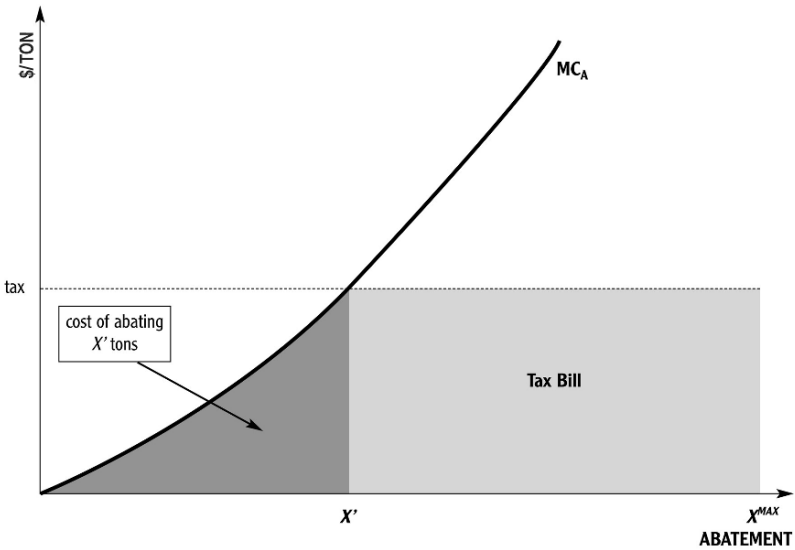


Figure 9.2 Choice of abatement by a cost-minimizing firm. The graph depicts a marginal abatement cost function for an individual firm. The cost of compliance is the sum of the tax bill and the cost of abatement—the two areas shaded in the figure. The least-cost abatement level X' is the level at which marginal abatement cost equals the tax.

Moreover, all this happens without direct intervention from the regulator. All the regulator does is set the tax. Given that tax, each firm independently chooses to abate at the level where its marginal abatement cost equals the tax. In doing so, it ensures that its marginal abatement cost is equal to that of every other regulated firm, so that the aggregate abatement is achieved at least cost.

Of course, for the regulator to successfully implement the tax, she must know how high to set the tax in order to achieve the desired policy target. In particular, the regulator must know the *aggregate*, or industry-level, marginal abatement cost curve. Such a curve traces out the cost of the last unit of abatement, as a function of the total abatement done by the industry as a whole. It is derived by summing up the abatement done by all the firms in the industry, for any given marginal cost. Figure 9.3 provides an illustration. In it, we have drawn the marginal abatement cost curves corresponding to firms A and B depicted in figure 9.1, along with the aggregate marginal abatement cost curve corresponding to the two firms put together.

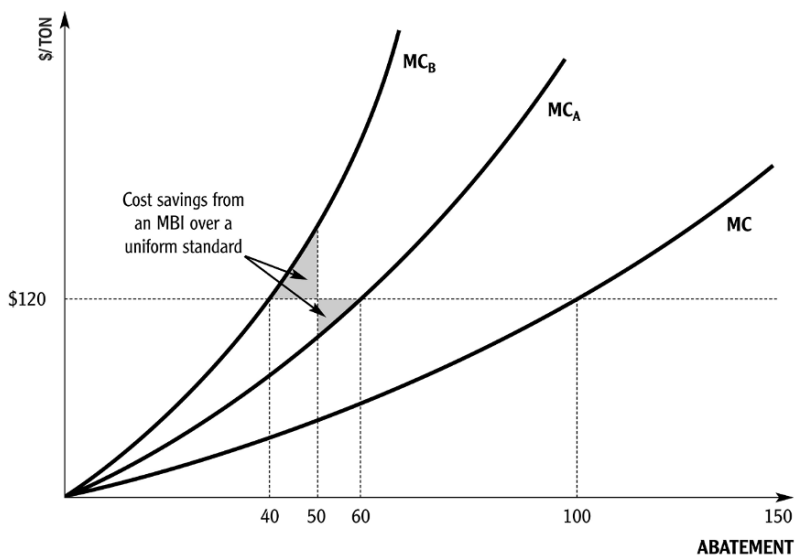


Figure 9.3 Individual and aggregate marginal cost curves. The aggregate curve (denoted MC) is the sum of the quantities abated by the two firms, at every level of marginal cost.

We have already noted the parallel with a supply curve in a market for private goods. In particular, one can think of the aggregate marginal abatement cost curve as the “abatement supply curve.”³ Recall that in the standard market setting, the supply curve tells us how much an industry will produce at a given price—or equivalently, how high the price of a good has to be to induce firms to produce a given quantity. In exactly the same way, the aggregate MC curve in figure 9.3 summarizes the relationship between the marginal cost of abatement and the total amount of abatement. Because the tax acts like a price on abatement, the total abatement induced by a tax can be found from the point on the horizontal axis directly below the intersection of the aggregate MC curve with the tax. In the case of our two firms, A and B, the required tax is \$120. That number corresponds to the height of the aggregate MC curve at 100 tons of abatement.

As in figure 9.1, we can depict the cost savings from using the emission tax (rather than a uniform standard) directly on the graph. In this case, the cost savings are represented by the two shaded areas. Close inspection should convince you that these two areas together correspond precisely to the shaded triangle ϵ that identifies cost savings in figure 9.1.

Emission Trading

Finally, let's consider a cap-and-trade policy. We can start by returning to figure 9.1. Once again, imagine that you are the regulator charged with limiting the combined emissions of firms A and B to 50 tons. Now consider the following policy. Suppose you start by assigning each firm 25 tons of emissions, as in the case of a uniform performance standard. But instead of requiring the firms to meet that standard individually, you give each firm 25 pollution allowances (worth 1 ton each) and allow them to trade. What will happen?

At the initial allocation of 25 tons of emissions each, each firm must abate 50 tons. At that level of abatement, firm A has a marginal cost of \$100, versus \$150 for firm B. Therefore, a trade will be in the interests of both firms. In particular, the manager of firm A could contact his counterpart at firm B and propose to sell her one allowance. Firm A would reduce its emissions by 1 ton and sell the resulting extra pollution permit to firm B. Firm A would be willing to do this for any price above \$100, and firm B would be willing to pay up to \$150. In the language of economics, there are "gains from trade." As long as the two firms can easily trade with each other, we would expect them to make such a deal.

These gains from trade remain as long as the marginal costs of the two firms are not equal. In other words, the two firms will have an incentive to keep trading until their marginal costs are equated. But this means that trading achieves the cost-effective allocation of abatement! In terms of figure 9.1, we would expect firm A to sell 10 pollution allowances to firm B. After trading, firm A would end up abating 60 tons (and receiving some payment from firm B), while firm B would abate only 40 tons (but would pay firm A for the 10 tons it was not abating). In terms of emissions, firm A would emit 15 tons (below its initial allocation of 25 permits) and firm B would emit 35 tons.

Moreover, note that this argument does not depend on how the allowances are initially allocated. Suppose, for example, that firm A receives only 10 allowances, and firm B receives 40. (Total allowable pollution is 50 tons, just as it was before.) Without trading, firm A would therefore be required to abate 65 tons, and firm B would have to abate only 35 tons. Once again, there are gains from trade. This time, however, the marginal cost of the last unit of abatement is higher for firm A than for firm B. In figure 9.1, firm A's marginal cost curve is above firm B's marginal cost curve when firm A abates 65 tons and B abates only 35 tons. Given this initial allocation, therefore, firm A would buy allowances from firm B. The

gains from trade would remain until the marginal costs were equal, at the cost-effective allocation of 60 tons of abatement by firm A and 40 tons by firm B. We conclude that as long as firms can easily trade with one another, the initial allocation of pollution allowances does not affect the final (equilibrium) allocation.⁴

We have described this cap-and-trade system in the context of only two firms, but the logic generalizes immediately to any number of firms. As long as all firms take the price of pollution allowances as given—that is, as long as no firm has the power to set prices in the allowance market—then each firm will abate up to the point at which its marginal abatement cost equals the allowance price. Below that point, reducing pollution is less expensive than paying for permits. Beyond that point, it would make more sense for the firm to buy allowances rather than abating. Of course, because the total number of allowances is capped by the regulator, the number of allowances bought by firms that choose to abate less than their initial allocations must be exactly balanced by the allowances sold by firms that abate more than required.

Because every firm abates until its marginal cost equals the allowance price, and all firms face the same allowance price, marginal cost must be equated across all firms. Just like an emission tax, therefore, a cap-and-trade policy is cost-effective. The market mechanism ensures that abatement is carried out by the firms that can do it at least cost.

One potential advantage of cap-and-trade systems is worth pointing out. The goal of environmental regulation is often defined in terms of the desired quantity of pollution abatement (or of environmental protection more generally). If so, a cap-and-trade system is easier to implement than a tax. This is because the connection between an emission tax and the resulting quantity of emissions is indirect: The regulator sets the price, and the actual level of pollution depends on the industry's marginal costs of abatement. Therefore, the regulator must know the aggregate marginal abatement cost curve in order to attain a particular level of abatement with an emission tax. In contrast, a cap-and-trade system determines the quantity of pollution directly; no information other than the policy target is needed.

Equivalence of Taxes and Emission Trading, Revisited

Recall from Chapter 8 that when marginal abatement costs are known with certainty, an efficient tax and an efficient cap-and-trade program yield the same price and quantity of emissions. This basic equivalence does not depend on efficiency; it holds regardless of how the abatement

target is set, as long as the regulator sets the level of the tax or the cap to achieve that target. For any choice of abatement target X , the marginal abatement cost curve provides the corresponding tax T that would induce the same quantity of abatement and vice versa. (To see this, note that figure 8.3 could be drawn for any choice of abatement target X .)

Of course, as we've already discussed, marginal abatement costs are unlikely to be known with certainty. In that case, the Weitzman Rule can help guide the choice of policy instrument, even if the desired abatement target is not necessarily efficient.

Cost Savings and the Degree of Heterogeneity

You may have realized by now that the cost-effectiveness advantage of market-based instruments depends on the variation (or *heterogeneity*) in the marginal abatement cost functions of different firms. To see this graphically, return to figure 9.1 and imagine redrawing the figure with firm A's marginal abatement cost function somewhat flatter than it is and firm B's marginal cost function even steeper. Doing so would increase the cost savings from a cap-and-trade program or emission tax relative to a uniform standard. (Recall that those cost savings are represented by the shaded area labeled c in the figure.) On the other hand, if firms were identical, the cost-effective allocation of abatement would coincide with a uniform standard, and there would be no cost savings from a market-based instrument. The intuition is simple: The greater the differences in abatement costs between the firms, the greater will be the gains from allowing them to trade.

Promoting Technological Change

The notion of cost-effectiveness we have been using is essentially a static one. If a regulator has a certain target in mind for pollution control, for example, market-based instruments can achieve that goal at lower cost than would be possible through uniform performance standards. They do this by reallocating pollution control from firms with high costs of abatement to firms with low costs—although such reallocation results not from any centralized coordination but from the actions of individual firms pursuing their own profits.

Our description of how market-based instruments work has essentially relied on a snapshot of the polluting industry at a particular point in time. For example, in the

The cost-effectiveness advantage of market-based instruments depends on the variation (or heterogeneity) in the marginal abatement cost functions of regulated firms.

previous section we just asserted that firm A has lower marginal abatement costs than firm B, but we did not ask where that cost differential comes from. We simply took the marginal abatement cost functions for granted. Nor did we ask whether firm B might be able to upgrade its pollution control technology—say, by investing in new equipment, thereby lowering its marginal abatement cost. In other words, we implicitly viewed abatement technology as fixed.

Cost-effectiveness is an important criterion, but it is limited by the static perspective we have just described. A complementary means of assessing policy instruments is to ask how well they promote technological change, such as the development and adoption of new, cheaper pollution control technologies. It turns out that market-based instruments outperform standards on this dimension as well. More precisely, the incentive to adopt new technologies with lower marginal costs is greater under an emission tax than under a performance standard. The incentive under a cap-and-trade system is less than under a tax, but it is still likely to be larger than under a performance standard. This is an extremely important advantage when we consider very long-run environmental policy problems such as those related to climate change. The invention and diffusion of new technologies will be essential to achieving meaningful greenhouse gas emission reductions over a century or more.

Incentives for Technology Adoption Under Various Instruments

Let's start by comparing an emission tax with a performance standard. Figure 9.4 depicts a firm whose initial marginal abatement cost function is denoted MC_0 . Another technology becomes available, with lower marginal cost MC_1 . Is the firm more likely to adopt the new method under the tax or under the performance standard? The gain to the firm from adopting the technology depends on the cost savings to the firm from having marginal cost function MC_1 rather than MC_0 . The greater these cost savings, the more likely the firm is to adopt the new technology.⁵

We can use figure 9.4 to illustrate how the reward from adoption varies under the two policies. To focus narrowly on the question of technology adoption, let's assume that the tax and standard are *equivalent* under the initial cost function MC_0 . In other words, the firm would choose exactly the same amount of abatement under the tax as it would be required

The incentive to adopt new technologies with lower marginal costs is greater under an emission tax than under a performance standard.




Figure 9.4 The incentive to adopt a new pollution control technology under a performance standard and emissions tax. The two policies are assumed equivalent under the original technology (MC_0). The new technology MC_1 has fewer costs; its adoption results in greater abatement under the tax, and hence greater cost savings to the firm.

to do under the performance standard. On the graph, this means that the tax (represented by the horizontal dashed line) intersects the initial marginal cost curve at the emission standard (the vertical dashed line at X_{STD}).

How much will the firm facing the performance standard save in abatement costs if it adopts the new technology? The cost savings from the new technology (not counting the capital cost of installing the equipment, which we have not specified) correspond to the area of the shaded triangle labeled *a* in the figure. Note that under the performance standard the firm has no incentive to abate more than it is required to do. Thus even if the firm adopts the lower-cost technology in that case, it will continue to abate the amount X_{STD} .

What is the reward from adoption if the firm is regulated by the tax instead? Under the tax, the firm abates up to the point that marginal cost equals the tax. Therefore, with the new technology the firm would abate all the way to X_{TAX} , which is greater than X_{STD} . Intuitively, a tax provides an incentive to abate as much pollution as is economically sensible. With a lower-cost abatement technology, it makes sense for the firm to abate more pollution, because it has to pay taxes on whatever it continues to

emit. Therefore, the increased abatement translates into additional cost savings. This is represented by the area of triangle *b* in the figure.

Intuitively, we can think of area *a* as “cost savings from doing what the firm was already doing under the old technology, only cheaper.” We can think of area *b* as “cost savings from increasing abatement, now that it’s cheaper, rather than paying emission taxes.” The firm gets area *a* from adoption under both the standard and the tax, but area *b* is gained only under the tax. We can conclude that the tax creates a greater incentive for the adoption of new abatement technologies.

A cap-and-trade system does not lend itself as well to a comparison like the one in figure 9.4, because the “cap” applies to the market as a whole, not to individual firms, and because different firms will necessarily have different cost functions (if they did not, there would be no gains from trade). Nonetheless, we can glean some intuition by comparing a cap-and-trade system to a tax. As we have noted several times, a tax fixes the price of abatement, whereas a cap-and-trade system fixes the quantity. This distinction has an interesting implication for technology adoption. As more and more firms in an industry adopt the new technology, the price of tradable allowances will fall, because abatement becomes less expensive on the margin. However, this fall in the price *reduces* the incentive for other firms to adopt the technology, relative to the constant incentive under a tax. (You can see this on figure 9.4 by shifting the tax line down, representing a fall in the price of pollution.) This should make sense: After all, if allowances become less expensive, firms will have less reason to look for a cheaper way of reducing pollution.

This feedback between technology adoption and the price of pollution means that cap-and-trade policies will typically provide less incentive than a tax for the adoption of new abatement technologies, assuming that the tax and cap-and-trade system would achieve the same initial level of abatement given the technology in place when the policy is implemented. (Of course, policymakers might well choose to tighten the cap, or adjust the tax, in response to new technologies.)

The effect of technology adoption on the allowance price also makes the comparison between a cap-and-trade system and a performance standard less straightforward. But the basic advantage of market-based policies continues to hold. By putting a price on pollution, such policies give firms a strong incentive to find cheaper ways of cutting emissions. In general, then, we expect adoption of cheaper pollution control technologies to be greater under either an emission tax or a cap-and-trade policy than under command-and-control.⁶

You may have noticed that so far we have concentrated on performance standards rather than on technology standards. Would a technology standard do better at promoting technological development? At first, you might think that it would have to. After all, under such a policy all firms have to adopt the required technology. Doesn't that have to mean that technology standards provide the greatest adoption incentive of all?

Two distinctions are key here. The first is between adopting a particular technology mandated by the government and adopting a lower-cost technology. Although we have described the gains from technology adoption in terms of just one new technology, the intuition applies even when there are many ways of reducing pollution. Faced with a menu of options, each firm will choose the one that suits it best, providing the greatest abatement cost savings in return for the investment needed. On the other hand, a technology standard requires firms to choose a particular technology, regardless of their other options. Moreover, there is no reason to think that the government will have better information than firms about the true costs of installing and operating various technologies.

The second key distinction is between adoption and innovation. Our discussion so far has essentially assumed that lower-cost technologies simply pop into existence, awaiting firms to adopt them. In fact, that ignores the crucial role of innovation: the development of new technologies by firms that hope to sell them. Because market-based instruments give polluters a continuing incentive to search for new and better ways of reducing pollution, they provide a constant spur to innovation. Technology standards, in contrast, run the risk of locking in a particular technology—whatever is the state-of-the-art at the time of the regulation—and thereby unintentionally dampening the incentives for firms to come up with new technologies.

Policies Aimed at Promoting Innovation

Although a strong incentive to generate and adopt new technologies is an important feature of market-based instruments, it is also an *indirect* one: It arises as a byproduct of putting a price on pollution. Especially in the case of very long-term challenges such as climate change, policymakers may also want to implement additional policies aimed directly at promoting technological innovation. Policies to promote technological change make sense from an economic perspective. Recall our discussion in Chapter 5 that research and development efforts by firms create a positive externality in the form of knowledge that benefits other firms. From the perspective of economic efficiency, therefore, firms will tend to do too *little* research

and development on their own, because they will not take into account the spillover benefits to other firms. This is especially true for basic scientific research, where the spillover benefits to other firms are generally large relative to the private benefits that the firm paying for the research can capture. For this reason, pharmaceutical firms devote vastly more resources toward applied research aimed at developing specific therapeutic drugs (which they can patent) than toward fundamental research on the underlying disease mechanisms.

This positive externality explains why governments typically play an active role in promoting research and development. Government support for R&D can take many forms: direct funding of government laboratories, competitive grants to research universities, subsidies to R&D efforts by firms, and so on.

Market-Based Instruments for Managing Natural Resources

In Chapter 7, we showed how open-access conditions lead to market failure in fisheries and other natural resources. As we saw there and in Chapter 8, efficient management of a fishery can be achieved by establishing property rights, either in fact or in effect. But just as in the case of pollution control, efficient management of a fishery (or other resource) requires a great deal of information—not only the marginal costs of harvesting but also the biological growth function of the resource. In the first part of this chapter, we saw how emission taxes and allowance trading can lower the cost of meeting a given target for pollution control, whether or not that policy goal is efficient. Similarly, we saw how such market-based policies could encourage the adoption and development of new technologies over time.

Now we look at how market-based instruments can be used to manage natural resources. We focus on fishery management as our example, but as always the lessons are general. In Chapter 10, for example, we will see how these same principles apply to resources as diverse as scarce water supplies and valued wetland ecosystems.

Cost-Effectiveness

The key intuition is that cost-effectiveness means the least-cost allocation of effort to achieve a particular goal. In the case of a fishery, the goal is defined in terms of the annual allowable harvest, and effort is the time and money fishers spend on catching fish and purchasing and maintaining their equipment.

Promoting Innovation through Prizes

On October 4, 2004, nearly 70 miles above the Mojave Desert in California, the privately built *SpaceShipOne* reached suborbital flight for the second time in 2 weeks. The payoff was more than good publicity: The team behind the spaceship won the \$10 million Ansari X-Prize. Since then, the X-Prize Foundation has awarded prizes for an ultra-fuel-efficient car, a new method of cleaning up oil spills, and a lunar lander; several other contests remain current. The success and visibility of the X-Prize have led a number of observers to call for innovation prizes to promote R&D into solutions to energy and climate challenges, such as zero-energy buildings or advanced batteries.

This idea has a long history. One of the earliest and best-known examples was intended to solve a problem that had vexed mariners for centuries: determining longitude at sea. Latitude (the distance north or south of the equator) can readily be estimated by the angle of the sun at its highest point in the sky; longitude is much harder. In 1714, the British Parliament passed the Longitude Act, offering a sum of up to £20,000 for a practical and reliable method of fixing longitude at sea (the size of the prize increased with the degree of precision). The prevailing wisdom was that the solution lay in the position of the stars in the night sky, and royal astronomers in Greenwich set about to make detailed star charts that navigators could use. Instead, the winner of the prize turned out to be a watchmaker named John Harrison, who built a clock that could keep accurate time despite the humidity of the tropics and the pitch and yaw of the open ocean. With such a clock set to local time at the Greenwich observatory, a sea captain could calculate his longitude simply by observing when it was noon at his location (and the sun was at its zenith) and computing the difference with Greenwich mean time.⁷

As the longitude example shows, innovation prizes can yield unexpected solutions that run counter to the conventional wisdom of experts. Indeed, that has often been the case. When the \$25,000 Orteig Prize was announced in 1919 for the first transatlantic flight, it was widely assumed that only a multiengine plane with a large crew could make the trip. Eight years later, Charles Lindbergh confounded expectations—and won the prize—with his daring solo flight.

Innovation prizes work best when there is a clear objective that can be concretely expressed and readily verified, when the upfront cost of pursuing a solution (which must be borne by the competitors) is not too high, and when success involves the provision of a public good (because firms should already have a strong economic incentive to develop innovations that they can fully profit from). Prizes can appeal to policymakers in part because government funds are spent only if a solution is found, reducing the risk of failure (or rather making prize participants bear that risk). Indeed, prizes can be a cost-effective way of generating private R&D investment: If multiple groups invest in pursuit of the prize, the total amount of investment may be much larger than the prize itself. For example, the Ansari X-Prize is estimated to have induced \$100 million in R&D despite paying out only \$10 million.⁸

In Chapter 7, we assumed for simplicity that the cost per unit of fishing effort—that is, the marginal cost of fishing—was constant (at least for a given size of the fish stock). In effect, this means that all fishers are identical and that the last fish they catch requires no more effort than the first. In reality, of course, fishers differ widely in ability and cost. Some individual fishers are better able to read the ocean floor or the habits of seabirds and other animals in order to locate stocks of fish. Large trawlers have lower costs (per unit of fish caught) than small one- or two-person boats. Moreover, it is reasonable to suppose that for any given fisher, the marginal costs of fishing increase with the total catch. At some point, for example, a fisher wishing to catch more fish must venture farther out to sea or to less promising (or less familiar) fishing grounds.

Why does such variation matter? Previously we saw how the degree of heterogeneity among marginal abatement costs determined the cost savings from market-based policies for pollution control. The principle is exactly the same in the case of a fishery: The more varied the costs of fishing effort, the greater the savings from a cost-effective allocation.

To see how this works in the case of a fishery, consider figures 9.5 and 9.6. Suppose there are two fishers in a fishery, who face the same price for fish (which we will take to be fixed) but have different marginal costs of fishing. These assumptions are depicted in figure 9.5, where we have drawn two marginal cost-of-fishing curves, labeled MC_A and MC_B in the figure, and a horizontal line representing the market price of fish. As the figure is drawn, fisher B is the low-marginal-cost fisher. Suppose that there were no restrictions on fishing. In that case, each fisher would earn marginal net benefits (equivalently, marginal net revenues) equal to the difference between the market price (the revenue on each unit of fish) and their own marginal cost. In plain English, the marginal net benefit represents the value to a fisher of each fish she catches. Because marginal costs rise while the price stays constant, marginal net benefits to each fisher fall as her catch increases. Moreover, because fisher B has lower marginal costs of fishing, she earns greater marginal net benefits.

Now take a look at figure 9.6. (This figure recalls both the two-firm pollution abatement model of figure 9.1 and the two-period resource extraction model of figure 6.3.) The two fishers are represented in the figure by their marginal net benefit curves, labeled MNB_A and MNB_B . The length of the horizontal axis represents the total allowable catch in the fishery—say, 100 tons of fish. The question is how to allocate that total catch between the two fishers. Suppose we start by dividing it equally. In this case, each fisher would harvest 50 tons. At that point, however,




Figure 9.5 Marginal fishing costs, the market price, and the resulting “marginal net benefits” for two fishers in a market. When each fisher catches fifty fish (the line shown on the graph), the shorter arrow represents the marginal net benefit to fisher A, while the longer arrow represents the marginal net benefit to fisher B, who has lower marginal costs.

fisher B values the last ton of fish she catches more highly than fisher A values the last ton of her own harvest. (We know this because when each fisher lands 50 tons, the marginal net benefit for B is greater than the marginal net benefit for A.)

Now consider what would happen if instead of allocating the catch equally, we gave each fisher an equal number of individual fishing quotas (IFQs) and allowed them to trade. Because the low-cost fisher B values a ton of fish more than her counterpart (on the margin), she will buy IFQs from fisher A. Indeed, as long as transaction costs are negligible, we would expect the two fishers to trade until their marginal net benefits are equated. This point corresponds to the intersection of the two *MNB* curves. The horizontal dotted line on the figure shows the price at which the last quota would trade. Moreover, the fishers will reach this point regardless of the initial allocation of the IFQs, just as we saw in the case of a pollution market.

The shaded triangle in figure 9.6 represents the total gains from trade, relative to the scenario in which each fisher lands half of the total catch. This is analogous to the cost savings from trading that we found in the

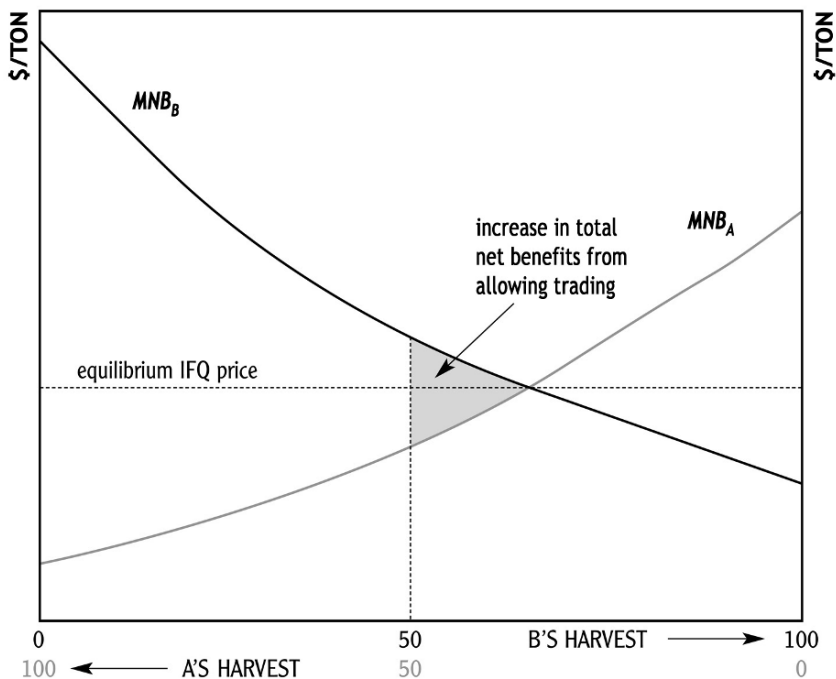


Figure 9.6 Marginal net benefits in a two-person fishery. The horizontal axis measures the total allowable catch, allocated between the two fishers. The share of the harvest caught by fisher B increases as one moves to the right. The share caught by fisher A increases to the left. At every point, however, the total catch is one hundred tons. The cost-effective allocation is where the two MNB curves intersect; the corresponding IFQ price is denoted by the horizontal dotted line. The shaded area represents the gains from an IFQ relative to an even allocation of the harvest. Under an IFQ, the low-cost fisher (B) catches more, while the high-cost fisher (A) catches correspondingly less. Since the marginal net benefits are higher for B, this represents a gain from the point of view of society.

case of pollution control (compare it with area c in figure 9.1). However, the triangle has a slightly different interpretation in this case. Instead of minimizing total costs, we are now interested in maximizing total net benefits. The net benefits to each fisher are the areas under their respective MNB curves. The gains from an IFQ system result because the low-cost fisher (fisher B in the figure) ends up with a larger share of the catch. Her lower costs of fishing translate into higher net benefits. The resulting increase in net benefits to the two fishers combined is shown by the shaded triangle.⁹

Incentives for Technological Change

In the case of pollution control, we saw how an emission tax (and to a lesser extent a cap-and-trade policy) could encourage the adoption and innovation of new, low-cost abatement technologies. What analogous incentives are created by market-based policies for managing a natural resource? It turns out that less theoretical work has been done on this point, but we can sketch out a few basic principles. First, an IFQ market (or a similar policy in another realm of natural resources) establishes secure property rights over the resource. In the absence of such property rights, as we saw in Chapter 7, each individual user has a strong incentive to ignore the effects of her behavior today on the state of the resource tomorrow. Therefore, the first-order dynamic incentive created by securing property rights is to give each participant a stake in the continuing value of the resource—encouraging them to act in ways that sustain the resource's productivity.

Second, an IFQ market changes investment incentives. As we saw in Chapter 7, under open-access conditions fishing effort will increase until the rents from the resource are exhausted. In the simple analysis of that earlier chapter, we measured effort in terms of the number of boats. More generally, however, we might think of effort as also encompassing the capital expenditures fishers make for their boats. For example, if a limit is placed on the total allowable catch, without restrictions on allowable gear, each fisher will have a strong incentive to buy the biggest, fastest boat possible, in order to take as many fish as she can before the total allowable catch runs out for the season. Earlier in the text, we described the tragedy of the commons in the Pacific halibut fishery, attempts to restrict the catch by shortening the fishing season create similar incentives. The result is overcapitalization: Fishers spend much more on equipment than is needed to catch the harvest, in an arms race with each other. In economic terms, each fisher sees only an incentive to increase her own share of the rent; in striving to do so, fishers as a group drive the rents to zero.

The security created by an IFQ market transforms these incentives. A fisher who is guaranteed a certain percentage of the catch has the proper incentive to invest enough to minimize her costs of harvesting that amount—but no more.

Other Considerations

Cost-effectiveness and the promotion of new technologies are the cornerstones of the economic argument for market-based instruments.

Choosing the best policy in a particular case may involve a number of other considerations, however. Although a full treatment is beyond the scope of this book, we highlight a few additional important considerations in this section.

Pollution Hot Spots

In our analysis so far, we have implicitly assumed that all emissions are created equal—in other words, that the social damages from emissions are the same across all firms. This is a convenient assumption in modeling, because it means we do not have to keep track of who is emitting what. For example, in the model of the first section, we assumed that the relevant measure of pollution was the sum total of emissions from firms A and B combined. Economists call this the *uniform mixing assumption*. If emissions from different sources uniformly mix in the atmosphere, then from society's point of view all that matters is total emissions, rather than the location of the emitters.

Is this a reasonable assumption for the real world? The answer—perhaps not surprisingly—is, “It depends.” In some cases, uniform mixing is a good description of the real world. For example, greenhouse gases such as carbon dioxide disperse widely in the upper atmosphere, regardless of where they come from. A ton of CO₂ emitted in Boston and a ton in Beijing have the same effect on atmospheric concentrations of CO₂ (and hence on climate change). In other cases, however, pollution accumulates locally near where it is emitted or in predictable cross-regional patterns rather than dispersing uniformly. As an extreme example, consider toxic waste, which poisons soil at a specific site, and may seep into groundwater nearby, but does not have effects outside a well-defined watershed.

The degree of mixing matters because the policies we have described rely on uniform mixing to work properly. A cap-and-trade approach is sound policy only if a cap on total emissions makes sense and if allowing firms to trade emissions will not lead to dangerously high concentrations of pollution in localized areas. The same caveat applies to an emission tax, which treats all sources the same (a ton of pollution is taxed the same regardless of its origin) and shifts pollution between firms just as a cap-and-trade does. A nationwide cap-and-trade system (or tax) limiting the aggregate quantity of toxic waste would be inappropriate, because it would ignore the question of how highly concentrated the waste was in any particular site. This is known as the hot spot problem: When pollution is nonuniformly mixed, cap-and-trade programs and emission taxes can result in hot spots of concentrated pollution if firms with high

damages from emissions also have high marginal abatement costs. Under a trading program or a tax, there is nothing to prevent emissions from individual sources from increasing, even as overall emissions fall.

What do we know about how significant the hot spot problem may be in real-world pollution control?

As noted earlier, this is a nonissue for the greenhouse gas emissions that are changing the global climate. The effect of these emissions on greenhouse gas concentrations in the upper atmosphere do not depend on the location of emissions. The U.S. sulfur dioxide trading program, while it operated, did not incorporate differences in marginal emissions damages due to power plant location; firms traded on a ton-for-ton basis, facing a single market permit price. However, SO₂ is not uniformly mixed: In the United States, it is commonly transported from west to east with prevailing winds, so the emissions from midwestern power plants end up in eastern states. Recent work suggests that had firms traded instead, using a system of exchange rates that captured these spatial differences in the impacts of SO₂ pollution, the program's net benefits would have increased by \$310–\$940 million per year, far more than the cost savings achieved by using a market-based policy rather than command and control for this pollution.¹⁰ Nitrogen oxide (NO_x) emissions, which contribute to the formation of ground-level ozone or smog, also create damages that vary with the location of emissions, but they have been regulated using tradable permits that do not fully capture this spatial heterogeneity.¹¹

This issue of nonuniform mixing turns out to be particularly significant for water pollution. In most cases, the location of an effluent source strongly affects pollution concentrations at various receptor points throughout a water body. Marginal damages from pollution can also vary with seasonal or daily weather conditions. Even if hot spots don't arise, simple allowance trading programs that focus only on total pollution may lead to wide gaps in the damages due to pollution from different areas. What should we do in such situations? One solution, of course, is to fall back on command-and-control approaches, such as performance standards. In some cases—for example, toxic effluent released by paper mills into streams—this may prove to be the most sensible policy. Another solution might be to use a market-based approach but prohibit some kinds of trades. For example, we might prevent polluters in a high-damage area

When pollution is nonuniformly mixed, cap-and-trade programs and emission taxes can result in hot spots of concentrated pollution if firms with high damages from emissions also have high marginal abatement costs.

Hot Spots and Pollution Trading: The Case of Nitrogen Oxides

A real-world hot spot problem is demonstrated by the cap-and-trade program for nitrogen oxides (NO_x) that has been functioning in the eastern United States since 1998. NO_x combines with volatile organic compounds, produced by both human and natural processes, to produce ground-level ozone, the main component of smog. Although cars and trucks account for the majority of NO_x emissions, large stationary sources such as electric power plants are also important contributors of this type of pollution. These stationary sources are the ones required to participate in the cap-and-trade program.

Importantly, the trading program operates as if each pound of NO_x causes the same damage, regardless of location; thus a power plant in North Carolina can trade on a pound-for-pound basis with a power plant in Maryland. The problem is that a pound of NO_x in North Carolina and a pound of NO_x in Maryland cause different damages: The former blows out to sea, and the latter contributes to urban smog in the heavily populated downwind areas of New Jersey and New York. Under a simple trading program, there is nothing to prevent the power plant in Maryland from emitting large amounts of NO_x and buying allowances from the North Carolina plant. Such a trade amounts to reducing air pollution over the Atlantic Ocean and increasing it over New York City—not a desirable exchange.

One possible solution (as suggested in the text) would be to define simple *trading ratios* between plants in different regions. One study has demonstrated how this might work.¹² The authors find that setting up source-specific trading ratios to account for differences in marginal damages would increase the total annual benefits of NO_x abatement quite significantly. But implementing these constraints on trade would also increase costs. In some cases, this cost increase might be outweighed by the increase in benefits from spatially differentiated trading; this is the case for the hypothetical water pollution trading program for the Ohio River discussed in the text. But in the NO_x case, it turns out that regulators underestimated what it would cost firms to comply with the original, undifferentiated trading regime. This underestimate, when combined with cost increases from implementing trading ratios, would actually cause the NO_x regulation to have net costs rather than net benefits. This analysis provides a great example of the complexity of pollution regulation in the real world. Although spatially differentiated trading may be desirable for many nonuniformly mixed pollutants, it is not desirable in all cases.

from buying allowances from firms in a low-damage area. This has its own problems, however. For one thing, it is unlikely to make much difference, as long as the market can circumvent the restriction.

Let's take an important water quality problem as an example.¹³ In the Upper Ohio River Basin, there are approximately seventy sewerage systems, including that of the city of Pittsburgh and many smaller cities, with combined sewer overflows. Sewer systems in these areas carry both sewage and, during and after a significant rainfall, stormwater. When the flow exceeds the system capacity, this sewage-stormwater mix is discharged directly into waterways in the basin. The impacts of this include pollution damages from bacteria, biological oxygen demand, and total suspended solids. Suppose we set up a cap-and-trade system in which the cap was a limit on total allowable sewer overflows. The marginal damages from sewer overflows vary greatly, depending on the receiving water characteristics, including flow and other hydrological aspects of the water's capacity to assimilate waste, as well as the exposed population.

To avoid hot spots, we could simply disallow trades between selected pairs of the seventy discharging sewerage systems, but this would be difficult to enforce. For example, suppose we forbid Pittsburgh from selling permits to another city, such as Morgantown, because the marginal damages from Morgantown's emissions are higher than Pittsburgh's. In this case Pittsburgh could still sell permits to the sixty-eight other cities in the basin, and some of those permits could then be purchased by Morgantown in separate transactions. All that would have changed is the transaction cost. And if the restriction did limit trade, it would reduce the gains from using a market-based instrument in the first place; the more divided the market, the smaller the potential gains from trade.

A differentiated market-based alternative would, instead, make it more expensive for Morgantown to buy permits from Pittsburgh. If the regulator knew the relative damages of pollution from those two places, she could institute a trading ratio—something like an exchange rate for pollution permits. A team of economists has estimated the system of exchange rates that would apply to the eight largest cities in the Upper Ohio River Basin, using an Environmental Protection Agency (EPA) water quality model to do so. Accounting for differences in the marginal damages from 1 kilogram of biological oxygen demand load, Morgantown would need to purchase 4.28 kilograms of overflow abatement from Pittsburgh to obtain credit for 1 kilogram of its own abatement. If all trades took place at these estimated exchange rates, then the problem of hot spots from nonuniform mixing of sewer overflows could be solved.

The analogous tax is even more straightforward. Rather than applying the same tax to every source, the regulator could apply a tax more closely tied to the variation in marginal damages. For example, the regulator could charge 4.28 times more for a kilogram of overflow in Morgantown than in Pittsburgh.

Several water quality trading programs now incorporate systems of trading ratios like this so as to avoid increasing emissions, and thus marginal damages, where they might cause the most harm.¹⁴ The point is that nonuniform mixing does not eliminate the case for market-based instruments per se. Rather, it points out the need to think carefully about the policy goal. If hot spots can arise, then a goal of reducing aggregate pollution may not make sense. If we redefine the goal, we can fashion a market-based policy to achieve it.

Monitoring and Enforcement

When we discussed cost-effectiveness, and even when we analyzed the incentives for technological change, we focused exclusively on abatement costs—that is, the costs to polluting firms (or individuals) from installing and operating pollution control equipment, and more generally from changing their production processes to reduce emissions (for example, by burning cleaner fuels). But in principle we care about minimizing all costs associated with environmental protection. After abatement, the prime source of costs is administrative, especially the costs of monitoring and enforcing compliance with environmental regulations.

For environmental regulations that target industries, these administrative costs are likely to be small relative to the costs of pollution abatement. For example, scrubbers to remove sulfur dioxide cost tens or hundreds of millions of dollars to install and several million dollars per year to operate thereafter. The costs of emission monitoring equipment are much smaller—a few hundred thousand dollars for installation, and on the order of \$50,000 in annual operating costs. Emission data is collected automatically by smokestack sensors and relayed electronically to computers at the EPA, which also helps to keep administrative costs low.

Regulating the behavior of individuals is a different story, largely because of the sheer numbers involved. A year's worth of emission and operation data from all fossil-fueled boilers at electric generating plants in the United States (there are 2,500 of them) fits comfortably on a single spreadsheet. In contrast, there are 65 million U.S. households burning oil or natural gas to heat their homes. Or consider automobile emissions. With 253 million registered automobiles in the United States alone,

the costs of monitoring and enforcing automobile emission standards—quite apart from the costs of installing and maintaining emission reduction equipment—make up a significant component of the overall cost of regulating automobile emissions.

One way to keep administrative costs down is to focus on technology and fuel inputs rather than the performance of individual sources of pollution. This helps explain why command-and-control approaches tend to be prevalent in regulating home heating equipment, automobile emissions, and the like. Monitoring emissions from every household furnace would cost an astronomical amount, even before taking into account the administrative costs of levying an emission tax or the transaction costs involved in trying to institute a household-level market for pollution allowances. In such cases, technology standards (like those imposed on new oil and gas furnaces) or input standards (such as restrictions on how much sulfur and other contaminants can be present in fuel oil) make a great deal of sense.

A second issue related to monitoring and enforcement concerns the rates of compliance with environmental regulation. A sometimes overlooked advantage of market-based instruments is that they may *increase* compliance. For example, transferable fishing quotas essentially give ownership of the resource to the people working the fishery, giving them incentives to help police their fellow fishers. If I pay for valuable rights to harvest from a fishery, I have a concrete financial incentive to make sure you don't catch more than you are allowed.

In the area of pollution control, command-and-control regulations have been undermined by ubiquitous delays, extensions, negotiated agreements with polluters, and the like. For example, the 1972 Federal Water Pollution Control Act required the EPA to announce regulations on effluent standards within a year. The EPA had still not completed the task by 1977, the date that sources were supposed to be in compliance with the regulation. That same year, Congress pushed back the date of compliance to 1984. Command-and-control regulation also raises the stakes for polluters, because installing new technologies can cost hundreds of millions of dollars. This gives them an incentive to sue regulatory agencies in an effort to stall or overturn environmental regulation, creating a major hurdle to the successful and timely implementation of those laws.

In contrast, the very nature of market-based instruments makes compliance easier, because firms have the option of paying a tax or buying allowances rather than installing new equipment or changing their practices. Remember that the total amount of pollution depends on the

emission tax or the pollution cap. If some firms opt to comply by paying for their pollution instead of reducing it, others will respond by doing much more than they would under command-and-control. Therefore, the flexibility built into market-based instruments can make pollution control easier to achieve and enforce, without leading to any more pollution.

Linking Policies in Different Jurisdictions

Throughout this book, our focus has been on policy instruments that address emissions (or natural resources) within single political jurisdictions. For many environmental issues, such as local air or water pollution, that is an appropriate focus. Even in the case of acid rain, in which sulfur dioxide emissions from power plants in the United States affected the health of lake and stream ecosystems in Canada, the policy instrument itself—an emission trading program for SO₂—involved only U.S. power plants.

With global climate change as an increasing focus of environmental policy, however, the interactions *between* jurisdictions have gained greater importance. Perhaps the most direct form of interaction is linking between emission trading systems. Just as countries can combine to form a free trade zone (in which they freely trade goods and services across borders) or a currency union (in which they share a currency), countries with cap-and-trade programs can also link their programs by agreeing to accept each other's allowances for compliance purposes. (We will see an example of this in Chapter 10.)

From the perspective of economic efficiency, linking brings clear benefits: Trading between two jurisdictions will create a common allowance price, ensuring cost-effective abatement for the two jurisdictions as a whole (in much the same way as the single price arising from trading between two firms ensures cost-effectiveness). By the same logic, linking will lower the total cost of meeting the combined emission target of the two jurisdictions. As a result, linking offers another potential advantage for emission trading over command-and-control regulation—although, much like free trade policies, it also introduces potential political challenges.

Is Command-and-Control Ever Preferable?

Throughout our discussion of policy design, in Chapter 8 and on into the current one, we have emphasized the arguments in favor of market-based instruments. Although such approaches have many advantages—they can achieve an efficient level of pollution control, are cost-effective, and promote technological change—they are not panaceas. If you have been

reading carefully, you will have noticed that in some settings command-and-control policies are preferable.

For example, when hot spot problems are severe, command-and-control approaches may be the only feasible option. Regulating the disposal of toxic wastes requires restricting the amount of waste and the means of disposal at specific sites, not simply taxing firms for dumping waste or imposing a nationwide limit on how much is created.

Occasionally, a certain control technology may be so effective and widely available, and production practices within an industry so similar across firms, that requiring the installation of that technology makes much more sense than regulating emissions. Double-hulled oil tankers are a good example. All shipping firms use essentially the same technology (large oceangoing tankers), meaning that there is likely to be little variation in abatement costs—and concomitantly little gain from allowing firms to trade control responsibilities. (This effect is compounded by the very high costs of cleaning up oil spills after the fact and the difficulty of recovering those costs from the firms responsible.)

Finally, as we saw in the case of administrative costs, the number of regulated entities may make market-based approaches impractical. Consider automobile emissions, for example. Catalytic converters (required on new cars sold in the United States since the mid-1970s) represent a classic technology standard, but they also are much more sensible than trying to regulate emissions from individual cars. This is partly because it is much cheaper to require automakers to install catalytic converters than to monitor emissions and partly because the costs of abatement are unlikely to vary all that much between different car owners, limiting the gains from greater flexibility to comply with regulation. Even when automobile emissions are regulated, it is through performance standards based on periodic tests rather than by taxes or cap-and-trade (which would require information on actual emissions).

Of course, technological advances can overturn these calculations of costs and benefits. When we wrote the first edition of this book, it would have been astronomically expensive to monitor automobile emissions directly. But that is changing. Several states have piloted remote sensing of tailpipe emissions, paired with license plate recognition technology, tracking real-time emissions from cars as they drive on roadways, particularly in metropolitan areas that face expensive mitigation requirements because of their failure to attain the Clean Air Act's local air pollution standards. Virginia and Texas use these methods to identify high emitters and notify them of the need for additional testing. Ohio and Tennessee

have integrated remote sensing into their standard emission testing programs, exempting some drivers who pass remote tests on roadways from the usual required drive-in emission inspection. Perhaps in a later edition of this book we will be discussing successful automobile emission tax policies!

Conclusion

This chapter has focused on the choice between market-based instruments and command-and-control policies. We started by distinguishing between *goals* and *means* in environmental policy. This distinction is useful because it allows us to analyze the design of policy (should we require a certain abatement technology or create an allowance market?) independently from the policy target (how much pollution should we allow?)

Regardless of how the policy goal is set, market-based instruments have two substantial advantages over technology and performance standards. First, they are cost-effective: A tax or cap-and-trade program can (at least in theory) achieve a given goal at least total cost. They do this by ensuring that all firms end up with the same costs on the margin because all firms face the same incentives, in the form of the tax or the allowance price. In contrast, a uniform performance standard (and to an even greater degree a technology standard) is a one-size-fits-all approach that ignores the different opportunities available to different firms. Although a firm-specific performance standard would be cost-effective in theory, it is likely to be infeasible in practice because of the enormous informational requirements it imposes on the regulator.

The second major advantage of market-based instruments is that they promote the adoption and innovation of new technologies. If firms pay for every ton of pollution they emit, for example, they have a strong incentive to look for new and better ways to reduce pollution.

There are a number of other arguments for (and against) market-based instruments, of course. These relate to the possibility of hot spots, heterogeneity among regulated firms, and the costs of compliance and enforcement. Although market-based instruments are not the answer to all environmental problems, they do represent a very powerful and widely applicable component of the environmental policy toolkit. In the next chapter, we shall see how they have been implemented in the real world.

10

Market-Based Instruments in Practice

We have now seen how environmental degradation and excessive resource extraction often result from absent or incomplete markets and how the best solutions to these kinds of market failures may be market principles themselves. Now we look at some examples of how market-based policies have been used in the real world.

We begin by describing three cases in detail: the market for sulfur dioxide (SO_2) emissions from power plants in the United States, the tradable individual fishing quota (IFQ) system for New Zealand's fisheries, and municipal drought pricing of water resources in the United States. In each of these cases, we describe the background to the regulations, discuss their performance (both in meeting policy goals and in achieving cost savings relative to conventional regulatory approaches), assess their distributional effects, and touch briefly on questions of compliance and enforcement. We then describe a number of applications of market-based policies to environmental and natural resource management, including greenhouse gas emission trading in the European Union, water quality trading, pay-as-you-throw programs to manage household trash disposal, and "banks" for wetlands and endangered species habitat.

The purpose of the chapter is not to catalog every application of market-based approaches to environmental management and document successes and failures.¹ Instead, we demonstrate specific cases in which market principles have been applied by governments to correct market failures. In some cases, these attempts have had great success, in others less so. We hope to leave you with the capacity to think broadly and creatively about the ways in which market principles can be used to promote

environmental protection by aligning the incentives of firms and consumers with the interests of society.

The U.S. Sulfur Dioxide Market

One of the most successful market-based environmental policies has been the U.S. sulfur dioxide allowance trading program set up by the 1990 Clean Air Act. This program, more than any other, has shown that market principles can improve environmental regulation and has served as a model for other environmental markets.

To understand how the program worked, it helps to start by describing the run-up to the eventual legislation. Throughout the 1980s, Congress debated reauthorization of the Clean Air Act, which had last been updated in 1977. One impetus was the growing concern over acid rain, caused largely by sulfur dioxide from power plants in the Midwest. Many of these power plants had been built in the 1950s or 1960s. As a result, they had been grandfathered into the laws passed in the 1970s, which covered only new sources. Indeed, the continuing role of these large power plants in emitting pollution exposed a fundamental flaw in the earlier approach. The logic behind the decision to focus on new sources was simple enough: Installing pollution control equipment at a new source is much less expensive than retrofitting existing plants. Over time—so the theory went—the standards would cover a greater and greater fraction of electricity generation, as existing plants were retired and replaced with new ones. But by making new power plants much more expensive to build, Congress had unwittingly made it much more attractive to keep old power plants in service. As a result, plants that were originally scheduled to last 30 years were still going strong, with no retirement in sight.²

If Congress were to start regulating these older plants, how should it do so? Earlier clean air legislation (in 1970 and 1977) had adopted command-and-control regulations, first setting a uniform emission standard, and then imposing a technology standard requiring power plants to install scrubbers. By the late 1980s, however, the idea of emission trading was moving from academic journals to the policy world. Officials at the Environmental Protection Agency (EPA) had begun to introduce market principles, letting polluters offset increased pollution at one facility with reductions in pollution elsewhere. The new ideas were even gaining ground with some environmental advocates, particularly the Environmental Defense Fund (EDF), whose director, Fred Krupp, described market-based policies as a coming “third wave” of environmental policy.

The 1990 legislation resulted from an unlikely alliance. On one side was the administration of president George H. W. Bush, who had

campaigned in part on a pledge to be “the environmental president,” going so far as to film a campaign spot near heavily polluted Boston Harbor to undercut his rival, governor Michael Dukakis of Massachusetts. At the same time, however, Bush needed to promote his bona fides as a business-friendly Republican. On the other side of the alliance was EDF. The two sides essentially struck a deal. EDF would endorse (indeed, would help to write) legislation proposed by the Bush administration that enshrined a market-based approach to SO₂ control, helping to give the Bush initiative credibility among moderate environmentalists. In exchange, the administration proposal would set of goal of reducing SO₂ emissions by 10 million tons per year from 1980 levels by the year 2000, rather than a weaker target of 8 million tons.

Congress passed the resulting legislation (with some modifications) as Title IV of the 1990 Clean Air Act Amendments. It contained two particularly notable provisions for sulfur dioxide regulation. For the first time, Congress exerted federal authority over emissions from plants built before 1971. And in bringing these existing power plants under federal regulation, Congress adopted a novel tack: a system of tradable allowances (or cap-and-trade policy). In the first phase of the program, which lasted from 1995 through 1999, 263 generating units at 110 power plants were required to participate in the allowance market. Although these made up less than one fifth of the total number of fossil-fired generating units, they were by far the largest and dirtiest; environmental advocates referred to them as “The Big Dirties.” Total pollution in this phase was capped at roughly 6.3 million tons of SO₂ per year (an annual reduction of about 3.5 million tons). This cap was divided up into the same number of allowances, each corresponding to 1 ton of pollution, which were allocated for free to the regulated power plants. In the second phase of the program, starting in the year 2000, virtually all power plants above a certain size were brought into the allowance market, with an overall annual cap of 9 million tons.³

Power plants were allowed to bank their allowances for later use. For example, an allowance handed out in 1996 could be used in that year or saved for a later date. (The reverse was not true: Allowances could not be borrowed from future years.) This feature provides additional flexibility for regulated firms to reduce their pollution over time.

Performance

Through 2008, the SO₂ allowance market was widely hailed as a tremendous success. Market activity was substantial, with the annual volume of trades equaling or exceeding the total number of allowances allocated

each year. In June 2008 alone, 250,000 tons of emissions were traded in this market, with an end-of-month price of \$325 per ton. Several brokerage companies competed to track and arrange bilateral trades. Forward markets, loans, swaps, and other financial derivatives also sprang up. Meanwhile, the market not only succeeded in meeting the cap but achieved *more* abatement than required in the first phase of the program. These early reductions, made possible by the banking provision, were encouraged by the effective tightening of the cap starting in the year 2000. Looking ahead, electric utilities foresaw that allowance prices would rise (and abatement would become more costly) as allowances became scarcer in later years. In response, utilities as a whole abated more than they were required to—roughly 2 million more tons per year than required, in fact, amounting to about one third of the total allocation. After the start of phase 2, they started to draw down the resulting allowance bank. From an economic point of view, this early abatement was a good thing, because more benefits from lower pollution were enjoyed earlier than would have been the case under command-and-control regulation with a similar target.

Ironically, the success of the SO₂ market contained the seeds of its eventual demise.⁴ In light of increasing scientific evidence demonstrating the health effects of fine particulate matter associated with SO₂ emissions, as well as the lower-than-expected costs of abatement under the 1990 program, the EPA under president George W. Bush tightened the SO₂ emissions cap in 2005. Under the new regulation, midwestern states contributing to violations of the Clean Air Act fine particulate standards on the East Coast were required to surrender three permits for every 1 ton of SO₂ emissions. A legal challenge by states and electric utilities resulted in a 2008 federal court decision vacating this new approach, however, on the grounds that the EPA could not modify the existing interstate SO₂ market to meet stricter standards in the absence of legislation from Congress. After further revisions by the Obama administration and an additional court challenge, the EPA introduced a new program with state-level emission budgets (limiting the operation of the market) and a new regulation on mercury and air toxics that, in effect, required all remaining uncontrolled coal plants to install scrubbers.

Although the SO₂ market has essentially become defunct as a result of these regulatory changes, we can still assess its economic performance. Let's start by comparing the estimated costs and benefits. The major categories of benefits from sulfur dioxide abatement include lower incidence of sickness and mortality caused by urban air pollution, reduction in the

acidification of aquatic ecosystems, and improvement of visibility in recreational and residential areas. Economic analyses have found that the lion's share of the benefits turn out to come from health effects rather than the ecological impacts of acid rain. Indeed, one authoritative analysis found that on a per capita basis for the northeastern United States, reduced sickness and mortality accounted for more than 85 percent of the benefits. Moreover, total estimated benefits outweighed estimated costs by roughly an order of magnitude. At a national level, the same study estimated that health benefits to the United States as a whole were \$3,300 per ton of sulfur dioxide reduced, compared with costs of about \$270 per ton.⁵

An interesting irony arises here. The impetus for reducing sulfur dioxide pollution was concern about acid rain and the attendant damage to ecosystems, not about human health impacts. At the time the legislation was written and passed, the contribution of sulfur dioxide emissions to urban air pollution was not widely understood. Indeed, the trading program itself was dubbed the "Acid Rain Program." If the benefits had been limited to reducing the ecological impacts of acid rain, however, the costs would have been comparable in magnitude to the benefits, with little net gain to society. This is a happy case of a positive unintended consequence. From an efficiency perspective, the allowance trading program did the right thing for the wrong reason.

Another way of assessing the program's performance is suggested by our discussion in Chapter 9. Taking the goal of the policy as given, how well did the cap-and-trade policy that was actually used perform, relative to alternative policies that might have been used instead? Table 10.1 provides some answers to this question, at least for phase 1 of the sulfur dioxide trading program.⁶ Several counterfactual scenarios (what could have happened but didn't) are given, along with the baseline scenario corresponding to the actual cap-and-trade program. As the table shows, the cap-and-trade program was more costly than the theoretical minimum cost of achieving the same emission reduction. In other words, if electric utility managers had had perfect foresight about the prices of low-sulfur coal and SO₂ allowances and had made cost-minimizing decisions based on that information, total abatement costs would have been less than half of what they actually were. This is a useful reminder that policies in the real world never work as perfectly as they do in theory.

However, it is more informative to compare the outcomes of the actual policy with estimates of what the outcomes would have been from other feasible policies, not just the theoretical best-case scenario. The table also shows that the cap-and-trade program was significantly less costly than a

Table 10.1

Estimated Costs of Various Alternative Policies to Achieve the Same
Emission Reduction as Phase I of the 1990 Clean Air Act Amendments

<i>Scenario</i>	<i>Estimated annual cost (millions)</i>	<i>Cost difference from baseline (millions)</i>	<i>Cost increase</i>
Theoretical least-cost outcome	\$315	-\$432	-57%
Baseline cap-and-trade program (actual policy)	\$747	—	—
Uniform emission rate standard	\$900	\$153	20%
Technology standard	\$2,555	\$1,808	242%

Source: Estimates taken from Nathaniel O. Keohane, "Cost Savings from Allowance Trading in the 1990 Clean Air Act," in Charles E. Kolstad and Jody Freeman, eds., *Moving to Markets in Environmental Regulation: Lessons from Twenty Years of Experience* (New York: Oxford University Press, 2007), 194–229.

uniform emission rate standard on the same power plants. The estimated savings of just over \$150 million represent about one fifth of the costs of the actual trading program. These cost savings, in turn, pale next to the cost savings relative to a technology standard requiring utilities to install scrubbers. Such a policy (along the lines of the 1977 Amendments) was actively considered in the debates leading up to the 1990 legislation. If it had been chosen, the costs of achieving the same amount of abatement would have been more than three times as high as they were: nearly \$2.6 billion per year, compared with a baseline annual cost of \$747 million. These large cost savings are due in large part to wide variation in abatement costs between electric utilities. For example, much of the steep reduction in sulfur dioxide emissions under the program has been due to greater use of low-sulfur coal from the Powder River basin in eastern Wyoming. Transportation (mostly by railroad) accounts for most of the cost of extracting and delivering coal to midwestern power plants. As a result, a power plant's geographic location is a key determinant of its abatement cost.

We can also frame cost-effectiveness another way: in terms of the pollution savings. Recall the role played by the environmental group EDF in supporting the allowance market. To a group like EDF, the cost savings from market-based policies offered an opportunity to secure greater environmental protection. The standard case made by economists (indeed, the case we made in Chapter 9) emphasizes cost-effectiveness: Market-based policies can achieve a given policy goal at less cost than

command-and-control. But such policies can be portrayed equally well another way: For a given total cost, a cap-and-trade policy allows more abatement. In the case of sulfur dioxide, the use of a cap-and-trade program translated into roughly 10 percent more abatement being done than could have been done for the same total cost with a uniform emission rate standard.

Finally, we can consider the effects of the cap-and-trade program on technological change. In line with what we would expect, based on Chapter 9, allowance trading seems to have boosted the incentive for electric utilities to adopt lower-cost technologies. Among electric utilities regulated by conventional emission rate standards, the cost of scrubbing appears to have had little impact on the decision to install a scrubber. In contrast, cost was an important consideration at power plants included in the allowance trading program. Meanwhile, evidence from patent data suggests that the SO₂ trading program spurred firms that design and build scrubbers to focus more on raising removal efficiencies than they had under previous command-and-control regulations (which did not reward increases in abatement beyond that required by the uniform standard). Both findings are consistent with theoretical predictions that market-based policies promote greater technological innovation.⁷

Distributional Implications

In the case of a cap-and-trade program like that for sulfur dioxide, we can assess the distributional impacts along two very different dimensions: How were the allowances allocated, and where did the pollution end up?

At a broad level, the more than 6 million allowances (during phase 1) and the 9-million-ton cap in phase 2 were given away for free to existing power plants. An alternative allocation mechanism would have been to auction the allowances.⁸ As we discussed in Chapter 8, an auction would have raised government revenue that could have been used to offset distortionary income and sales taxes. Instead, that revenue was handed to the electric power plants that participated in the market. In effect, the 1990 Clean Air Act Amendments created a new scarce resource and gave away the rents. A quick calculation can give you a sense of the total value of the scarcity rents involved. The average allowance price during phase 1 was \$135; multiplying by 6.3 million allowances implies a total annual value of \$850 million. We saw earlier that the estimated annual cost of abatement was about \$750 million, implying rents on the order of \$100 million a year.

What about the consequences of allowance trading for the distribution

of pollution? As we noted in Chapter 9, a potential concern with market-based instruments arises when pollutants are not uniformly mixed. As it turns out, SO_2 is not uniformly mixed. Emissions from power plants in Ohio and Indiana travel downwind to the urban centers of the Northeast and contribute to acid rain in the Adirondacks. Emissions from power plants in Maryland or Delaware, on the other hand, are much more likely to blow out to sea. In theory, the SO_2 trading program could have made matters worse—even as it lowered total pollution—if it had led to the reallocation of pollution from Delaware to Ohio. Indeed, this issue was voiced by critics of the program.

Retrospective analysis suggests that much larger net gains from the SO_2 trading program might have been possible, had regulators developed a *differentiated* trading program based on the marginal damages from emissions at each source location rather than allowing trading on a ton-for-ton basis.⁹ However, although some redistribution of pollution did occur, by far the largest reductions in emissions took place in Ohio, Indiana, and other midwestern states—precisely the states whose emissions were of greatest potential concern.¹⁰ The reasons for this are varied, but three stand out. First, the required abatement implied by allowance allocations was largest in these regions, because they were also the biggest emitters before the program started. In other words, even if electric utilities had simply used up the allowances given to them, without trading with each other, abatement would have been greatest in these midwestern states. Second, abatement was cheaper for many midwestern plants than for other power plants. Relative to power plants in eastern states like West Virginia and Georgia, midwestern plants have easy access to low-sulfur Wyoming coal. Finally, state-level regulations may have played a role. Even after the introduction of the national trading program, some states still imposed limits on pollution from individual power plants, in order to protect local air quality. And states such as Ohio and Indiana sought to protect their in-state high-sulfur coal industries by encouraging power plants that participated in the allowance market to install scrubbers. As this discussion makes clear, there was nothing in the cap-and-trade program itself that prevented hot spot problems from intensifying. Rather, a combination of factors, including luck, made the difference.

Compliance and Enforcement

Compliance and enforcement did not pose major hurdles for the SO_2 trading program. We have noted already (in Chapter 9) that the costs of monitoring, though real, were roughly two orders of magnitude less than

the costs of abatement. Moreover, these monitoring costs would have been incurred by any approach based on emissions, such as a uniform performance standard.

Observers typically note that compliance with the trading program was “100 percent,” meaning that participating utilities indeed retired as many allowances as needed to cover their emissions. Because the program allowed for a truing-up period early each year (firms had until March 1 to ensure that their allowance holdings covered their previous year’s emissions), and utilities were fined \$2,000 per ton for noncompliance (far above prevailing allowance prices), this high rate of compliance is hardly surprising. Nonetheless, the smooth experience with the allowance trading system may have broader implications for market-based policies.

Individual Tradable Quotas for Fishing in New Zealand

In 1986, New Zealand set up what has become the world’s largest market for tradable IFQs. At the time, concerns about overfishing created a sense of crisis, especially for an island nation heavily dependent on its natural resources. A market-based approach to fishery management also appealed to the government at the time, which pursued a larger agenda of market reform, including privatizing industries in some sectors and reducing subsidies in others.

Each year, the government sets a total allowable catch (TAC) for each species–region, based on a biologically determined maximum sustainable yield (MSY).¹¹ Quotas are freely allocated based on a fisher’s average catch in a set of preceding years. Fishers can buy and sell IFQs, which represent the right to fish a percentage of the TAC in perpetuity.¹² With few exceptions, these quotas cannot be traded across regions, species, or years. When the system began, it covered twenty-six different species. By the mid-1990s, IFQ markets covered 85 percent of the commercial catch within the exclusive economic zone extending 200 miles out to sea from New Zealand. By 2009, ninety different species were included in the program.¹³ Coastal waters are spatially divided into quota management regions, with markets in each region for each relevant species, generating hundreds of separate markets.

Performance

The first question about the performance of an IFQ market is: Did it help reduce overfishing and restore the stock? From its beginning, the program mandated a sharp reduction in the catch. In 1986, the TACs were only a quarter to three quarters of what they had been before the

Market-Based Policies for Greenhouse Gas Emissions

Greenhouse gases present an interesting case for market-based policies, and not only because of the increasing focus on climate change. Several factors suggest that market-based instruments should be a natural fit. CO₂ emissions can be readily measured, either directly (from smokestacks) or indirectly (because vehicle emissions can be calculated from the carbon content of the fuel). Marginal abatement costs vary widely across sectors and different types of greenhouse gases. Hot spots are not a concern. Creating incentives for technological innovation is crucial. At the same time, the fact that greenhouse gas emissions come from virtually every economic sector, in every country around the globe, creates its own design challenges.

By 2008, market-based instruments for greenhouse gases seemed ascendant. Emission trading was the centerpiece of the 1997 Kyoto Protocol, which took effect in 2008 on emissions from all developed countries except for the United States. Interest in the global carbon market was booming, prompted by a Kyoto program called the Clean Development Mechanism, under which emission reduction projects in developing countries could generate “offset credits” to be sold in compliance markets. The European Union’s cap-and-trade program for carbon dioxide was entering full operation, having completed a 3-year pilot phase. And both candidates in that year’s U.S. presidential election, Democrat Barack Obama and Republican John McCain, were calling for an economy-wide cap-and-trade program that would reduce U.S. greenhouse gas emissions 80 percent by the year 2050.

Then the momentum stalled. Comprehensive climate legislation sponsored by Henry Waxman and Ed Markey passed the House of Representatives in 2009 but never came to a vote in the Senate. Opponents of climate policy declared that cap-and-trade was dead. Russia’s Kyoto target (which was calculated from a 1990 baseline) ended up far above its emissions, thanks to the severe economic downturn following the collapse of the Soviet Union; the resulting “hot air” essentially ended the prospects for robust trading under the Kyoto system. Even the EU’s system has run into criticism; as a result of the economic crisis of 2009–2012, large supplies of Clean Development Mechanism credits, and the effects of complementary policies to promote renewable electricity generation, the price of EU allowances fell below €5. As recently as April 2013, the *Economist* magazine ran a headline asking, “Carbon Trading: ETS, RIP?”

To paraphrase Mark Twain, the death of cap-and-trade (and other market-based instruments for greenhouse gases) was widely exaggerated. As the map in figure 10.1 illustrates, dozens of political jurisdictions around the world have implemented emission trading systems or carbon taxes, with others poised to follow suit.¹⁴

Market-Based Policies for Greenhouse Gas Emissions *continued*

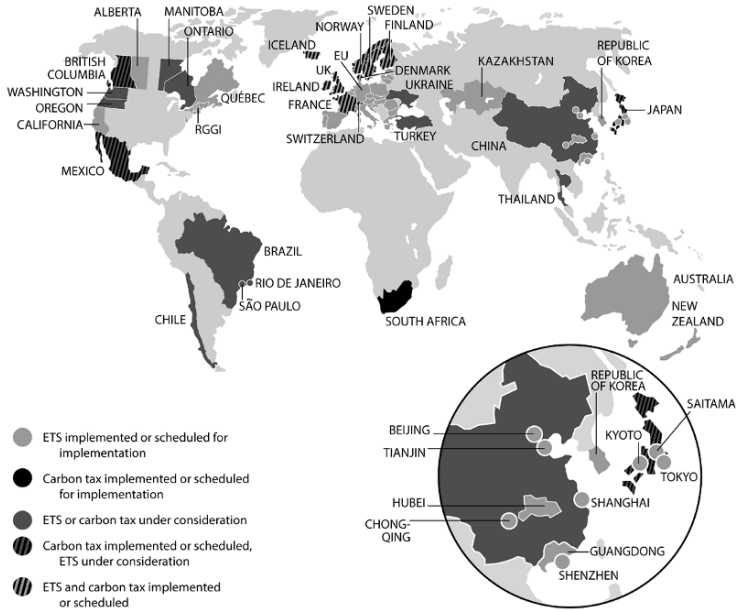


Figure 10.1 A map of current and planned carbon pricing programs around the world, including emission trading and carbon taxes. Note that soon after this map was published, Australia abandoned its emission trading program. *Source:* World Bank, *State and Trends of Carbon Pricing 2014* (Washington, DC: World Bank, 2014).

The European Union's Emissions Trading System (EU ETS), discussed in the text, is still the largest emission trading system in the world (as of this writing). But other markets are worth mentioning as well. California has a robust emission trading system with the broadest coverage of any system in the world, including not only the electric power sector and industrial facilities but also transportation fuels, adding up to nearly 85 percent of the state's greenhouse gas emissions. California has also fully integrated its market with that of Quebec, creating the first internationally linked emission market (recall our discussion of linkage in Chapter 9). Elsewhere in the United States, the nine northeastern and mid-Atlantic states that are members of the Regional Greenhouse Gas Initiative (RGGI) have had a cap-and-trade program in place for CO₂ emissions from the electric power sector since 2008, and they successfully revised and strengthened the program in 2013.

Market-Based Policies for Greenhouse Gas Emissions *continued*

Perhaps most importantly, market-based policy instruments are being put in place in emerging economies. In 2015, trading began in South Korea's emission trading system, whose first phase (from 2015 to 2017) covers roughly 60 percent of the country's annual greenhouse gas emissions, including power generation, steel, petrochemicals, and other sectors. China has implemented pilot emission trading programs in seven cities and provinces, which have a combined population of 250 million people and account for more than 25 percent of the country's gross domestic product. The country has announced plans to implement emission trading at a national level in 2017. Mexico instituted a partial carbon tax on fossil fuels as part of a sweeping energy reform package passed in 2012; the tax, which is roughly US\$3 per ton of CO₂ emissions, varies by fuel (and exempts natural gas) but is expected to raise roughly US\$2 billion in annual revenue. Chile enacted legislation to put a US\$5 per ton tax on CO₂ emissions from the power sector starting in 2018 and has discussed a possible later transition to an emission trading system. Other countries, most notably South Africa, have also announced plans to proceed with a carbon tax.

program started. (Where a particular fishery fell within that range depended on biological status and management goals.)¹⁵ Available evidence suggests that few if any species populations are worse off as a result of the IFQ system (especially in comparison to the alternative of continued open access), and some show significant positive signs of recovery. One analysis of the rebound of regulated fish stocks in New Zealand examined 149 of the 179 individual stocks (defined by species and region) that were governed by the system in 1993.¹⁶ Only thirteen of these stocks were estimated to be smaller than the stock that would generate the maximum sustainable yield. Another thirteen were estimated to be at or above this level, and the status of the remaining stocks was not determined. TACs for the fish stocks deemed to be below the MSY-supporting level have been reduced over time. This performance stands in contrast to the well-publicized crashes of fish stocks in other parts of the world over the same period of time.

As with the SO₂ cap-and-trade program, we can also assess the performance of the New Zealand IFQ program relative to alternative command-and-control policies, such as nontradable permits. No study has estimated the cost savings directly. But an analysis of market activity and IFQ prices can tell us a great deal about how well the market is performing, in terms of its ability to minimize costs.¹⁷ A recent study of the New Zealand IFQ

market found that 70 percent of quota owners had taken part in a market transaction by 2000. Quota sales were highest early in the program, as the initial allocation was redistributed among fishers. This is consistent with increased efficiency: Presumably, more profitable producers bought their way into the program, and less profitable producers sold their quotas and exited.

The variation of prices around their annual mean can also tell us something about how smoothly the market is operating. In a well-functioning market, the price of an IFQ would reflect the present value of the stream of expected future net revenue from future harvests. Economic theory would predict a single price within a species-region, just as in any well-functioning competitive market. In reality, transaction costs and other imperfections can lead to price variation even for an identical good. (Think, for example, of gasoline prices, which often vary widely within a neighborhood despite the fact that the underlying product is essentially identical.) In the New Zealand IFQ market, quota prices varied by as much as 5 percent around their mean value, comparable to price dispersion in other real-world markets. Moreover, price dispersion has been trending downward since the market's inception in 1986, suggesting that market frictions are diminishing.

Trends in quota prices give a measure of the long-run performance of the program. A prime goal of fishery management, after all, is to raise the economic value of the fishery. Remember the basic intuition that IFQs work by establishing property rights to the fishery. Just as house values rise with neighborhood revitalization, the value of a fishing quota should rise over time as the health of the fishery improves. This is exactly what appears to be happening in the New Zealand market. Quota prices, adjusted for inflation and controlling for a range of outside factors, have risen at a rate of 5 to 10 percent over the course of the program, with higher rates observed in markets that saw greater initial reductions in the allowable catch and correspondingly higher consolidation of quota ownership.

Distributional Implications

One of the strongest objections to IFQ markets is that they encourage consolidation of the fishery, in other words, a reduction in the number of fishers operating in the market. This occurs partly because of the overall decline in the allowable catch that often accompanies a new market, just like under any fishery regulation. But consolidation tends to be greater under an IFQ system, which allows less profitable fishers to exit the industry, while more profitable fishers enter or expand to capture a greater

share of the market. From an efficiency perspective, the exit of high-cost fishers and their replacement by low-cost fishers is exactly the purpose of the market. After all, falling costs boost the economic value of the resource, leading to higher profits for the fishers who remain in the industry, lower prices for consumers, or both.

The implications of consolidation for equity are mixed, however. Critics argue that IFQs give an unfair advantage to large firms that can take advantage of scale economies to reduce their fishing costs below those of small-scale fishers.¹⁸ They also lament the decline of a way of life that in many cases has been handed down over several generations. On the other hand, it is not clear that the distributional implications are negative. After all, any fishers who exit the fishery do so by choice and earn a profit from selling their quotas as a result. In contrast, restricting harvests without allowing trading may deny fishers the choice of whether to remain or exit, and those who exit do so without compensation. Indeed, any adverse distributional impacts of IFQ systems must be compared against the consequences of some other form of restriction on harvests (which will also harm fishing communities) or of continued open access (which is likely to bring on the collapse of the fishery).

IFQ markets can be designed to address some of these concerns. For example, the initial allocation of quotas provides a way to address distributional issues. Recall from our discussion of tradable pollution allowances in Chapter 9 that the *equilibrium* allocation of quotas in a market is essentially independent of the initial allocation (as long as transactions costs are low). If equity is a concern, fishery managers can hand out a disproportionate share of quotas to small-scale fishers or those who have had a long history fishing the resource. If consolidation per se is a concern, it can be limited by policies that impose maximum quota holdings or the like. However, unlike changes to the initial allocation, such constraints on the final allocation of quotas come at a substantial cost in terms of efficiency.

How does all this play out in the case of New Zealand? Some consolidation has certainly happened. Between 1986 and 2003, the number of quota owners in the IFQ system declined by 37 percent, primarily in species-regions most severely overcapitalized before the introduction of markets.¹⁹ Many of those who have exited the markets are small-scale fishers, but many of these remain as well.

Another distributional issue in New Zealand's IFQ system concerned its impact on the traditional fishing communities of native New Zealanders, the Maori. The initial design of the IFQ market excluded the Maori

altogether. In response, the Maori successfully sued the government in the late 1980s, arguing that they had been disenfranchised. As a result of this litigation, the government allocated shares of the TAC to Maori communities in 1992 (10 percent of all existing IFQs, plus 20 percent of the TAC for any new fish stock regulated by the program) and provided funding for the purchase by Maori of one-half of New Zealand's largest fish company.²⁰

Compliance and Enforcement

Traditional fishing regulations are famously difficult to monitor and enforce. For example, limiting fishing seasons or particular fishing areas results in steep increases in the concentration of effort within allowable times and places. Limits on the types of allowable fishing technologies create strong incentives for the development of alternative fishing gear that is equally, and in some cases even more, productive. Gear restrictions require on-boat inspections, and the enforcement of fishing seasons and area restrictions requires that regulators monitor fishing activity at sea.

Like traditional fishery regulation, market-based regulation requires monitoring and enforcement to ensure compliance. Within New Zealand's IFQ program, quota holders pay levies per metric ton of quota share to support the cost of managing and administering the program, including enforcement costs.²¹ Enforcement under an IFQ program is mainly a matter of auditing paperwork and the exchange of fish between boats and fish purchasers where fish are landed—easier than monitoring at sea. Penalties for violating the IFQ regulations in New Zealand include forfeiture of fishing quotas, seizure of property, and exclusion from the fishing industry.

In comparison to what may occur under traditional regulations, fishers' attempts to maximize the value of their catch, given their quota, under a market-based system can lead to practices such as high grading and price-dumping. High grading involves the discard of lower-valued members of the quota species (for example, smaller fish) so as to fill the quota with the most profitable fish. High grading of snapper was a problem in the early years of New Zealand's IFQ program, but vigorous enforcement has ended the practice.²² Price-dumping occurs when fishers discard their catch when fish prices fall, to leave room to fill their quota on more profitable days when prices are higher. It is difficult to determine how common such practices are. One analyst surveyed New Zealand fishers in 1987, at the start of the program, and found that 40 percent thought

that enforcement of the IFQ system was an important problem, and 66 percent were concerned about high grading. By 1995, however, the comparable figures were 21 and 25 percent, respectively.²³

A related concern involves the effects of fishery management on other marine life, a problem known as bycatch. Fishers concentrating on a particular valuable species often end up catching other unwanted fish or marine animals in their nets. (The most famous example is the dolphins caught in seine nets formerly used to catch tuna in the Eastern Tropical Pacific.) There is reason to believe that the comprehensive nature of the New Zealand IFQ system may indirectly alleviate the bycatch problem. The quota markets cover so many species that fishers have taken significant measures to reduce bycatch, in order to avoid the need to purchase additional quota or pay fees to regulators for species caught outside a fisher's quota portfolio. Nonetheless, bycatch is a significant issue in some New Zealand fisheries. The important question is how bycatch under a market-based system compares to that which we would observe under traditional fishing regulation. As yet, the data from New Zealand are insufficient to make this comparison.

Municipal Water Pricing

Cities in arid U.S. states such as Texas and California have struggled to manage water scarcity in the face of growing populations, high household water demand (for swimming pools and landscaping, among other things), and the increasing cost of acquiring and developing new water supplies. During droughts, cities typically implement voluntary or mandatory limits on residential water consumption. For example, cities may restrict certain activities such as watering lawns, or require homeowners to install water-saving devices such as low-flow shower heads. As you have probably realized by now, these kinds of blanket limitations on water consumption are not how an economist would approach the problem of managing scarce water resources. What do market-based solutions look like in this context?

A key economic insight from our discussion of natural resources in Chapters 6 and 7 is that the price of a resource should reflect its scarcity. Setting higher water prices during droughts (when scarcity is greater) could achieve the same reductions in water use as the kinds of policies mentioned earlier, at lower total cost. Moreover, such prices promote the efficient allocation of water among competing demands. Higher prices would ensure that scarce water flowed to the highest-valued uses. Relative to one-size-fits-all restrictions, consumers with high willingness to

pay for water would end up consuming more, and consumers with low willingness to pay would cut back on their use. We have emphasized the critical role played by cost heterogeneity in pollution control or fishery management. In the context of water pricing, benefit heterogeneity plays an analogous role. The greater the differences between water consumers and the more varied their marginal willingness to pay for the resource, the greater the gains will be from allocating water by prices or markets.

Performance

As a baseline for assessing the performance of drought pricing, let's start by considering the effectiveness of command-and-control alternatives. Here, the evidence is mixed. In general, requiring customers to install specific water-conserving technologies does reduce consumption—but typically by much less than the manufacturing specifications for the conservation technology would predict.²⁴ (One possible reason is that consumers change their behavior after installing the new technologies, such as starting to take longer showers after installing a low-flow showerhead or doing more loads of laundry after they buy a water- and energy-efficient front-loading machine.) A comprehensive study of conservation programs in California found that public information campaigns, retrofit subsidies, water rationing, and water use restrictions all contributed to reductions in residential water use, with the more stringent, mandatory policies having stronger effects than voluntary policies and education programs.²⁵

Price increases have been used infrequently as a drought management tool. During an extended drought in California from 1987 to 1992, a handful of municipal water utilities implemented price increases to reduce water demand, achieving aggregate demand reductions of 20 to 33 percent with very substantial price increases.²⁶ In principle, it would be straightforward to design a drought pricing system that would result in the same aggregate water savings as an existing nonprice approach. The key piece of information needed in each case would be the price elasticity of water demand, a measure of how sensitive consumers are to changes in the price of water.²⁷ Estimates of residential water price elasticities in the short run usually range from -0.3 to -0.6 , meaning that a 10 percent price increase can be expected to reduce demand by 3 to 6 percent.

In the absence of extensive empirical experience, analysts have used data on actual water use to run simulations of what would happen under hypothetical (but plausible) market-based policies. A recent study of thirteen California cities found that under a wide range of assumptions, a modest water tax (a price increase) would be more cost-effective than a

technology standard (a mandatory low-flow appliance regulation).²⁸ Another study of eleven urban areas in the United States and Canada compared residential outdoor watering restrictions with drought pricing.²⁹ For the same level of aggregate demand reduction as that implied by a regulation allowing households to use water outdoors (for watering lawns and washing cars) only two days per week, the establishment of drought pricing in each city would result in welfare gains of approximately \$96 per household per summer drought. This represents about 29 percent of the average household's total annual water bill in this study.

In the long run, higher prices for water will lead to land use patterns, investments, and consumption decisions that take account of water scarcity. For example, we would expect households to plant fewer green lawns and install front-loading washing machines (which use much less water than top-loaders) more often in cities where water prices are high. As in the case of market-based policies to reduce air and water pollution, prices provide a strong incentive for technological change that lowers the marginal cost of water conservation.

Distributional Implications

The main distributional concern with a price-based approach to urban water management arises from one of the central features of a market. To an economist, one of the virtues of markets is that they allocate resources according to who is willing to pay the most for them. But willingness to pay strikes some people as an unfair criterion for allocation, because it is strongly influenced by ability to pay. What you are willing to pay for something depends in part on how much money you have to spend. (Recall our discussion of willingness to pay versus willingness to accept in Chapter 3.) This sense of unfairness may be especially acute when we are dealing with resources that satisfy basic needs, such as water for drinking and bathing. These distributional impacts are illustrated by the results of the study of U.S. and Canadian cities mentioned earlier. Not surprisingly, raising prices during a drought, rather than restricting consumption, would result in more water being consumed by wealthier households with large lots and less by poorer households. In welfare terms, as well, poorer households would be hit harder: The reduction in consumer surplus as a percentage of income is larger for low-income households.

Concerns about distributional equity could be addressed by combining drought pricing with income transfers. Higher water prices would yield substantial profits for the utilities, profits that would have to be returned to consumers in some form (because water utilities tend to be

subject to strict price regulation and oversight). In the case of residential water use, higher prices during droughts could be offset by rebates (unrelated to water use) to low-income households, which could appear automatically on residential water bills. (Note the parallel with the use of IFQ allocations to meet distributional equity goals in the case of fishery management and with the use of air pollution allowance allocations to meet political and distributional goals.)

Compliance, Monitoring, and Enforcement

Applying market principles to urban water management raises fewer concerns for monitoring and enforcement than do existing command-and-control approaches. Under outdoor watering restrictions during a drought, utilities rely on neighbors to report illegal watering, a notoriously ineffective system. Additionally, requirements for the installation of indoor water-saving fixtures cannot be monitored and enforced without utility representatives entering private homes. In contrast, drought pricing involves simply changing the price of water. Because almost all households pay according to how much water they use, and their water consumption is already metered, there's no need for additional monitoring and enforcement. Indeed, cheating in the drought pricing scenario would require households to figure out how to consume water off the meter. Of course, where metering is not prevalent, a drought pricing policy would require meter installation for all users, but so would any attempt to charge people according to the amount of water they consume.

The European Union's Emissions Trading System

As we discussed briefly in Chapter 8, the European Union Emissions Trading System (EU ETS) is the world's largest emission trading program, covering more than eleven thousand power stations, factories, and other stationary facilities (as well as civil aviation) accounting for roughly 50 percent of CO₂ emissions and 40 percent of total EU greenhouse gas emissions.

The EU ETS was implemented as a central means of complying with the greenhouse gas emission targets taken on by the EU Member States under the Kyoto Protocol signed in 1997.³⁰ By the early 2000s, the United Kingdom and Denmark had each begun to implement their own cap-and-trade programs (a voluntary program in the case of the United Kingdom, a cap on emissions from the electricity sector in Denmark), triggering concerns about a patchwork of emission trading systems within the continent. (Attempts to establish an EU-wide carbon tax in

the early 1990s had failed for legal and institutional reasons: EU law requires unanimity on the part of the member states for fiscal decisions, and this proved impossible to achieve. Many Member States resisted the idea, motivated in part by the desire to retain authority over their own systems of taxation.)

The system has been implemented in phases. After an initial pilot phase from 2005 to 2007, the EU ETS was launched in earnest in 2008, with this second phase lasting until 2012 (coinciding with the first compliance period under the Kyoto Protocol). Phase 3 started in 2013 and will run to 2020. In addition to retiring allowances (known as European Union Allowances [EUAs]), regulated entities may comply with their obligations by using “offset” credits generated under the Kyoto Protocol’s Clean Development Mechanism, which allows projects in developing countries to receive credits for estimated emission reductions relative to business as usual. After the initial pilot phase, full banking of allowances has been allowed.

At the most fundamental level, the EU ETS has met its goal: Total emissions among covered facilities have been less than the cap. Determining the amount of abatement achieved as a result of the system has been surprisingly difficult, however, partly because of a paucity of good data but mostly because of the challenge of disentangling the effects of the EU ETS from the global recession in the wake of the 2008 financial crisis. The second phase (2008–2012) of the EU ETS coincided with a deep decline in economic activity, and this continued into the first years of the third phase (2013–2020); emissions over this period would have fallen even without the EU ETS, complicating attribution. Nonetheless, analyses that have sought to estimate a “counterfactual” level of emissions have generally concluded that the EU ETS achieved reductions on the order of 3 percent of emissions (roughly 60 million tons per year) in the first 18 months of phase 1, when allowance prices were in the range of €30–15 per ton, with somewhat less abatement in phase 2. Most of the emission reductions appear to have occurred in the electric power sector, as a result of switching from coal to natural gas.

Another important measure of performance is innovation. As we noted in Chapter 9, economic theory suggests that putting a price on emissions should encourage firms to increase investment in research and development for new technologies. A careful study of innovation under the EU ETS did indeed find an increase in innovation (as measured by patents for low-carbon technologies) among firms that were regulated by the trading system. However, because the set of firms regulated by the EU ETS

accounted for a small share of overall patenting activity, the overall effect on low-carbon innovation has been small.³¹

Not surprisingly, one of the most politically contentious issues has been the allocation of emission allowances. In the first two phases, a large proportion of allowances were allocated for free to covered entities, including electric utilities. In countries that had deregulated their wholesale electricity markets, power generators were able to pass on the opportunity cost of emission allowances in the form of higher electricity prices even though they received the allowances for free. Although the resulting windfall profits were limited to the electric power sector and concentrated in a few countries, they generated predictable outrage. Partly as a consequence, in phase 3 the EU began using auctioning to allocate most allowances, eliminating the windfall profits problem. However, industries that face foreign competition from countries without comparable climate policies will continue to receive some free allocation, motivated in part by concerns about “carbon leakage,” the possibility that stringent climate policies in the EU could drive industries to move their operations overseas, increasing emissions elsewhere and undermining the gains from the EU’s actions.

Another focus of attention has been the price of allowances. In the early years of the program, allowance prices were quite volatile, largely as a result of how the pilot phase was designed. Because EU Member States lacked actual data on emissions before the EU ETS began, they set their initial emission caps based on estimates of what emissions were likely to be; when new data became available showing that the system was over-allocated, the allowance price dropped sharply from more than €30 to around €15 per ton. In addition, because the EU had deliberately specified that allowances from the pilot phase could not be used in subsequent years, the price inevitably collapsed to zero by the end of the period.

Since the start of phase 2, allowance prices have been less volatile, but they remain lower than initially expected: After staying around €15 per ton for most of 2009 to 2011, the price fell to less than €5 in 2013 and remained below €10 per ton in early 2015. The low prices are attributable to a number of factors, including the economic recession (real economic output in the EU as a whole fell by more than 4 percent from 2008 to 2009 and did not fully recover to its pre-recession level until 2014), the availability of a large supply of offset credits available under the Clean Development Mechanism and an absence of significant demand for those credits outside the EU, and the effects of other policies such as “feed-in tariffs” to subsidize renewable electricity generation.

What should we make of this price history? On one hand, the low price of EUAs has contributed to the modest levels of abatement and innovation discussed earlier. In particular, allowance prices have been too low to support the commercial deployment of major new abatement technologies, such as technologies to capture and store carbon dioxide emissions from power plants. Those concerns have led many observers to suggest that a higher price would have been preferable.

On the other hand, one of the causes of the lower-than-expected allowance price was the economic downturn after the financial crisis. Less economic activity means lower emissions, which in turn means lower demand for allowances and a lower price. As the EU's experience illustrates, allowance prices in emission trading systems are procyclical: They will be lower during periods of low economic growth. Although this feature may be frustrating for policymakers (and environmental advocates) hoping to see more significant abatement, it cushioned the effect of the EU ETS on regulated sectors during a severe downturn and probably helped make it more politically resilient. An emission tax that had remained high during the deepest recession in decades would have imposed much higher costs on a fragile economy and might not have survived politically. (Recall our discussion about setting prices versus setting quantities in Chapter 8.)

At any rate, the lower-than-expected prices have been a source of concern within the EU and have led to a number of reforms. The first was simply to tighten the cap: In 2014, EU leaders agreed on a more stringent 2030 target, which also implies steeper annual reductions after 2020. In addition, the EU delayed releasing some allowances into the market for a few years, a move known as backloading. (Economic theory suggests that backloading by itself should not affect the price of allowances because it changes only the timing of allowance releases into the market and not the overall amount, and there is already a large bank of allowances. Indeed, allowance prices do not appear to have risen significantly as a result.) Finally, the EU has proposed a measure called a market stability reserve, to be implemented no later than 2021, when phase 4 begins. Under this approach, allowances will be withheld from auction and placed into a reserve if the surplus of available allowances rises above a pre-specified level; allowances will be released from the reserve and added to auctions if the surplus falls below a lower bound. (This approach can be thought of as analogous to how the Federal Reserve of the United States seeks to influence interest rates by buying or selling Treasury bonds to influence the money supply.) It will be interesting to see how these reforms play out in the EU ETS over the coming years.

Water Quality Trading

There are about three dozen water pollution trading programs (usually called water quality trading programs) in place around the world, with all but a handful of these in the United States.³² These cap-and-trade programs are small and low-profile relative to their counterparts for air pollution, discussed at length in this book. In contrast to the SO₂ trading program or the CO₂ trading programs in place in the United States and elsewhere, water quality trading programs typically involve bilateral trades: Sources of the targeted water pollutant negotiate directly with each other (with a regulator's careful supervision) rather than through a traditional market with a single, well-known price. The higher transaction costs that result from this style of trading may be one reason for the relative lack of success in water quality trading, relative to programs for air pollution. For some programs, an intermediary (either a state regulatory agency or a specially designed clearinghouse institution) may convert the abatement activities of diffuse nonpoint sources of pollution (such as farms that implement changes in their land management practices, like planting buffer strips that separate livestock or crop fields from streams) into a uniform credit currency to be purchased by point sources (such as municipal sewage treatment plants or industrial point sources of pollution). To date, two active programs, Australia's Hunter River Salinity Trading program and the Pennsylvania Nutrient Credit trading program, have established true exchange markets, where buyers and sellers trade uniform credits at transparent prices.

The Hunter River program has been in operation since 2004, reducing the concentration of salts in the river from agricultural irrigation, disposal of brine from coal mining, and water diversions from cooling in electricity generation (which concentrates salts in the water replaced in the river). The program restricts saline discharges from these sources to periods when the river's flow is high and to concentrations compatible with each facility's credit allocation. If discharges exceed a facility's credits, they may purchase credits from other facilities. Because flow conditions can change rapidly, trading is accomplished online, in real time, through a central website. The alternative for participating sources would have been the construction and maintenance of larger saline water reservoirs, at much greater expense.

In 2010, the U.S. EPA imposed a pollution "budget" on the six states whose rivers and streams discharge to the Chesapeake Bay under the Clean Water Act's Total Maximum Daily Load (TMDL) program. The

water pollution problem addressed by this program is nutrient concentrations, which cause excessive aquatic vegetation and eventual decomposition, which deprives deeper waters of oxygen, creating hypoxic or “dead” zones. Three states—Pennsylvania, Virginia, and Maryland—opted to implement water quality trading programs to reduce compliance costs for the abatement required by the Chesapeake Bay TMDL.

As noted earlier, Pennsylvania’s program is a traditional market. Thus far, trades have been made by municipal sewage treatment plants, counties, industrial point sources, and several brokers or aggregators of credits for nonpoint source abatement by farms. Trades are facilitated through online auctions, although some bilateral negotiations also take place.

This trading program, and the others set up to address the problem of hypoxia in the Chesapeake Bay, are too young to fully assess in terms of their effectiveness and cost-effectiveness relative to other regulatory approaches. However, a prospective study has estimated the potential cost savings from water quality trading across all of the states covered by the Chesapeake Bay TMDL. If sources are allowed to trade only with other point sources, within a river basin, and within a state, compliance costs could be reduced by \$78 million per year, or about 20 percent relative to no trading. That potential cost savings would increase to about 50 percent if trading were allowed watershed-wide across state and basin boundaries, among all sources.³³

As discussed in Chapter 9, regulators must deal with nonuniform mixing of pollution in water quality trading programs, establishing complex matrices that describe the effects of effluents by each potential polluter on pollution concentrations at various points downstream. For example, say that we are trying to reduce the dead zone in the Gulf of Mexico, and we have determined that the two main causes are farming in midwestern states upstream of the Mississippi Delta and municipal sewage treatment plant discharge from cities in the delta. The effects on water quality in the gulf of these upstream and downstream effluents, of nitrogen and phosphorus in different forms and quantities, will differ by source and even by season. We cannot simply set a cap on the two pollutants in the delta region and allow polluters to trade, without first thinking about the ratios at which each pair of sources should be allowed to trade. For this reason and others, markets for water quality problems are more complex than for carbon dioxide emissions or for fishery management.

However, one of the biggest barriers to the expansion of these programs in the United States and elsewhere is not related to the issues we have raised thus far: transaction costs and nonuniform mixing of water

pollutants. Instead, it is a much more basic problem that has nothing to do with markets but everything to do with the success or failure of market-based policies: who is included and who is excluded from the cap on water pollution. In particular, agricultural nonpoint source pollution is essentially unregulated by the Clean Water Act, creating a de facto property rights distortion that hampers the ability to attain water quality goals (because this source of pollution is the primary remaining cause of impairment in U.S. rivers and streams) and increases compliance costs for the sources that participate. In some cases (including the Chesapeake Bay trading programs described earlier), farmers are enticed to participate by the promise of financial gain. Because they tend to have low-cost abatement opportunities on their land (because of the absence of direct regulation), they face significant demand for credits for their abatement activities from regulated sources that have climbed up their own marginal abatement cost curves due to decades of increasingly stringent regulation. However, establishing reliable data on the effectiveness of on-farm abatement activities and dealing with the issue of liability for emission reductions have both been significant barriers to increased trading involving farms. This problem is similar to the problem of incorporating carbon offsets from countries outside any greenhouse gas emission cap into CO₂ trading programs in countries and other entities that do face a cap, and it is a significant challenge to further expansion of trading in this context.

Waste Management: “Pay as You Throw”

Market-based approaches have also been used to manage solid waste. Some waste products have high recycling value and may not even enter the actual waste stream. We rarely observe piles of copper piping set at the curb by households on trash collection days. Similarly, Alcoa (the aluminum giant) sponsors local drop-off centers for aluminum recycling—not as a charitable effort but because it is much less costly to produce aluminum from scrap metal than from virgin ore. However, the bulk of municipal solid waste eventually ends up as trash, with legal and illegal disposal costs (such as pollution externalities, collection costs, and landfill space) borne by the community as a whole.

One estimate suggests that the marginal cost of garbage collection and disposal by the public sector for an American household is \$1.03 per bag. But until recently, the private marginal disposal cost for these households was zero.³⁴ Prompted by rising landfill costs and incineration fees (in part due to widespread community opposition to the siting of new waste disposal facilities “in their backyards”), many U.S. communities in the

1980s and 1990s began experimenting with volume-based waste disposal charges, also known as pay-as-you-throw systems. By internalizing some marginal waste disposal costs, these market-based approaches create incentives to minimize waste volume through recycling, composting, and reducing demand for products with excessive packaging. The programs can take many forms, including requirements for the purchase of official garbage bags, or of stickers to attach to bags, not to exceed a specified volume; periodic disposal charges for the number and size of official city trash cans collected at the curb; and charges based on the measured weight of trash. In 2006, more than seven thousand U.S. communities had implemented some form of pay-as-you-throw disposal.³⁵

A study of a pay-as-you-throw program in Charlottesville, Virginia found that charging 80 cents per garbage bag resulted in a 37 percent decrease in the number of bags.³⁶ However, this effect was offset by two factors that are common to such programs. First, the reduction in the weight of trash (a better indicator of disposal cost than volume, because garbage trucks compact trash bags anyway) was much smaller: only 14 percent. This phenomenon is known as the Seattle Stomp, after the supposed exertions of Seattle residents who responded to volume-based charges by compacting their trash themselves. Second, charging for trash creates an incentive for illegal dumping. In the Charlottesville case, the true reduction in weight was only 10 percent once illegal disposal was taken into account. A broader study of twenty U.S. metropolitan statistical areas found that the effects of unit disposal pricing were unclear and almost certainly much smaller than the effects of simply providing curbside recycling pickup. A more recent study of two hundred New Hampshire towns, sixty-two of which were using these market-based solid waste management programs, suggests that their introduction reduces municipal solid waste generation very significantly and that increasing the per-bag disposal cost has an additional marginal effect on waste reduction.³⁷ Overall, the empirical evidence suggests that pay-as-you-throw policies do reduce the volume of municipal solid waste and that reducing barriers to recycling and increasing enforcement of illegal disposal regulations may be important complements to these policies.

Habitat and Land Management

Some of the most novel applications of market-based policies have concerned land use management and habitat preservation. We discuss three prominent examples along this new frontier: tradable development rights,

wetland mitigation banking, and flexible mechanisms for endangered species preservation. At the end of this section, we discuss some common concerns that arise in this realm of market-based policy.

Tradable Development Rights

In developing countries, conflicts between conservation and development goals can be particularly intense. For households mired in poverty, the short-term cost of land use restrictions can be very high. Conversion of forests and other ecologically valuable lands to agriculture is pervasive. Attempting to balance these concerns, some countries have implemented a variety of market-based policies attempting to reweight relative returns to land uses, making the socially desirable land uses more competitive. For example, since 1965 land in southern Brazil has been subject to a requirement that each parcel of private property remain in native or regenerated forest, with mixed results. Recent changes exempt some landowners from this requirement, allowing them instead to offset the loss of forest to development on one parcel by preserving another parcel elsewhere. This policy, known as tradable development rights (TDRs), was first introduced in the Amazon region in 1998, with a requirement that the offset occur within the same ecosystem and with lands of greater or equal ecological value. Similar systems were implemented in the Brazilian states of Paraná and Minas Gerais around the same time. Simulations for Minas Gerais indicate that allowing such trading lowers the cost to landowners of protecting a given amount of forest. Naturally, the potential gains from trade increase with the geographic scope of trading, but this also increases the heterogeneity of incorporated forest areas (perhaps diminishing real substitutability) and the costs of monitoring and enforcement.³⁸

In developed countries, TDRs have been used to control urban growth (and sprawling development), as an alternative to traditional zoning regulations. Since the 1970s, about 140 U.S. communities have implemented TDRs, with many other potential programs in the pipeline. Calvert County, Maryland (near Washington, D.C.) adopted a TDR program in 1978 to preserve farmland on the urban fringe. An estimated 13,000 acres of farmland had been preserved by TDR sales through 2005, with some signs of influence on housing density in the region.³⁹

Wetland Mitigation Banking

Wetlands provide a rich set of ecosystem services, including water purification, groundwater recharge, flood control, and habitat for many species

of fish, birds, and mammals. Accordingly, they present a classic public goods problem. There are no markets for these services, and wetlands have been rapidly depleted in many parts of the world by conversion to competing land uses, such as urbanization and agriculture. Some conversion may reflect efficient land use change, as wetlands give way to highly valued uses; by the same token, in other instances the costs of losing wetlands surely outweigh the benefits from development. The key point is that in the absence of markets for wetland services, the social value of wetlands in these transactions is essentially ignored.

Since the early 1990s, the two federal agencies in the United States that share responsibility for wetlands (the Army Corps of Engineers and the EPA) have experimented with a market-based approach to this problem, known as *mitigation banking*. To secure a federal permit to convert wetlands to other land uses, a developer must compensate for the lost wetlands by preserving, expanding, or creating wetlands elsewhere. Between 1993 and 2000, developers filled 9,500 hectares of wetlands, with federal permission, and restored or created 16,500 hectares in mitigation. To meet the demand for mitigation, wetland banks have sprung up. The banks work in many different ways, but the general idea is that developers may fill or drain wetlands in one area in exchange for the purchase of credits for wetlands restoration or creation through a central broker. As of 2005, there were 405 approved wetland banks operating in the United States, nearly twice the number 4 years earlier. Of those 405 banks, 75 were sold out, that is, they had exhausted their credits. Another 169 banks were awaiting approval. Moreover, 70 percent of the operating banks were private commercial wetland banks, set up by private entrepreneurs in order to sell mitigation credits on the open market (rather than banks created to offset a specific development). Demand for these credits appears high. In North Carolina, for example, mitigation credits fetch \$30,000 to nearly \$60,000 per acre; similar prices have been reported in other states.⁴⁰

It is worth pointing out the important role played by current regulation in establishing the baseline for trading. The impetus for mitigation banking arose from a policy goal of no net loss of wetlands. In effect, mitigation banking in the United States is a cap-and-trade policy in which the cap on new wetlands loss (at least in principle) is zero.

Flexible Mechanisms for Endangered Species Preservation

A similar policy, known as *conservation banking*, has been developed to mitigate the destruction of endangered species habitat. Fittingly enough, the origins of conservation banking can be traced to a trade very much

in the spirit of Ronald Coase. In 1993, the Bank of America foreclosed on a parcel of land in southern California that had low value to developers but abundant coastal sage scrub habitat, home to the coastal California gnatcatcher, a songbird that had recently been designated as threatened under the federal Endangered Species Act (ESA). Around the same time, the state's highway department, CalTrans, wanted to build a highway project through similar habitat elsewhere in the state. The U.S. Fish and Wildlife Service approved a trade: CalTrans purchased the land from Bank of America and placed a conservation easement on it, in return for permission to proceed with the highway construction. In 1995, such trades were enshrined in a state policy modeled after federal wetland mitigation banking.⁴¹

The federal government's role in conservation banking stems from the underlying authority granted by the ESA. In 2003, the program became a national one when the U.S. Fish and Wildlife Service (USFWS) issued an official regulatory guidance approving the use of conservation banking to mitigate the destruction of endangered species habitat. As in wetland mitigation banking, developers can purchase credits from an approved conservation bank as a means of offsetting adverse impacts on threatened or endangered species. According to the USFWS, approximately forty-five conservation banks had been approved by the end of 2004, mostly in California.

An even more widely used flexible mechanism under the ESA is the Habitat Conservation Plan (HCP). Since 1983, the ESA has allowed the USFWS to issue an "incidental take permit," allowing landowners to engage in land development activities otherwise prohibited under the act, as long as they have an approved HCP. The use of these plans ramped up slowly at first, with only fourteen completed between 1983 and 1992. The Clinton administration added the "no surprises rule," ensuring that landowners' obligations under an HCP would not change even if future conditions changed (for example, if species recovery expanded the habitat present on their private property). Between 1994 and 1997, 225 HCPs were approved, and by late 2012, almost 700 were in place, affecting 40 million U.S. acres and hundreds of species. The first empirical analysis of this approach demonstrated that between 1990 and 2004, species covered by HCPs were more likely to show improvement in their recovery status, by USFWS standards, and less likely to be declining or classified as extinct than those without HCPs.⁴²

Finding ways to cost-effectively engage private landowners in species and habitat conservation is essential to the success of these endeavors,

because 80–90 percent of known U.S. endangered species populations are on private land. Over the past two decades, studies have confirmed anecdotal stories suggesting that landowners may engage in pre-emptive habitat destruction to avoid the discovery or designation of critical species habitat on their land, given the cost of the restrictions that such designation imposes.⁴³ Thus, flexible mechanisms such as HCPs and conservation banking may not only improve the cost-effectiveness of species preservation efforts; they may be essential to its *effectiveness*.

Prospects and Remaining Issues

As these examples illustrate, the application of market principles to land management is an active frontier in environmental policy. Despite the excitement and activity, however, this area is fraught with potential pitfalls. One vexing issue is how credits for wetlands or endangered species habitats can be generated. For example, California's conservation banking policy grants credits not just for creating new habitat but also for preserving existing habitat. Some critics, including the nonprofit EDF, have argued that this approach threatens to undermine the system. If new development in one location can be offset by a promise to preserve habitat elsewhere, how do we know that anything has been gained by allowing the trade? If preservation would have occurred anyway, perhaps because it takes place on land with low value for development, then the result is a net loss of habitat relative to what would have happened without trading.

Another important drawback involves ascertaining whether land in a bank is equivalent to the land being developed. In the case of carbon dioxide emissions, or a natural resource such as a fishery, equivalence is easy to determine. A pound of carbon dioxide from one automobile is equivalent to a pound from another; a boatload of fish caught by one fisher is equivalent to the same amount landed by another.

The situation is very different in the cases of natural forests, wetlands, or endangered species habitat. Land is a nonuniformly mixed resource; in plain English, not all wetlands are created equal, and this fact makes a market-based approach less appropriate. Consider a wetland in a particular location, which provides a specific portfolio of biophysical services. Many of those ecosystem services are location-specific: Coastal wetlands provide nurseries for shrimp and other shellfish, for example, and wetlands farther inland help remove contaminants from freshwater flows, providing cleaner drinking water. Even if it is possible to regenerate similar services elsewhere, a different set of people will benefit. For example, some of

the greatest pressures on wetlands occur in urban areas, where values for other land uses are highest. Therefore, we might expect that mitigation banking would shift the balance of wetlands from urban to rural areas. If the new wetlands truly provide the same services as the old ones, that shift might be the desirable effect of reallocation by a market. But if they don't provide similar services, trading may diminish rather than increase social welfare. One analysis of wetland mitigation programs suggests that generating credits from wetland banks in compensation for converting natural wetlands for development may have resulted in as much as an 80 percent loss in both acreage and function, relative to a baseline of no conversion. Other analyses are less pessimistic but emphasize the importance of linking credits to carefully defined functional ecological criteria, rather than task-oriented engineering outputs such as canal filling and grading.⁴⁴

One way to get around the problem of unequal function between converted and created or restored wetlands, and nonuniform mixing in other market-based approaches to land conservation, might be to establish trading ratios of the sort we discussed in Chapter 9. However, determining the proper ratios is much more difficult for land use than for air and water pollution, because of the sheer number and variety of ecosystem services and spatial concerns that must be taken into account. Moreover, measuring the ecosystem services from a particular site is a daunting problem, let alone comparing services from different sites. Another way to alleviate the problems associated with land use trading is to impose limits on how much trading can be done. For example, developers in the United States are supposed first to avoid impacts on wetlands, and then to minimize unavoidable impacts, before they seek permission to compensate for those impacts by purchasing mitigation credits. Such restrictions on trading inevitably limit the gains from trade that can take place, a cost that should be weighed against the benefits of preventing inappropriate trades.

From an economic perspective, developers should incorporate the social costs of the damages to environmental amenities from forest, wetlands, or endangered species habitat lost to development. At the same time, there are social benefits from providing flexibility in how and where these ecosystem services are provided, in order to make room for valued development. These competing concerns suggest a potential role for market-based policies for land use, but balancing them takes care. Governments play an important role in ensuring the quality and equivalence of trades before they take place and in monitoring and enforcing the maintenance of ecosystem services afterward.

Conclusion

These are a handful of the many creative ways in which policymakers have sought to use market principles to correct market failures, some more successfully than others. As the costs of environmental regulation and natural resource preservation become more apparent, use of these market-based approaches is likely to increase. A solid understanding of how these approaches have been implemented in the real world is therefore essential to understanding environmental policy.

We looked at the distinction between ends and means in Chapter 9. Now we can see that market-based policy instruments have found their broadest application as ways of implementing environmental policy goals that have been determined without explicit regard to efficiency. To take just one example, the 10-million-ton reduction goal for sulfur dioxide emissions enshrined in the 1990 Clean Air Act Amendments was reached through political wrangling, with only a cursory reference to the marginal costs and benefits. (And as we saw, subsequent analyses of that program have found that the benefits far outweigh the costs, implying that despite its ambitious scope, the program was not stringent enough from an efficiency perspective.)

From an economic perspective, any environmental policy that ignores efficiency is a distinctly less satisfying approach, because it gives up on the primary goal of maximizing social welfare. However, as we saw in Chapters 2 and 3, the very idea of efficiency and of benefit–cost analysis stirs up political controversy. Therefore, it is hardly surprising that economic analysis has exerted its strongest influence on how policies are designed rather than on their goals.

11

Sustainability and Economic Growth

Microeconomics, the subject of Chapters 2 through 10, examines how households and firms make choices and interact at the scale of individual markets. When we ask whether the owner of a natural resource will incorporate scarcity into her extraction decisions, or whether the manager of a steel mill will take account of the damages from pollution, or how a government policy will shape the incentives of firms and individuals, we are exploring questions of microeconomics. Macroeconomics takes a more top-down view, focusing on economy-wide phenomena—the sum of millions of micro-level actions by households and firms. In this chapter, we address the intersection of economic growth—a macroeconomic phenomenon—and the natural environment.

You may wonder why we focus on economic growth. What is the link between growth and social welfare, the measure we have discussed throughout the text as an appropriate gauge for environmental policy choices? Economic growth can reduce the pain of the tradeoffs necessary in every social decision-making context. Enlarging the pie allows greater potential for environmental quality and other things you may care deeply about; thus, it has the potential to increase social welfare. However, just as efficiency is neither a necessary nor a sufficient criterion for sound public policy, growth is just one piece of the puzzle of social welfare. Economist Joseph Stiglitz has compared looking at economic growth to measure well-being with looking only at a firm's revenues, or a family's income, to determine well-being at an individual level. We are really interested in the balance sheet, which reports not just flows of revenues and expenditures but stocks of assets and liabilities.¹

In the same way that governments have the power to correct market failures and help promote efficient outcomes, governments have the capacity to spur economic growth and help direct its course. Many of the best places to live in the world are not simply places experiencing significant economic growth but places where income is equitably distributed, education and health care systems are good and widely accessible, and environmental quality is high. Just like households, countries may grow but dissipate much of their income for consumption rather than investment in these kinds of assets and institutions. Nonetheless, economic growth is a central topic in macroeconomics and the focus of much discussion among policymakers and economists. Thus it is worth focusing on growth and its implications for the environment, keeping in mind the caveats discussed earlier.

We begin our discussion of economic growth and the environment with the debate over the degree to which the finite availability of some natural resources can be expected to act as a limit to growth. We then define the concept of *sustainability* in economic terms, contrasting economic definitions with other common definitions. This is followed by a discussion of *green accounting*, the practice of including additions to and decreases in natural capital in the calculation of traditional measures of economic growth. We end the chapter with some reflections on the viability of economic growth as a global goal, applying competing concepts of sustainability.

Limits to Growth?

In Chapter 6, we solved a two-period petroleum extraction problem in which we assessed the impact of the limited petroleum stock on efficient extraction. Our analysis suggested that private owners of nonrenewable natural resources would consider the limited stock in determining how much to extract each year; thus, the resource would be extracted efficiently. In addition, we developed the concept of marginal user cost, a measure of scarcity that incorporates economic factors and the physical limits of resource stocks. How are the decisions of millions of individual firms and consumers regarding nonrenewable resource extraction reflected in the global economy? The fact that many minerals that historically have been critical inputs to economic growth are available in limited quantity in the earth's crust has given rise to a contentious debate over the possibility that we will run out of some or all of those minerals, with drastic consequences for global economies.

The Limits of Assumptions

This dire scenario was the subject of the book *The Limits to Growth*, based on a system dynamics model developed at the Massachusetts Institute of Technology in the early 1970s.² This large-scale computer model was designed to simulate likely future outcomes for the world economy. The model assumed continued exponential economic growth; fixed stocks of nonrenewable resources; no substitution between nonrenewables and abundant inputs; no changes to the world's basic physical, economic, social, and political institutions; and no technological change. The conclusions of *The Limits to Growth* and subsequent models were grim. They foresaw two possible outcomes. Either the world's nations would immediately exercise the self-restraint necessary to bring economic growth almost to a halt, thereby avoiding collision with the earth's natural limits, or the global economy would collapse within 100 years. The collapse of the economic system would be due to scarcity, given no new resource stocks; to pollution, should known resource stocks double and be consumed; or to population growth, should stocks double and pollution be controlled.

The early 1970s was a time of significant worry in the United States and other countries, given the oil shocks of that decade, high rates of inflation, rapid population growth, and increased attention to the problem of pollution in industrialized and developing countries. But the *Limits* view actually has its roots in some very old economic models. In the 1800s, Thomas Malthus, David Ricardo, and others were very concerned with the limited ability of the earth's resources (particularly land) to support fast-growing populations. And recall our discussion in Chapter 6 of Stanley Jevons, the nineteenth-century economist who worried about coal depletion and also hoarded writing paper, anticipating a future shortage of trees.

Cause for Optimism

Modern economists' answer to the *Limits* models and their contemporary counterparts (e.g., "peak oil") has been called economic optimism. Like the *Limits* view, the optimistic view is

The Limits to Growth model of the early 1970s predicted that without immediate constraints on growth, global economies would collapse in as little as 100 years because of natural resource scarcity, pollution, and population growth. The model made very restrictive assumptions about resource scarcity, substitution, and technological and social change.

represented by a particular approach, that of Julian Simon, a population economist who died in 1998.³ But Simon was not the only dissenter: The *Limits* arguments are disputed by most of the discipline of economics. In the words of one analyst, the only general conclusion one would reach from examining long-run economic growth models is that there is no general conclusion.⁴ The conclusions of such models depend critically on what they assume about some key parameters. Long-run economic growth depends on the growth of inputs, the rate and direction of technological change, and the degree to which different inputs to production can substitute for one another. Thus, the question of whether the finite availability of some inputs will slow or stop economic growth is actually an empirical question about the relative influence of these factors in the real world. The *Limits* models make restrictive assumptions about these factors that are not supported either by analysis of historical data or by general consensus regarding likely future trends.

The bottom line in the optimistic view is that, although resource scarcity may exert a small drag on global economic growth, the benefits of technological change have thus far outpaced the influence of these resource constraints and will probably continue to do so. In addition, real economic and political systems respond to scarcity, unlike the systems simulated by the *Limits* model. For example, scarcity causes prices to rise (as in the two-period oil extraction model we discussed in Chapter 6), decreasing demand, making new natural resource stocks worth exploiting, and acting as an incentive for the development of new technologies. Consider, for example, the boom (beginning in the mid-2000s) in the extraction of oil and gas from deep shale formations, discussed in Chapter 6. Increasing incomes over time tend to result in lower pollution levels, not higher levels as assumed by the *Limits* models. Income increases

Economic models generally suggest that, although resource scarcity may exert a small drag on global economic growth, the positive effect of technological change has thus far outpaced the influence of resource constraints and will probably continue to do so.

They are optimistic in comparison to the Limits to Growth model.

also tend to slow population growth, as evidenced by the very small rates of population growth—in some cases, shrinkage—observed in Western Europe and the United States and declining rates of growth in some rapidly developing economies.

How Scarce Are Natural Resources, in Economic Terms?

In Chapter 6 we discussed the fact that physical measures of natural resource reserves are insufficient economic

Growth, Income, and the Environment

Over the past several decades, economists have observed an inverted U-shaped relationship between national income and environmental quality, focusing on pollution in particular, although some have looked at other environmental disamenities such as deforestation. Pollution may increase during early periods of development, reach a maximum, and then fall as incomes rise beyond that turning point. From this evidence, some have hypothesized a causal income–pollution relationship, known as the environmental Kuznets curve (EKC).⁵ Although evidence of lower local and regional pollution levels in higher-income countries is robust, the causal EKC hypothesis is more controversial.⁶

Early studies appeared to confirm an EKC pattern for air pollutants such as particulate matter and SO_2 .⁷ But these early results have not held up to better data and further scrutiny. Both the inverted U-shape and per capita income turning points established by these early studies appear to be highly sensitive to slight data variations and the specification of statistical models.⁸ Economic growth is accompanied by many other changes, including changes in the scale, composition, and technologies of production, as well as political and civil institutions, which are all associated with changing pollution levels.⁹ These factors and others confound the ability of researchers to estimate the causal relationship between income and environmental quality at the heart of the EKC hypothesis.

Whether the pattern is real or not is a matter of particular importance, as some have interpreted the EKC as a reason to suggest that economic growth will eventually repair much of the damage from early exploitation of resources such as clean air and water. As noted earlier, there is little empirical evidence to support the idea that economic growth alone is sufficient to induce efficient pollution control. In addition, the EKC says nothing about natural resource extraction, about global CO_2 emissions with few direct local effects, or about the role of trade in supporting consumption of polluting goods and services. For example, one analysis of U.S. states suggests that CO_2 emissions related to consumption peak at much higher incomes than those related to production; as incomes grow, higher consumption levels may simply be supported by the export of emission-intensive production.¹⁰

measures of scarcity. One example of this is the petroleum reserves-to-production ratio, known petroleum reserves divided by annual production, which presumably gives the number of years to petroleum exhaustion. What has happened to this number over time? Table 11.1 lists the world reserves-to-production ratio periodically from 1980 to 2013. The number does not decrease monotonically, as one would expect in a scenario in which the world was using up a finite quantity of known reserves. Instead,

it increases and decreases, indicating movement along both the vertical and the horizontal dimensions of the McKelvey diagram we encountered in Chapter 6. In 1980, the ratio of reserves to production was 28 years, suggesting that world petroleum resources would be exhausted in 2008; 5 years after that “deadline,” in 2013, the reserves-to-production ratio was 50 years. Tracked over time, this statistic is characterized by periods of stasis or decline, punctuated by increases due to discovery of new reserves or exploitation of known reserves using new technologies.

The best way to measure scarcity from an economic perspective would be to examine trends in resource rents, which should increase as stocks dwindle. In practice, however, data on resource rents are not readily available. If resource rents are simply the difference between price and marginal extraction cost, why are they difficult to calculate? Many of the resources we have discussed in this book—clean air, clean water, some forests, and some fish—are either freely available at a price of zero or are priced at levels that ignore scarcity rent. In these cases, prices reveal no information about scarcity. Even for natural resources for which we do believe that prices capture scarcity—privately owned minerals traded in markets—it may be difficult to separate marginal extraction cost from other costs. Finally, in most countries marginal extraction cost is proprietary information, data that private firms need not reveal to economic analysts. As a result, even in markets for privately owned nonrenewable resources, analyses attempting to assess economic scarcity must use indirect observation.

Attempts to assess the scarcity of nonrenewable resources through resource prices, or through scarcity rents reconstructed by statistical means, reveal that resource prices are either trendless or decreasing over time. Much attention has been paid to the bet between biologist Paul Ehrlich and economist Julian Simon. In 1980, Simon challenged Ehrlich to choose a list of any five metals, worth a combined \$1,000. If the 1990 inflation-adjusted price of the package was higher than \$1,000, Ehrlich would win; if the value of the package in 1990 was lower than \$1,000, Simon would win. The stakes were tied to the change in the package price. Ehrlich chose copper, chrome, nickel, tin, and tungsten. In 1990, Ehrlich sent Simon a check for \$576.07; the prices of all five metals had fallen dramatically over the decade.

Such tremendous price variation is reasonably common in the short and medium term. But in the long term, inflation-adjusted average prices do not appear to trend upward (as we would expect, were scarcity really

Table 11.1
Reserves-to-Production Ratios for
Petroleum, 1980–2013

<i>Year</i>	<i>Ratio (years)</i>
1980	28
1985	32
1990	41
1995	39
2000	36
2005	41
2010	42
2013	50

Source: Calculated by the authors from data published by the U.S. Energy Information Administration, “International Energy Statistics,” available at <http://www.eia.gov/cfapps/ipdbproject/IEDIn dex3.cfm>, accessed February 11, 2015.

a binding constraint).¹¹ In fact, real prices for many important nonrenewable natural resources have actually declined over time.¹²

Why might scarcity not appear to be increasing, in economic terms? After all, we are certainly consuming resources at a fast pace: In 2013, the world consumed about 91 million barrels of oil per day. Substitution possibilities and technological change provide the likely explanation.¹³ In the words of one analyst, “Technological progress is holding scarcity at bay.”¹⁴ Think of the many examples of this phenomenon. New seeds and chemical fertilizers have outpaced the need to cultivate marginal lands due to growing populations in many countries. Advances in finding and extracting oil have countered the need to drill deeper, access less permeable formations, and develop wells in harsher climates. Fiber optics substitute for copper as a means of information transmission. Recycled aluminum substitutes for aluminum newly processed from bauxite. Ceramics replace tungsten in cutting tools. Better irrigation technologies substitute for some nonrenewable groundwater in agricultural production. Abundant sources of energy such as solar and wind power, although relatively expensive today, act as ceilings on the future prices of oil, coal, and natural gas. (After all, to the extent that renewable energy sources are a substitute for fossil fuels, no one will be willing to pay more for power generated from coal than that from the sun or wind.)

How Large Is the Impact of Scarcity on Economic Growth?

Even if prices are falling over time, scarcity may still dampen economic growth. For many nonrenewable resources, it may be that resource rents are rising as stocks fall, but the contrary effects of substitution and technological change may simply mask this indicator of scarcity because of their stronger downward pressure on prices. If some inputs to important economic processes are physically limited, and those resources do not have abundant perfect substitutes, then we should be able to detect the negative impact of scarcity on economic growth through careful empirical analysis. Two analyses of note have estimated this negative impact. Nordhaus (1992) examines the influence of the limited availability of a set of nonrenewable resources, plus a set of renewable resources (including clean water and clean air) on economic growth. He estimates that the scarcity of this set of resources can be expected to slow global economic growth by a combined 0.31 percent per year between 1980 and 2050 (see table 11.2). Weitzman (1999) estimates that the limited availability of fourteen minerals important to economic growth causes a decrease in global consumption of about 1 percent per year.

How does the measured drag on growth imposed by scarcity compare to the positive force exerted by substitution and technological change? Weitzman estimates the positive welfare impacts of technological change to be about forty times the negative welfare impacts of nonrenewable natural resource depletion. Thus far, technological change has overwhelmed scarcity in the process of economic growth.

The historical evidence economists have offered cannot be considered the final word on whether economic growth at current levels is sustainable into the indefinite future. The economic argument, like arguments in other disciplines, is based on historical data and a consensus of analysts' best guesses about future scenarios. But the economic evidence overwhelmingly supports the optimistic view that nonrenewable resource scarcity will not be an important obstacle to continued economic growth in the long run.

Sustainability, in Economic Terms

We have come to the conclusion that economists are not particularly worried about running out of any specific nonrenewable natural resource. Is this reconcilable with any potential definition of sustainability?

The word *sustainability* is most commonly associated with the definition offered by a report of the United Nations World Commission on

Table 11.2

Estimated Drag on Economic Growth from Limited Resources, 1980–2050

<i>Source of drag</i>	<i>Impact on world growth rate, 1980–2050 (percent)</i>
<i>Market goods</i>	
Nonrenewable resources	
Energy fuels	–0.16
Nonfuel minerals	–0.03
Renewable resources	
Land	–0.05
<i>Environmental goods</i>	
Global warming	–0.03
Local pollutants	–0.04
<i>Total</i>	–0.31

Source: Adapted from William D. Nordhaus, “Lethal Model 2: The Limits to Growth Revisited,” *Brookings Papers on Economic Activity* (2): 1–59 (1992), table 3, p. 31.

Environment and Development (the Brundtland Commission, named for chair Gro Harlem Brundtland) in 1987. The commission defined the term as follows: “Sustainable development is development that meets the needs of the present without compromising the ability of future generations to meet their own needs.”¹⁵ As a goal to strive for, this definition of sustainable development is vague. Do “needs” refer to specific resources per se or just the ability to sustain specific levels of human welfare? Are the needs of future generations less than ours, the same, or greater? Are needs to be considered in absolute terms or per capita?

Related concepts were topics of economic research well before the Brundtland Commission. But the Brundtland report prompted increased debate among economists, like analysts in many fields, to think about useful definitions of sustainability for decision making, especially in a policy context. At the Woods Hole Oceanographic Institute in 1991, Robert Solow, the 1987 Nobel laureate in economics, offered a lecture titled “Sustainability: An Economist’s Perspective,” which forms the basis of the economic definition of sustainability we will discuss here.¹⁶

The economic definition of sustainability requires that we leave future generations the capacity to be as well off as we are. It does not require the preservation of any particular resource; rather, it hinges on the investment of resource rents, substitution possibilities, and technological change.

Sustainability, according to Solow, means leaving to future generations “the capacity to be as well off as we are today.” The concept is essentially about distributional equity between generations, and it requires that we avoid “enriching ourselves by impoverishing our successors.” In doing so, we must think carefully about what we use up and what we leave behind in an attempt to preserve an intergenerational balance sheet.

How can this be reconciled with the prospect of depleting global oil reserves, or any natural resource stock? Economic sustainability does not require the preservation of specific resources, species, things, or places. This is not to diminish the value of such preservation efforts. Even solely on the grounds of efficiency, many individual resources, species, things, or places may be worth preserving for their own sake, because their value to people is greater than the opportunity cost of preserving them. But the larger concept of sustainability, by Solow’s definition, is a generalized goal not specific to any resource. Natural capital, human capital, and physical capital are all, to some degree, interchangeable. Thus, if we deplete fossil fuels to drive economic development in the twenty-first century, we must create, in the process, enough capital of other types to replace that lost value, leaving future generations the capacity to be as well off as we are today through our consumption of fossil fuels.

Sustainability, Substitution, and Technological Change

The economic definition of sustainability, like the economic arguments about depletion of nonrenewable natural resources, hinges on substitution possibilities and technological change. Clearly, some natural resources and environmental amenities have no substitutes. If humans value these resources highly (including both use and nonuse value, as defined in Chapter 3), then depleting them would not be consistent with Solow’s economic definition of sustainability.

In most cases, however, substitution possibilities are a matter of degree. Seawater is a substitute for freshwater in industrial cooling processes if firms invest in the materials necessary to prevent rapid corrosion of capital equipment from saltwater exposure; some coastal electric generating plants use seawater for cooling, especially in arid regions. Given a large enough investment in desalination equipment, seawater can also be a substitute for fresh drinking water. Middle Eastern countries including Saudi Arabia, Israel, and Kuwait derive much of their citizens’ drinking water supply from the sea, and drinking water desalination plants have been constructed in freshwater-scarce parts of the United States, including Tampa, Florida, and Santa Barbara, California.

We are not suggesting that seawater is a perfect substitute for freshwater; it takes seawater, plus some other inputs such as energy, to produce the quantity of freshwater desired for drinking water and other uses. Without freshwater, life on Earth would cease to exist. Obviously, a decision by world nations to consume the entire global supply of freshwater to fuel economic growth would not be economically sustainable. However, a decision by agricultural communities in the Great Plains of the United States to draw down nonrenewable groundwater supplies, fueling agricultural production and thereby economic development in the region, might be economically sustainable under some conditions. A decision by communities in northwest India to deplete groundwater to provide safe drinking water to impoverished households might as well. In the same sense, drawing down reserves of nonrenewable oil, coal, and natural gas might be economically sustainable under some conditions.

Substitution possibilities are easier to identify and accept when we discuss resources such as fossil fuels, water, and metals than they are when we discuss individual species, landscapes, and other natural resources and environmental amenities for which we have a particular affinity as human beings. Perhaps it is high nonuse value that separates these “low-substitution” goods and services from those for which we can easily imagine substitution possibilities. Even in some of these cases, substitution can sometimes be a matter of degree. The collapse of the cod fisheries in the North Atlantic, if not reversible, is surely an ecological disaster, and regionally it is the cause of much economic dislocation. But although there may be no substitute for cod as a component of its resident ecosystems, populations in the United States, Canada, and Europe have substituted cod with other fish, purely for the purpose of human consumption.

Investment of Natural Resource Rents

Substitution and technological change are only part of the story when it comes to economic sustainability. In addition to substitution possibilities and technological change, economic sustainability also speaks to the role of investment of rents from natural resource depletion.

An example may illuminate this concept a bit. In the 1960s, oil deposits were discovered in the North Sea, off the coast of the Netherlands. The United Kingdom and Norway are the two major oil-producing countries in the North Sea. These two countries have taken very different approaches to the depletion of this valuable, nonrenewable natural resource. Rising oil prices in the 1980s made large-scale exploitation of the North Sea fields economically feasible. By the mid-1980s, the United Kingdom

was producing millions of barrels of oil per day. The government of prime minister Margaret Thatcher supported a policy of rapid extraction, making the United Kingdom a net oil exporter and generating substantial tax revenues, used largely to support current consumption and to lift the country out of a long economic recession.

In contrast, exploitation by Norway has occurred at a slower pace. In 1990, Norway established a Petroleum Fund, which receives tax revenues from oil companies extracting from North Sea fields and royalties for licenses to explore. The fund is owned by the citizens of Norway and administered by the Norwegian Central Bank. In February 2015, the value of the fund's portfolio was more than \$6.5 trillion.

We are not suggesting that Norway's reinvestment in nonoil capital has been sufficient to offset the loss to future generations resulting from depletion of its North Sea fields (we do not have the data to make such an assertion). But the contrast illustrates one aspect of Solow's definition of sustainability: Consuming the rents from resource extraction is probably not economically sustainable, but investing those rents is, at the very least, a step in the right direction. It is not depletion of natural resources, then, that damages the capacity of future generations to be as well off as we are but rather depletion of the value of the total global capital stock—natural, physical, and human.

Problems with the Economic Definition of Sustainability

The economic concept of sustainability offered by Solow requires a good deal of knowledge about the future. To put it into practice, we must know something about the tastes and preferences of future generations and perhaps something about the future technologies that will be available to achieve welfare gains. It is difficult to imagine how we would determine whether we are leaving future generations the capacity to be as well off as we are without this information.

For example, to residents in industrialized countries, the value of environmental quality and natural resource amenities has increased with incomes. Should we assume that, as global incomes continue to grow, this trend will continue? In this case, future citizens will place a higher value on these goods and services than we do today, and public policy decisions that affect future environmental quality and natural resource amenities must take this into account.

Predictions about the rate and direction of technological change, even by experts, are notoriously shaky. Lord Kelvin, in 1895, is said to have claimed, "Heavier-than-air flying machines are impossible." Robert

Millikan, 1923 Nobel laureate in physics, noted in his acceptance speech that “there is no likelihood man can ever tap the power of the atom.” For overestimates of the pace of technological change, one need only watch an episode of the futuristic cartoon *The Jetsons* or the movie *2001: A Space Odyssey* to get an inkling of just how large the margin of error is in forecasting future technologies.

It is not uncommon to read public documents that project the benefits and costs of policy interventions, particularly for climate change, out to the year 2100 or further. A modicum of humility suggests that we are as blind in these predictions as analysts were in 1900 in thinking about what the world would be like today. Things change rapidly, often in directions we cannot anticipate even a few years in advance. Constraining the direction and magnitude of those changes with respect to the goal of sustainability, based on today’s myopic perspective, might result in tremendous (and unpredictable) losses in future welfare.

Given the difficulty in predicting future tastes and preferences and the future pace and direction of technological change, it may seem that Solow’s definition of sustainability is no more workable as a basis for policy decisions about the environment and natural resources than the Brundtland Commission definition we offered earlier. A partial response to this criticism is that we may be equally likely to err on both sides of these uncertainties about future preferences and the path of technological change. Thus, if we are acting in expectation, we may come reasonably close to the right path in managing the global capital stock with an eye to the future. Regardless of its direct workability as a basis for policy decisions, the economic definition of sustainability has much to offer in the way of insights for current decision making.

Insights of Economic Definitions of Sustainability for Environmental Policy

Economic definitions of sustainability highlight some insights for current natural resource management and environmental protection policies. Those insights are the importance of correcting the negative externalities that arise from pollution and natural resource extraction, the consideration of sustainability as a problem of intergenerational equity, and the potential conflict between intergenerational and intragenerational equity concerns.

Getting Prices Right

The first major insight we draw from the economic perspective on sustainability is that it requires dynamic efficiency, the subject of previous

chapters in this book. Prices tell producers and consumers about the economic value of a good, service, or natural resource amenity—its value in use and the opportunity cost of its consumption, including relative scarcity. So far we have framed our discussion of externalities in microeconomic terms. But their importance is magnified in the macroeconomic context. If individual firms and consumers do not bear the full social marginal cost and benefit of production and consumption, then the aggregate consequences of their actions will diminish welfare. Without policies in place to correct environmental externalities, economic growth is unlikely to be sustainable.

Sustainability as a Problem of Intergenerational Equity

The second major insight of economic sustainability is the importance of investment rather than consumption of resource rents. If we deplete specific natural resources in the process of economic growth, we must leave equivalent capital assets of other types to future generations so that their welfare is not diminished by this depletion.

We can think of this notion of capital investment as Pareto efficiency, with an intertemporal dimension.¹⁷ In Chapter 3, we discussed the concept of Pareto efficiency. A Pareto efficient policy is one that makes at least one party better off, and none worse off, than they were before the policy was enacted. If some parties are made worse off by the policy, then the “winners” must compensate the “losers,” which is possible as long as the policy results in net gains to society. Without those income transfers, the policy might pass a benefit–cost test, but it would not be Pareto efficient. Economic sustainability is Pareto efficiency across generations. If, by applying the rules of static and dynamic efficiency we have described in this book, we determine that depleting a natural resource stock is efficient, we must then consider whether its depletion will result in a net gain or a net loss to future generations. In the case of a net loss to the future, we are obligated to provide compensation, in the form of returns to an investment of resource rents or some other form. Under these conditions, natural resource depletion may be economically sustainable.

On a hopeful note, the course of global economic development throughout history, though full of short-term regional ups and downs, has been quite positive from the perspective of economic sustainability. That is, past generations have left us the capacity to be at least as well off as they were. If we consider the continually rising standards of living in most parts of the world, one could argue that past generations were “too generous” in this regard.

Conflicts between Intergenerational and Intragenerational Equity

Thinking about economic sustainability as intergenerational Pareto efficiency brings us to a third insight. A concern for sustainability raises the inherent conflicts between intergenerational and intragenerational equity. If sustainability is about equity between generations, can it tell us anything about equity within generations? Economist Tom Schelling has noted the paradox of considering the expenditure of large sums today to “purchase” a climate less influenced by anthropogenic carbon dioxide emissions tomorrow, when those funds could be used to purchase increased quality of life for today’s poor.¹⁸

This is a particularly stark point, given that most future beneficiaries of actions taken today to prevent global climate change will reside in developing countries and that future residents of these countries will almost certainly be wealthier than today’s residents. Such policies, then, may amount to billing poor residents of developing countries today for welfare improvements accruing to wealthier residents of these countries tomorrow. Note that this paradox does not require that developing countries actually pay out of pocket for climate change policies today—only that they pay in terms of reduced resources available for mitigating the impacts of poverty and environmental degradation today. One dollar spent on improving environmental quality and natural resource amenities for future generations is one dollar not spent on improving welfare (environmental or otherwise) today.

Keeping Track: Green Accounting

A consideration of economic sustainability forces us to think more carefully about whether we are acting as good stewards of the world’s capital stock. If sustainability requires that we consume and invest such that future generations have the capacity to achieve our own level of well-being, it also requires some method of keeping track of whether this preservation of the global capital stock is actually occurring. Adjusting standard indicators of economic growth to reflect natural resource stocks and the state of the environment is an important step in keeping track of our progress toward that goal.

Traditional indicators of economic growth are measures of goods and services produced by labor and property either within a country (gross domestic product [GDP]) or supplied by a country’s nationals (gross national product [GNP]). These gross measures can also be converted to

net measures by subtracting capital depreciation, creating net domestic product (NDP) and net national product (NNP).¹⁹

Each of these traditional measures of economic growth excludes non-market activities, such as household production (cooking, cleaning, child care, and home improvements, for example) and black markets for goods and services. Most importantly for our purposes, they also exclude environmental services, such as clean air and water, and the value of natural resource stocks, such as oil, coal, forests, and fish. Many economists have suggested that a true measure of economic growth should account for changes in the value of these assets. A measure of growth that includes the depreciation of not only the stock of physical capital but also that of natural capital is often called “green NNP.”

Should Traditional Measures Be Changed?

The debate over whether the traditional measures of economic growth were sufficiently inclusive began well before the phrase *sustainable development* was first uttered. The fathers of national income accounting in the United States, the process that undergirds the calculation of growth measures, were well aware of its shortcomings. Simon Kuznets, who received a Nobel Prize for his work in developing the U.S. national income and product accounts, did not intend these growth measures to become measures of social welfare, as they have often been used. Income accounting was designed with many important goals in mind, including the provision of indicators of an economy’s performance over time, measurement of savings and investment, and tracking of business cycles.

The information produced by this process is invaluable to governments seeking to implement policies to promote or control the pace and direction of economic growth. Nonetheless, A. C. Pigou noted, with irony, that “If a man marries his housekeeper or his cook, the national income is diminished.” More to the point of this chapter, he states,

It is a paradox, lastly, that the frequent desecration of natural beauty through the hunt for coal or gold, or through the more blatant forms of commercial advertisement, must, on our definition, leave the national dividend intact, though, if it had been practicable, as it is in some exceptional circumstances, to make a charge for viewing scenery, it would not have done so.²⁰

This paradox identified by Pigou has wide-ranging implications. If the owner of an oil well in Texas pumps out her remaining reserves and sells

the oil, the value of her sale is added to U.S. NNP, the depreciation of the physical capital used at the well (such as a drilling rig) is subtracted from U.S. NNP, but nothing is recorded to account for the value of the oil no longer beneath the ground in Texas.

The mining of 1 ton of coal from a U.S. mine increases GNP by \$17; subtracting the depreciation of mining equipment and the 1-ton decrease in coal reserves would bring that number down to about \$5.50; and the number would be even lower were we to consider the pollution externalities from mining and burning coal.²¹ The value of privately owned domestic livestock lost to disease is deducted from traditional measures of NNP; commercial fishery depletion is not.

We have spent a good deal of time in this book describing how economic theory treats natural resources as capital assets in determining efficient extraction rates. The exclusion of these assets from measures of economic growth is inconsistent with this practice. Even where marketed natural resource commodities make important contributions to national output—oil, fish, and timber are good examples, particularly in some developing countries—additions to and subtractions from the stocks of these resources are not included in the calculation of NNP.

If these shortcomings of traditional measures were recognized from the outset, why has their calculation not been adjusted to account for these excluded portions of economic activity? There are two main reasons why this is the case. First, measurement is difficult; the values of nonmarket goods and services are, by definition, difficult to capture. Even for marketed goods and services that are currently excluded (such as minerals, fish, and timber), the ease of measuring changes in stocks varies greatly. For this reason, most of the initial experiments with green accounting that we will discuss in the next section have started with the easiest categories, minerals and timber.

Second, as a measure of economic growth, traditional NNP is likely to be strongly correlated with a true, all-inclusive measure of NNP. However, if the economic value of excluded goods and services has increased over time, the inherent bias in using traditional NNP as a proxy for true NNP has probably grown as well. For example, the value to the U.S. population of recreational opportunities in wilderness areas has increased markedly. In some cases, these uses conflict directly with commercial use

A measure of growth that includes the depreciation of not only the stock of physical capital but also that of natural capital is often called "green NNP."

of such lands. Commercial timber extraction and mining generate additions to NNP, but no subtractions are made to represent the opportunity cost of potential recreation and other values lost when wilderness is developed. Over time, this omission has grown in importance, because of the growth in these values, making NNP an increasingly biased measure of U.S. economic activity.

Experiments with Green Accounting

Some of the experiments with greening the national income accounts have taken place in industrialized countries, where the idea began to take root in the 1970s. Two economists at Yale asked the question, “Is Growth Obsolete?” and attempted to incorporate things such as traffic congestion, crime, and, to some extent, natural resource depletion, into U.S. economic growth measures.²² Their comprehensive welfare measure increased by about 42 percent between 1929 and 1965, less than one half of the increase in traditional per capita growth measures over this time but still a substantial increase. Around the time that Norway began to exploit its North Sea oil deposits, Norwegian analysts focused on the issue of how to account for the depletion of this resource in calculating national economic growth estimates. Norway has since estimated green NNP, accounting for depletion of oil, fisheries, forests, and even clean air (by including some air pollutant emissions).

The 1980s also saw significant attention paid by analysts in industrialized countries to natural resource depletion in the developing world. This was a time of worldwide public concern over issues such as tropical deforestation, particularly in the Amazon basin. A book called *Wasting Assets* estimated that the phenomenal rates of economic growth posted by Indonesia in the 1970s and early 1980s would have been halved had they taken into account the depletion of timber, oil, and other resources.²³ In the aftermath of this work, many developing countries became interested in adjusting NNP to account for natural resource depletion. In 2012, the United Nations Statistical Commission developed a standard accounting framework for natural capital. Experiments are ongoing in the Philippines, Namibia, and other developing countries.

A relative latecomer to this process, the United States made its first official effort to incorporate resource depletion into the national income and product accounts in 1994. This initial effort included only selected mineral commodities, including oil, gas, and coal. Opposition from the U.S. Congress resulted in the suspension of further efforts on green

accounting, despite a highly favorable report by a National Research Council external review.²⁴ The congressionally imposed ban on expanding the national income and product accounts to include natural resource depletion remains in place in 2015. But that has not stopped academics from studying how green accounting might affect our understanding of the U.S. economy. A recent analysis suggests that oil- and coal-fired power plants, along with some other major sources of air pollution, may cause damages that exceed the value of their economic output.²⁵

General Conclusions from Theory and Experience

There are some general conclusions to be drawn from past and current experience with greening the national income and product accounts. First, the difference between green NNP and traditional NNP tends to vary by country and over time. An economy's greater dependence on resource extraction tends to create a greater difference between the two measures. This is especially true of green NNP calculations that include natural resource extraction but exclude environmental externalities such as air and water pollution. For example, in the calculations of NNP that were done for the United States in 1994, subsurface mineral depletion appears to have been approximately counterbalanced by exploration and discovery of new resources. In contrast, in the late 1980s the green NNP calculated for Indonesia by Repetto and others departed quite drastically from measures of traditional NNP. The picture for the United States might be quite different during an earlier period of its development, when natural resource extraction made up a larger share of total economic activity.

Second, traditional NNP not only leaves out natural resource depletion and environmental externalities but also omits the contribution to NNP of technological progress. Rough calculations indicate that including technological progress in the calculation of NNP would result in a substantial upward correction, perhaps by as much as 40 percent.²⁶ Including both resource depletion and technological progress in national income and product accounting, used to compare national economic conditions across countries and to characterize the direction and pace of global economic growth, would be an "almost practical step toward sustainability."²⁷

Note that searching for a way to keep track of what is happening to the world's natural capital within individual nations' annual balance sheets will leave out changes in some very important assets: global commons, such as oceans, stratospheric ozone, and the upper atmosphere (currently

GDP versus “Genuine Wealth”

A group of economists and ecologists have compared actual rates of change in traditional GDP, 1970 through 2000, with those of their own estimated measure of “genuine wealth,” which incorporates changes in the stocks of commercial forests, oil and minerals, and air pollution, as well as population growth, investment in human capital through education, and technological change.²⁸ In the analysis, investments in education and technological change can increase the magnitude of growth estimates relative to the traditional measures, whereas natural resource depletion and rapid population growth can decrease them.

Rates of change in estimated per capita genuine wealth over this period in the United States, Bangladesh, India, Nepal, and Pakistan are positive but significantly lower than the growth rate of traditional per capita GDP, suggesting that education and technological investments have not outweighed the negative effects of resource depletion and population growth. For the United Kingdom, the growth rate of estimated per capita wealth is approximately the same as that of traditionally measured per capita GDP. In China, the more comprehensive wealth measure has grown at a rate substantially higher than traditional GDP. In sub-Saharan Africa and the oil-exporting Middle East/North Africa region, the more comprehensive wealth measure actually falls between 1970 and 2000, in contrast to positive measures of economic growth using traditional metrics.

The authors point out that concluding from this analysis that the poor countries, particularly in sub-Saharan Africa, are consuming too much and that the rich countries, such as the United Kingdom and the United States, are following a more sustainable path would be a mistake. In fact, the poor countries suffer from too little consumption, as well as too little investment; the rich countries may, in contrast, be growing through the import of natural resources and resource-intensive products from poor countries (thereby avoiding depleting their own resources).

a major receptacle for the world’s greenhouse gas emissions). In addition, estimation of green NNP on the national scale masks the depletion of natural resources through trade; resource-adjusted NNP figures for poor countries that deplete forests or minerals to fuel economic growth will be diminished by this practice, but the same figures for rich countries that import these products will not. Estimates of *global* green NNP, taken regularly, would be necessary to illuminate the tradeoffs taking place between natural and other assets in the pursuit of economic growth.

Trade, Growth, and the Environment

The removal of barriers to trade was one of the hallmarks of the late twentieth century. Regional agreements such as the North American Free Trade Agreement (NAFTA) took shape, along with further evolution of the Global Agreement on Tariffs and Trade (GATT) and then the World Trade Organization (WTO). Economic integration has continued in the twenty-first century. Open economies generate higher levels of social welfare—this is a bedrock economic principle. How does the phenomenon of trade interact with natural systems?

Economic theory suggests that efforts by some countries to internalize the costs of pollution through environmental regulation will alter international trade patterns, resulting in the export of dirty industries to countries with less stringent regulations. This has been called the *pollution havens hypothesis*. The underlying theory and intuition behind this concept are strong, yet through the mid-1990s, studies found no evidence of the impacts of environmental regulations on trade patterns.²⁹ This might be due to several factors. First, environmental regulatory compliance represents a small share of production costs for most industries. Second, many countries' largest trade partners have similar levels of environmental regulatory stringency. Third, some industries simply are not very mobile; power plants, for example, may spend a lot to comply with air pollution regulations, but in most cases it would be inefficient or even impossible for electricity generators to relocate to neighboring countries in response to regulation. These factors make it difficult for researchers to detect the “signal” of regulatory costs in data on production and trade.

Additional statistical complications are at work as well: unobservable heterogeneity and endogeneity.³⁰ By *heterogeneity* we mean that unobserved industry and country traits can be correlated with the likelihood of regulation and the export of pollution-intensive goods. This poses the danger of seeing in the data a causal link between trade and pollution where none really exists. By *endogeneity* we mean that the direction of causation between trade, regulation, and pollution is unclear. Environmental regulation often appears in countries with substantial international trade. Does trade influence pollution regulation? If so, this is the opposite effect suggested by the pollution havens hypothesis, which posits that regulation influences trade.

Recent studies that surmount these challenges support the hypothesis that environmental regulation may contribute to the export of some highly mobile industries, between trading partners with significant differences in regulatory stringency.³¹ Small effects have also been measured in trade between U.S. states.³² Taken together, although these results show that there is some effect of regulatory costs on firm location, the effect is small and restricted to a handful of industries.

Trade, Growth, and the Environment *continued*

Note that trade can also influence rates of natural resource extraction in many ways.³³ If local producers of fish, timber, and other resources suddenly gain access to world markets for these goods, and world prices are higher than local prices, greater exploitation may ensue. For example, increased trade between Europe and the United States in the nineteenth century may have accelerated the slaughter and near-extinction of the American bison.³⁴ Without regulatory and property rights structures in place to deal with externalities and public goods, increased pressure on natural resources from trade can reduce welfare.

Are Economic Growth and Sustainability Compatible?

Among the competing definitions of sustainability are many that conflict with the economic definition we offered earlier. In fact, the economic definition has been called “weak sustainability,” in contrast to “strong sustainability,” which adds the additional requirement that natural resource stocks not be further depleted nor environmental quality further degraded in pursuit of economic growth. This constraint amounts to a restricted view of substitutability between natural and physical or human capital.

In their concept of sustainability, economists see substitution and technological change as the rule and limits to these forces in supporting economic growth as exceptions. An alternative view, a perspective often labeled *ecological economics*, holds binding scarcity as the rule and sees the ability of substitution and technological change to overwhelm scarcity in pursuit of economic growth as the exception—something possible, perhaps, over the first few millennia of human development but not indefinitely.³⁵ Adherents of this view see serious conflicts between continued economic growth and sustainability. Ecological economists worry that the pursuit of economic growth and development for their own sake overlook the critical importance of the appropriate scale of economic activity.

These are two very different worldviews. One may subscribe to either of these views on sustainability and still draw important lessons from the economic definition of sustainability offered in this chapter. For example, getting prices right, justified on efficiency grounds alone, is also a step in the direction toward either of these concepts of sustainability. The issues of intergenerational and intragenerational equity we raise (such as the contrast between consuming and investing resource rents) are also relevant to either view of sustainability.

Conclusion

In this chapter, we have explored the links between the environment and the macroeconomic phenomenon of economic growth. We began with a big-picture look at the issue of nonrenewable resource scarcity that we approached from the level of individual firms in Chapter 6. Some have postulated that the physically limited stocks of nonrenewable resources such as oil and copper will eventually act as serious brakes on economic growth. We examined the historic evidence for the negative impact of scarcity on economic growth and found that this impact, though measurable, is much smaller than the positive impact of substitution and technological change.

We developed an economic definition of sustainability, which requires leaving “the world the capacity to be as well off as we are today.” The definition hinges on substitution possibilities and technological change. Although there are some serious complications in applying this rule directly, it offers significant insights for public policy regarding the environment and economic growth. First, it offers additional support for getting prices right and sorting out the market failures initially discussed in Chapter 5. This is a step in the right direction, whether one adheres to the economic definition of sustainability or definitions that encompass stronger conservation principles with respect to specific natural resources. The economic definition of sustainability also speaks for investment over consumption of the rents from natural resource extraction, and it prompts us to think more clearly about the tradeoffs between achieving goals of intragenerational and intergenerational equity.

Since the creation of the practice of national income accounting, economists have worried that growth measures such as the rate of change in net national product omit many portions of the economy, including natural resource depletion and environmental degradation. We determined in this chapter that more comprehensive measures are needed, particularly if we are to apply the economic concept of sustainability, which forbids us from depleting the total global capital stock, including natural, physical, and human capital.

Although economic growth is not synonymous with growth in social welfare, it does have the capacity to increase welfare, making somewhat less painful the necessary tradeoffs in “sacrificing some of one good thing for more of another.” For example, as incomes rise, countries are better able to afford improvements in health, education, and environmental quality; whether they choose to invest in these, or in something else, can have critical consequences for social welfare in the long run.

12

Conclusion

We have now come to the end of our journey. Rather than try to summarize all that we have discussed, we will consider some of the broader implications of economic analysis for environmental policy.

What Does Economics Imply for Environmental Policy?

In June 2014, the U.S. Environmental Protection Agency (EPA) released the details of its proposed Clean Power Plan, a set of federal regulatory guidelines (that states will need to comply with) aimed at reducing carbon dioxide emissions from power plants by 30 percent, from a 2005 baseline, by 2030. Not surprisingly, the plan was received differently by those on various ends of the political spectrum. In an op-ed in *USA Today*, the coal industry decried the new set of rules as “overzealous,” harmful to middle-class and lower-income Americans, and a danger to the reliability of U.S. electrical grids.¹ The Union of Concerned Scientists calls the plan a “historic opportunity” and an “affordable solution with substantial benefits for our economy, our health, and our children’s future.”²

Such debates are common in the environmental realm (and indeed in other areas of policy, such as health care). In a sense, there is a good reason for that: Environmental regulation can impose costs and can also generate very substantial benefits. Of course, advocates on both sides may make exaggerated claims about the impacts of regulation for political effect. Rigorous economic analysis can help cut through those debates. EPA’s regulatory impact analysis of the Clean Power Plan, for example, suggests that the benefits of reducing emissions (using the social cost of carbon we discussed in Chapter 3, as well as the co-benefits from reductions in other air pollutants such as particulate matter) outweigh the costs (increased

investment in low-carbon electricity generation, in greater energy efficiency, and in other means of reducing emissions).³

Nor is economic analysis of the regulation confined to calculating costs and benefits. Indeed, much of the discussion around EPA's power plant regulations involves how states should meet the proposed federal guidelines: Energy efficiency requirements? Plant-level standards? Emission trading? The principles outlined in this book can also inform those decisions.

It is instructive to place this debate over U.S. climate change regulation on the continuum of public discussions about environmental policy. Astute journalists have recently unearthed old newspaper articles from the early days of the Clean Air Act, citing similar debates.⁴ An important difference between those debates in the 1970s and those occurring today is the role of economic analysis, which has increased in scope and quality. In the process, economic analysis itself is increasingly used by both opponents and proponents of regulation and is a target for debate.

This is all the more reason to study environmental economics. Economic theory provides strong arguments in support of active (if carefully designed) government policies in the environmental realm. For the most part, economic analysis suggests that environmental amenities are underpriced and that renewable natural resources are overexploited. For example, subsidies for fishing or the extraction of timber from public lands are inefficient, in part because they result in environmental losses. Studies of U.S. sulfur dioxide regulation have found that benefits vastly exceed costs, implying—if anything—that regulation ought to be *more* rather than less stringent.⁵

In other cases, of course, economic analysis arrives at conclusions that an environmental advocate might object to. For example, benefit-cost analyses of the U.S. Clean Water Act suggest that the act's original goal of zero emissions to all U.S. water bodies, its focus on regulating point-source pollution, and its primary reliance on technology standards have resulted in a federal regulation with substantial net costs.⁶ Many economists would argue that suburban sprawl, despite its aesthetic shortcomings, may in fact be the socially desirable reflection of individuals' willingness to pay for certain amenities (such as quiet streets and large yards).⁷

The underlying point is that environmental economics, when applied correctly, is not "green" or "brown" but neutral. Economists treat natural resources and environmental amenities, in effect, like any other assets. On one hand, this approach brings environmental quality onto the balance sheet, ensuring that it is not given a value of zero in public debates over

whether to extract oil from pristine wilderness areas or whether to allow the construction of new housing in a forested or agricultural urban fringe. On the other hand, treating natural resources and environmental amenities in the same framework as other goods and services reflects an underlying assumption that substitution possibilities are just as relevant to environmental and natural resources as they are to other assets. To be sure, there may be many specific resources (such as the Grand Canyon, or portions of the Amazon rainforest) that should be preserved for their own sake even on efficiency grounds. But in many other cases it will be efficient to convert forested lands to urban or agricultural use, to deplete a nonrenewable groundwater aquifer, or to tolerate some pollution in exchange for the services provided by a polluting industry. Efficiency and cost-effectiveness analysis in all of these cases should serve as a starting point for discussion about environmental and natural resource management policies—not the final word.

The Roles of Firms, Consumers, and Governments

We read a lot these days about the power of consumers to affect the practices of firms whose activities may result in pollution or other natural resource damages, and about “corporate social responsibility” and its implied voluntary measures to reduce the environmental impacts of economic activity. There are certainly instances in which firms improve their environmental performance beyond what is required by law, and there are cases in which consumer pressure has resulted in substantial improvements in environmental and natural resource outcomes. Nonprofit environmental advocacy organizations also play an important role in this process.

However, an important message of this book is that as long as the markets for environmental amenities are incomplete, consumer pressure and voluntary efforts by companies to reduce their impacts on the environment will not be sufficient to achieve the efficient level of pollution control or efficient natural resource management practices. The incentive structure created by the “wrong prices”—prices that do not reflect the full social cost of engaging in an environmentally damaging activity—is simply too powerful.

For example, millions of Americans are Sierra Club members, but does each member contribute annually an amount equal to his or her full willingness to pay for wilderness preservation and the organization’s other ambitious goals? Given our discussion of public goods and free riding, one might suspect that the answer is “no” and that there are probably many

nonmembers of the Sierra Club who have some willingness to pay for the services it provides. Similarly, many firms trumpet their activities in environmental stewardship, and many of these efforts are sincere and generate substantial environmental benefits. But we cannot rely on these voluntary activities alone to correct the substantial market failures that contribute to inefficient environmental degradation and resource depletion.

Where markets are absent or otherwise incomplete, well-designed public policies, such as market-based instruments for pollution control, are needed to correct these incentives and get the prices right. These policy instruments use market principles to correct market failures and align private incentives with public ones. This is not to say that government regulation is always the answer to market failure in the environmental realm. In fact, in this volume we have noted a number of examples of government failure—the role of government in worsening environmental outcomes by subsidizing resource extraction, for example. But from the economic perspective, government is a necessary central authority with the power to create and enforce property rights structures that promote resource stewardship rather than excessive depletion, tax citizens for the provision of public goods, and implement and enforce regulations that internalize the external costs of pollution. Without this, the voluntary actions of firms and active participation of citizens can ameliorate the impact of incomplete markets on environmental degradation, but they cannot eliminate it.

Some Final Thoughts

Good economic analysis formalizes and makes transparent the difficult compromises inherent in decisions about the use and management of natural resources and environmental amenities. Applying economic principles to environmental policy choices comes as naturally to economists as doing so in choices about other aspects of the economy. The “environment” is not separate from the “economy” in the framework we have offered in this text; indeed, environmental problems cannot be fully understood without a basic intuition for how markets function and how they fail.

Correcting market failures in this realm, as in any other, is efficient. It is good for the economy, in part because of the benefits generated by the resulting environmental improvement. This is not to say that there are no tradeoffs to be made. Using cap-and-trade, taxes, and other market-based approaches to correct market failures can reduce the total costs of environmental regulation and natural resource management, but they

cannot eliminate them. As we have seen, rigorous economic analysis can also inform an understanding of those tradeoffs. Internalizing the climate change externality of CO₂ emissions from power plants may increase the cost of electricity and energy-intensive goods such as cement and aluminum while securing benefits in the form of lower future damages from climate change. Even if regulation has benefits greater than costs, those benefits and costs may not be equally distributed. Reducing SO₂ from power plants may reduce employment in high-sulfur coal mining regions while reducing acid rain and providing cleaner air in downwind areas. Establishing property rights over an open-access fishery may drive some high-cost fishers out of fishing altogether, even as it raises the overall returns to the local economy.

Economic analysis, combined with careful consideration of equity issues, shrewd political strategy, and other inputs, will help students of environmental studies to make better decisions about environmental policy and to better interpret the consequences of others' decisions. Although efficiency and cost-effectiveness are not the only criteria that contribute to sound environmental policies, they can help us make conscious choices about how much of one good thing must be sacrificed to have more of another, just as we do in daily decisions about our own household budgets. In this text, we have shown that economics can make vital contributions to both the analysis of environmental problems and the design of possible solutions. We hope the tools we have introduced help illuminate the environmental issues you will approach throughout your coursework and your career.

Discussion Questions

Chapter 2

1. Explain why it would not be desirable (from the point of view of economic efficiency) to eliminate all emissions of sulfur dioxide from electric power plants. Now suggest a setting in which zero pollution might be efficient.
2. In figure 2.3, abating X^{MAX} units of pollution would result in total benefits of pollution abatement greater than the total costs. Is this amount of pollution abatement efficient? Why or why not? Refer to both figure 2.3 and figure 2.7 in your answer.
3. Consider the case of Aracruz Celulose, S.A., the paper pulp manufacturer described in this chapter. Suppose a study finds that the marginal benefit of reducing chlorinated organic compounds in the effluent from pulp mills is \$0.50 per kilogram of AOX. (For simplicity, assume that this number does not change with the amount of pollution.) Which pulping technology would be efficient in this case? What would be the resulting annual costs and benefits associated with pollution abatement?
4. Explain the typical shapes of the total abatement cost and benefit curves, as well as the marginal benefit and cost curves.
5. Imagine a policy to reduce emissions of greenhouse gases that would yield a marginal net benefit, in present value, equal to \$1 million today and \$2 million next year. Would this policy be dynamically efficient? Why or why not?
6. Advocates for wilderness protection have often criticized the common practice of setting aside “rock and ice” (i.e., alpine areas) as protected land, while more productive bottomland areas (along

river valleys, for example) are typically left in agriculture or other intensive use. In general, the bottomland areas would also provide richer species habitat than the alpine areas. What are the likely pros and cons of this approach from the perspective of economic efficiency?

Chapter 3

1. How might we estimate the marginal benefits of preservation of the California condor, the endangered species discussed in chapter 2? Imagine that we estimate the marginal benefits of condor preservation to be \$200 per bird. Are there policy measures that would be excluded if we applied the equimarginal principle to condor preservation? Explain.
2. The economic benefits of an environmental policy such as the reduction of sulfur dioxide emissions from power plants are measured by the collective willingness to pay of human beings. Policies like this may have ecological benefits, such as the effects of reductions in acid rain. Discuss the degree to which measuring the economic benefits of a pollution reduction policy will capture ecological benefits.
3. Why do economists generally prefer revealed preference approaches to environmental benefit valuation over stated preference approaches? Are there cases in which stated preference approaches would be recommended?
4. Explain the difference between use and nonuse value, with reference to a particular environmental policy in which you may be interested, such as greenhouse gas emissions reductions or endangered species preservation.
5. Should environmental and natural resource management policies be put to a strict benefit–cost test? Why or why not?
6. Some environmental laws in the United States explicitly prohibit the use of benefit–cost analysis in some areas of environmental policy. For example, the Clean Air Act declares that air quality standards are to be determined purely on the basis of protecting public health with “an adequate margin of safety” and forbids the administrator of the Environmental Protection Agency from considering costs in setting standards. Can you provide a critique of such an approach, from the perspective of economic efficiency? What might be the consequences of such an approach? Now take a step back. From

your own perspective, do you think such an approach is advisable? Why or why not?

7. Contrast the economic perspective on endangered species preservation with the perspective that individual species have infinite value. What are the implications of each perspective for public policy?
8. Economic analyses indicate that reducing timber extraction in the U.S. Pacific Northwest in the 1990s, to preserve old-growth habitat for the endangered northern spotted owl, had significant net benefits. Yet many logging communities experienced significant economic dislocation as a result of these policies. Discuss the links, if any, between economic efficiency and distributional equity in this case. Was reducing timber extraction a Pareto improvement?

Chapter 4

1. Explain why a demand curve can be considered to be a marginal benefits curve and why a supply curve is equivalent to a marginal cost curve. Refer to a specific environmental problem in your answer.
2. Describe a situation in which you would expect the free market to result in an efficient outcome with respect to environmental quality or natural resource management and one in which you would not.

Chapter 5

1. Describe the efficiency loss that results when the social costs of pollution are external to the private costs of producing a good such as electricity. How do the market price and quantity of electricity with pollution externalities compare to the efficient price and quantity? What are the potential gains to society from regulating pollution?
2. Examine figure 5.1 and imagine that the marginal damages of pollution were much flatter than the curve represented here. How would this change the magnitude of the deadweight loss from pollution in this case? What are the intuitive implications?
3. Contrast a pure public good, such as biodiversity preservation, with an open-access resource, such as some fisheries and groundwater aquifers. How do these classes of goods and services differ, and what are the implications for environmental policy?
4. Many people contribute money to environmental advocacy organizations. Can we measure the benefits of the services provided by these organizations by summing up contributions? Why or why not?

5. Explain the link between public goods and positive externalities.
6. In the international environmental treaty “game” described by figure 5.5, both countries would be better off if each contributed to the cleanup of a shared pollution problem, but this is not what we expect to happen. Why is this better outcome unlikely to occur?

Chapter 6

1. If we take the world’s known reserves of oil and divide this total quantity by average annual oil consumption, we obtain the reserves-to-use ratio, the number of years that remain before exhaustion of our oil resources. Explain why this ratio paints a misleading picture of oil scarcity.
2. The Hotelling Rule states that marginal user cost rises at the rate of interest. Explain the intuition behind this result.
3. Under what conditions would you expect the extraction rate of a nonrenewable natural resource to depart from the dynamically efficient rate?
4. Oil and endangered species are both natural resources with high economic value. Yet a private landowner in the United States might react very differently to the discovery of an oil well on her property than she would to the discovery of an endangered species population. Explain this difference, using economic concepts.
5. Economist Robert Solow has said that “the monopolist is the conservationist’s friend.” Explain this in the context of nonrenewable resource extraction.

Chapter 7

1. The models we explore in Chapters 6 and 7 (nonrenewable resources, fisheries, and forests) all refer to the concept of economic rent. Define rent and explain its relationship to economic efficiency.
2. Explain the relationship between the biological timber rotation, the Wicksell rotation, and the Faustmann rotation. Which one is economically efficient, and why?
3. Forests generate nontimber benefits. Explain the effects of the following nontimber values on the efficient timber rotation: (a) new growth provides habitat for white-tailed deer, (b) old-growth forest provides habitat for the endangered red-cockaded woodpecker, and (c) decaying trees provide habitat for an important insect species. How should policymakers incorporate competing forest values such as these into forest management decisions?

4. Explain the link between the establishment of property rights and deforestation in tropical countries. Given our discussion in this chapter, does this imply that large-scale privatization of land in tropical regions is a solution to deforestation?
5. A fishery will always be economically overfished before it is biologically overfished, and therefore the economically efficient level of fishing effort is lower than the biologically efficient level. Explain why this is true.
6. Under open access, fishing will occur until the total benefits are exactly equal to total costs, and net benefits are equal to zero. Why don't fishers stop entering the fishery before this happens?
7. Explain precisely why reducing fishing effort from the open-access equilibrium would be efficient. What would be the effect on fishing communities in the short run? In the long run?
8. Imagine that a national government reduces the marginal cost of fishing effort through a subsidy, such as a fuel tax exemption. How would this subsidy affect the open-access equilibrium level of fishing effort, represented in figure 7.6? Would the net benefits of the fishery still be zero in this case?
9. History offers many examples of small groups of farmers collectively managing shared irrigation systems, yet there are also many examples of inefficient depletion of groundwater aquifers due to agricultural irrigation. Use the concepts of common property and open access to describe a potential explanation for these two very different phenomena.

Chapter 8

1. The Coase Theorem suggests that, under some conditions, private bargaining will resolve negative externalities. Does this mean that pollution problems should be left for the market to solve? Why or why not? Discuss this from the perspective of efficiency, as well as equity.
2. According to Ian Parry and Kenneth Small, the efficient gasoline tax in the United States is somewhat less than \$1 a gallon (see Parry and Small, 2005, cited in the Further Reading for Chapter 8). Of this amount, only 6 cents is related to global climate change from carbon dioxide emissions. Some people might argue that 6 cents is far too little to change people's behavior, for instance by inducing them to drive less or buy more fuel-efficient cars (and even a dollar might not make much difference). If the Pigovian tax turns out not

- to make much difference in what people do, is it still efficient? Why or why not?
3. How would a landing tax in a fishery work to reduce the amount of fish caught?
 4. Many economists have concluded that the marginal benefits associated with reducing greenhouse gas emissions are flat, because each additional ton of carbon dioxide (or its equivalent) has the same effect on global warming. At the same time, there is uncertainty about the future marginal costs of controlling greenhouse gases. What do these two assertions imply about the choice between a tax and a tradable permit system to reduce carbon dioxide emissions, on the grounds of efficiency?
 5. Much of the political debate surrounding the use of emission trading concerns the allocation of the pollution allowances—in particular, whether to auction the allowances or give them away for free. Explain why the method of allocation does not affect how much pollution firms end up controlling, assuming that transaction costs are low. What do you think would happen if transaction costs were high?

Chapter 9

1. Explain the intuition behind why market-based instruments (emission taxes and tradable permits) are cost-effective, whereas uniform standards are generally not.
2. Redraw figure 9.1 using an even flatter marginal abatement cost curve for firm A and a steeper one for firm B. What would happen to the cost-effective allocation? What would happen to the size of the cost savings from a market-based instrument, relative to a uniform standard? Explain.
3. The 1977 Clean Air Act Amendments required new electric power plants to install scrubbers in order to remove sulfur dioxide emissions. Such an approach is often called a “technology-forcing” approach and is typically promoted as a way of ensuring that polluters install the most advanced or best available abatement technology. From an economic perspective, what kind of incentive does such a policy provide for the development and adoption of new technologies?
4. The cap-and-trade program for carbon dioxide emissions under California’s AB 32 climate legislation includes both a maximum emissions permit price (a “safety valve”), and a minimum permit

price (a “price floor”); the two together are often described as a “price collar.” How might you expect these two additions to the standard cap-and-trade model described in Chapter 8 to affect eventual emissions under the policy? How might they affect the cost-effectiveness of cap-and-trade, relative to a uniform performance standard? What about long-run incentives to adopt new pollution abatement technologies?

5. Why are hot spots a potential problem with market-based instruments? Why would location-specific taxes (or trading ratios in a cap-and-trade program) help alleviate the problem?

Chapter 10

1. Why was the U.S. sulfur dioxide allowance trading program widely considered to be a success? In your answer, be sure to discuss the policy’s environmental performance, cost-effectiveness in comparison to other potential policies, compliance and enforcement, and distributional implications.
2. One of the strongest objections to market-based fishery management is the possibility of consolidation. Describe this phenomenon and explain why it may occur under an individual fishing quota (IFQ) system. To what degree has consolidation occurred in New Zealand fisheries managed by IFQs? Is it a concern for efficiency, distributional equity, or both? The government of New Zealand has taken some measures to reduce the impact of consolidation. Describe these measures, and discuss the equity–efficiency tradeoff they imply.
3. Draw an analogy between a cap-and-trade program for air pollution emissions and a market for water consumption permits during a drought. How would these two policies be similar? How would they differ? Be sure to address the source of the potential gains from a market-based approach in each of these cases.
4. Numerous successful tradable permit systems for air pollution control have emerged in the United States in the past two decades, but experiments with water quality permit trading have been much less successful. Offer some potential reasons for this difference.
5. Is a pay-as-you-throw policy for solid waste management a Pigouvian tax? Why or why not? Have these policies been successful?
6. What are some of the barriers to large-scale application of market-based policy approaches to land management and species preservation? To what degree can these barriers be overcome?

Chapter 11

1. The *Limits to Growth* model and economic models differ significantly in their assumptions. Describe these different assumptions and the resulting differences in what the models suggest about the limits on future economic growth posed by the finite availability of important nonrenewable resources such as oil and coal.
2. The world's supply of oil is being depleted much faster than the rate of natural regeneration. From an economic perspective, can this be efficient? Can it be sustainable?
3. Economist Robert Solow describes green accounting as an “almost practical step toward sustainability.” Why is it *almost* practical? How might green accounting promote sustainability from an economic perspective?
4. If we were to try to implement the economic concept of sustainability, we would face some important sources of uncertainty. Describe these areas of uncertainty and how they might limit our ability to implement sustainable policies.
5. What are the most important insights of economic sustainability for current policies regarding natural resources and the environment?
6. In 1976, the state of Alaska established the Alaska Permanent Fund, valued at about \$54 billion in 2015, which primarily uses the returns from investing the proceeds of the sale of oil to provide Alaska residents with dividends (averaging \$1,365 over the past 15 years). The Permanent University Fund in Texas (valued at \$17 billion in 2014) uses proceeds from the sale of oil leases and royalties on state land as endowment funds for several state universities. Assess these policies from the perspective of economic sustainability.

References

Chapter 1: Introduction

1. Intergovernmental Panel on Climate Change, Working Group I, *Climate Change 2013: The Physical Science Basis* (New York: Cambridge University Press, 2013), 161–162.

Chapter 2: Economic Efficiency and Environmental Protection

1. These arguments and many more are summarized in Hunt Allcott and Michael Greenstone, “Is There an Energy Efficiency Gap?” *Journal of Economic Perspectives* 26(1): 3–28 (2012). One commonly cited source on the existence of the energy efficiency gap is McKinsey & Co., “Unlocking Energy Efficiency in the U.S. Economy,” available at http://www.mckinsey.com/client_service/electric_power_and_natural_gas/latest_thinking/unlocking_energy_efficiency_in_the_us_economy, accessed February 27, 2015.
2. See Nathaniel O. Keohane, Benjamin Van Roy, and Richard J. Zeckhauser, “Managing the Quality of a Resource with Stock and Flow Controls,” *Journal of Public Economics* 91: 541–569 (2007).
3. We have given a somewhat loose explanation for why marginal cost equals the slope of the total cost function. A precise explanation requires calculus: Marginal cost is (by definition) the first derivative of the cost function, which in turn is the slope of the total cost function. But even if you haven’t had calculus, we can make the discussion a bit more rigorous. For any given change in abatement, the change in total cost is the same as the change in the height of the cost function. To see this, refer back to the cost function in figure 2.1. Choose two nearby levels of abatement, which we shall call point *A* and point *B*. Suppose we increase abatement of some pollutant from *A* to *B*. The change in total cost that results is obviously the total cost at *B* minus the total cost at *A*. For such large changes, we need to take into account the curvature of the cost function. As the distance between *A* and *B* gets smaller, however, the curvature matters less and less. Indeed, for very small changes in abatement the cost function is almost a straight line, and the change in height between *A* and a nearby point can be measured by the slope of the function at *A* times the distance between the points. In particular, the change in total cost

- between A and $A + 1$ (assuming that we are measuring units so that one unit is sufficiently small) equals the slope of the function at A . But we defined marginal cost as the change in cost from one more unit of abatement. Thus the marginal cost at point A is just the slope of the total cost function at that point.
4. The cost figures for Aracruz are derived from estimates in Jackie Prince Roberts, "Aracruz Celulose, S.A.," Harvard Business School Case 9-794-049 (January 1995). In particular, the numbers in the table draw on the estimated costs for upgrading Mill A and from current authors' estimates of 4 kg/ton AOX for standard pulping, 0.8 kg/ton for ECF, and 0.04 kg/ton for TCF. Following Roberts, we have applied a 10 percent discount rate to capital costs in order to combine them with variable costs.
 5. This discussion is based on Nathaniel O. Keohane, Benjamin Van Roy, and Richard J. Zeckhauser, "The Optimal Management of Environmental Quality with Stock and Flow Controls," John F. Kennedy School of Government, Faculty Working Paper #RWP05-042, June 2005, available at <https://research.hks.harvard.edu/publications/workingpapers/citation.aspx?PubId=3037>.
 6. These events are summarized from Scott Barrett, *Environment and Statecraft: The Strategy of Environmental Treaty-Making* (Oxford: Oxford University Press, 2003), Chapter 8: 221–253.
 7. Many other scenarios were analyzed also, including a 50 percent reduction in U.S. CFC consumption and production, with no mitigation efforts by other countries. All of the scenarios analyzed, at discount rates from 1 to 6 percent and using alternative assumptions about the value of averted skin cancer mortality, had very significant domestic net benefits. The EPA study is described at length in James K. Hammitt, "Stratospheric-Ozone Depletion," In: Richard D. Morgenstern, ed., *Economic Analyses at EPA: Assessing Regulatory Impact* (Washington, DC: Resources for the Future, 2003), 131–169.
 8. In fact, summing up the areas of these rectangles would only give an approximation of the actual area under the curve. But as we made the rectangles narrower and narrower (and drew more and more of them to fill in abatement from zero to X_i) the total area of those rectangles would be a better and better approximation of the area under the curve. Eventually, as the rectangles became increasingly narrow—"in the limit," as mathematicians say—this approximation would become perfect.
 9. To be more precise, if the efficient or optimal policy in one period is independent of the policies in all other periods, the problem to be solved is a static efficiency problem. If optimal policies are correlated across periods, the efficiency problem to be solved is dynamic. Note that a problem can be static in this sense even if it occurs over a long period of time, as long as the benefits and costs do not change from period to period.
 10. The three explanations for the time value of money given in the text correspond, respectively, to what economists call *pure time preference*, *the marginal utility of income*, and *the rate of return on capital*. The often-cited Ramsey equation essentially states that along an optimal policy path, the social discount rate implied by time preference and the marginal utility of income equals the rate of return on capital.
 11. For some very useful treatments of discounting in an environmental context, see Goulder and Stavins (2002), Brennan (1999), and Portney and Weyant (1999).

Chapter 3: The Benefits and Costs of Environmental Protection

1. If waste products can be reused or sold as inputs to others' production processes, these benefits of a regulation would be "netted out" of costs in a benefit–cost analysis. For example, some municipal wastewater treatment plants sell nutrients such as nitrogen and phosphorus, derived from treated sewage, for use as fertilizer. A regulation that required increased removal of these nutrients from a plant's effluent before its disposal in a river would have a cost associated with the pollution abatement and also a benefit associated with selling the recovered nutrients. These kinds of negative regulatory costs have given rise to a literature that suggests that environmental regulation yields net benefits even without considering the *environmental* benefits—that because of increased profits in industries related to pollution control and stimulation of product and process innovations in the regulated industries, firms' profits increase, on average, with regulation. This vision of a regulatory "free lunch," though appealing, is not shared by most economists. Although these kinds of impacts may occasionally occur, they do not prove a general pattern. Regulation is justified economically when the benefits exceed the costs; by definition, this occurs if the costs themselves are negative, but true negative costs are likely to be very rare. For views on both sides of this debate, see Michael E. Porter and Claas van der Linde, "Toward a New Conception of the Environment–Competitiveness Relationship," *Journal of Economic Perspectives* 9: 97–118 (1995); and Karen Palmer, Wallace E. Oates, and Paul R. Portney, "Tightening Environmental Standards: The Benefit–Cost or No–Cost Paradigm?" *Journal of Economic Perspectives* 9: 119–132 (1995).
2. Analysts may need to decide how to attribute capital investments over time and how to prorate costs that serve multiple functions. For example, the cost of an environmental engineer or lawyer may be spread across many of a firm's environmental, health, and safety objectives, not just compliance with a single regulation. For descriptions of well-accepted methods regarding these and other cost estimation issues (such as how to deal with uncertainty), see Chapter 8 in National Center for Environmental Economics, U.S. Environmental Protection Agency, *Guidelines for Preparing Economic Analyses* (Washington, DC: U.S. EPA, 2010, updated 2014), available at: [http://yosemite.epa.gov/ee/epa/eer.nsf/vwAN/EE-0568-50.pdf/\\$file/EE-0568-50.pdf](http://yosemite.epa.gov/ee/epa/eer.nsf/vwAN/EE-0568-50.pdf/$file/EE-0568-50.pdf), accessed February 18, 2015.
3. If you have taken introductory microeconomics, you will realize that firms' ability to pass through these regulatory cost increases to consumers in their prices will depend on the relative elasticities of demand and supply.
4. Note that price increases are not the only potential source of losses in consumer surplus from environmental regulations. Losses in consumers' utility—the well-being they obtain—from using regulated goods and services can also occur. Consider, for example, many consumers' dislike of early compact fluorescent lamps (CFLs), relative to the incandescent lightbulbs they have replaced.
5. Jon Creyts, Anton Derkach, Scott Nyquist, Ken Ostrowski, and Jack Stephenson, "Reducing U.S. Greenhouse Gas Emissions: How Much at What Cost?" McKinsey & Company (2007).
6. For a thorough discussion of top–down versus bottom–up economic models, see Terry Barker, Leena Srivastava, et al., "Sector Costs and Ancillary Benefits of

- Mitigation,” Chapter 9 in B. Metz, O. Davidson, R. Swart, and J. Pan, eds., *Climate Change 2001: Mitigation*, Contribution of Working Group III to the Third Assessment Report of the Intergovernmental Panel on Climate Change (Cambridge, UK: Cambridge University Press, 2001), 561–599.
7. Stephen J. Davis and John Haltiwanger, “Gross Job Flows,” Chapter 41 in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Volume 3 (Amsterdam: Elsevier Science, 1999), 2711–2805.
 8. For an excellent discussion of the social cost of carbon, its usefulness for regulatory analysis, and the challenges involved in estimating it, see Richard L. Revesz, Peter H. Howard, Kenneth Arrow, Lawrence H. Goulder, Robert E. Kopp, Michael A. Livermore, Michael Oppenheimer, and Thomas Sterner, “Improve Economic Models of Climate Change,” *Nature* 508: 173–175 (April 10, 2014).
 9. U.S. Government Interagency Working Group on the Social Cost of Carbon, “Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866” (May 2013).
 10. On the United States, see Richard Morganstern, William Pizer, and Jhih-Shyang Shih, “Jobs versus the Environment, an Industry-Level Perspective,” *Journal of Environmental Economics and Management* 43(3): 412–436 (2002). On the United Kingdom, see Matthew A. Cole and Rob J. Elliott, “Do Environmental Regulations Cost Jobs? An Industry-Level Analysis of the UK,” *B.E. Journal of Economic Analysis and Policy: Topics in Economic Analysis and Policy* 7(1): 1–25 (2007).
 11. See Reed Walker, “The Transitional Costs of Sectoral Reallocation: Evidence from the Clean Air Act and the Workforce,” *Quarterly Journal of Economics* 128(4): 1787–1835 (2013).
 12. For discussion of the role of employment impacts in benefit–cost analysis, see Michael Livermore, Elizabeth Piennar, and Jason A. Schwartz, “The Regulatory Red Herring: The Role of Job Impact Analyses in Environmental Policy Debates,” Institute for Policy Integrity, New York University School of Law, 2012; and Eric Posner and Jonathan Masur, “Regulation, Unemployment, and Cost–Benefit Analysis,” John M. Olin Program in Law and Economics Working Paper no. 571, University of Chicago Law School, 2011.
 13. The original idea for this method came from Harold Hotelling, a statistician who made many other contributions to environmental economics, particularly with respect to the nonrenewable resource management issues we will discuss in Chapter 6. In 1947, the director of the U.S. National Park Service wrote a letter to leading economists and others (including Hotelling), asking them to suggest ways to estimate the value of the park system to the American public. Hotelling replied with a letter suggesting that the travel cost an individual incurs to visit a site could be exploited to estimate an implicit price for that site’s services.
 14. See Lucija Muehlenbachs, Elisheba Spiller, and Christopher Timmins, “The Housing Market Impacts of Shale Gas Development,” *American Economic Review* (in press).
 15. The sources of bias in CV, and the debate that has played out in the economics profession and elsewhere on its potential usefulness in policymaking, are summarized in Timothy C. Haab, Matthew G. Interis, Daniel R. Petrolia, and John C. Whitehead, “From Hopeless to Curious? Thoughts on Hausman’s ‘Dubious to

- Hopeless' Critique of Contingent Valuation," *Applied Economic Perspectives and Policy* 35(4): 593–612 (2013).
16. For a summary of the CV analysis done for the State of Alaska, see Richard T. Carson, Robert C. Mitchell, Michael Hanemann, Raymond J. Kopp, Stanley Presser, and Paul A. Ruud, "Contingent Valuation and Lost Passive Use: Damages from the Exxon Valdez Oil Spill," *Environmental and Resource Economics* 25: 257–286 (1993). For a collection of articles criticizing CV, presented at a conference sponsored by Exxon, see Jerry A. Hausman, ed., *Contingent Valuation: A Critical Assessment* (Amsterdam: North-Holland, 1993).
 17. See Catherine L. Kling, Daniel J. Phaneuf, and Jinhua Zhao, "From Exxon to BP: Has Some Number Become Better than No Number?" *Journal of Economic Perspectives* 26(4): 3–26 (2012); and Jerry Hausman, "Contingent Valuation: From Dubious to Hopeless," *Journal of Economic Perspectives* 26(4): 43–56.
 18. Richard D. Morgenstern, "Decision Making at EPA: Economics, Incentives, and Efficiency," draft conference paper in *EPA at Thirty: Evaluating and Improving the Environmental Protection Agency* (Durham, NC: Duke University Press, 2000), 36–38.
 19. Robert Costanza, Ralph d'Arge, Rudolf de Groot, Stephen Farber, Monica Grasso, Bruce Hannon, Karin Limburg, Shahid Naeem, Robert V. O'Neill, Jose Paruelo, Robert G. Raskin, Paul Sutton, and Marjan van den Belt, "The Value of the World's Ecosystem Services and Natural Capital," *Nature* 387: 253–260 (May 15, 1997).
 20. See Michael Toman, "Why Not to Calculate the Value of the World's Ecosystem Services and Natural Capital," *Ecological Economics* 25: 57–60 (1998), 58. Curiously, Costanza et al. acknowledge this very point at the beginning of their article. They write, "It is trivial to ask what is the value of the atmosphere to humankind, or what is the value of rocks and soil infrastructure as support systems" (255). But as we explain in the text, they end up confusing marginal and total values.
 21. Summaries of many important regulatory impact analyses, and the politics and economics surrounding them, are included in Richard D. Morgenstern, ed., *Economic Analysis at EPA: Assessing Regulatory Impact* (Washington, DC: Resources for the Future, 1997).
 22. For an interesting discussion of benefit–cost analysis as applied to the environment, see "Cost–Benefit Analysis: An Ethical Critique," by Steven Kelman, with replies by James DeLong, Robert Solow, and Gerard Butters, John Calfee, and Pauline Ippolito. *AEI Journal on Government and Social Regulation* (January/February 1981): 33–40. Reprinted in Robert N. Stavins, ed., *Economics of the Environment: Selected Readings*, 5th ed. (New York: Norton, 2005), 260–275.
 23. This discussion is based on Albert L. Nichols, "Lead in Gasoline," in Richard D. Morgenstern, ed., *Economic Analyses at EPA: Assessing Regulatory Impact* (Washington, DC: Resources for the Future, 1997), 49–86.
 24. Kenneth J. Arrow et al., "Is There a Role for Benefit–Cost Analysis in Environmental, Health, and Safety Regulation?" *Science* 272: 221–222 (April 12, 1996). Reprinted in Robert N. Stavins, ed., *Economics of the Environment: Selected Readings*, 5th ed. (New York: Norton, 2005), 249–254.

25. William Nordhaus, "Discounting and Public Policies That Affect the Distant Future," in Paul Portney and John Weyant, eds., *Discounting and Inter-generational Equity* (Washington, DC: Resources for the Future, 1999).
26. See Solow's response to Kelman, cited in note 22.
27. See Dallas Burtraw, Alan J. Krupnick, Erin Mansur, David Austin, and Deirdre Farrell, "The Costs and Benefits of Reducing Air Pollutants Related to Acid Rain," *Contemporary Economic Policy* 16: 379–400 (October 1998).
28. This discussion is based on Andrew Metrick and Martin L. Weitzman, "Patterns of Behavior in Endangered Species Preservation," *Land Economics* 72(1): 1–16 (February 1996); and Metrick and Weitzman, "Conflict and Choices in Biodiversity Preservation," *Journal of Economic Perspectives* 12(3): 21–34 (Summer 1998).
29. It is worth pointing out that strict Pareto efficiency is much less attractive as a general welfare criterion for evaluating overall outcomes, such as a particular equilibrium achieved by a market economy. Indeed, dramatically unequal distributions of income can be deemed Pareto efficient as long as distributing money from the rich to the poor would make at least one rich person less well off, as Amartya Sen has famously pointed out. See Amartya K. Sen, *Collective Choice and Welfare* (San Francisco, CA: Holden–Day, 1970). Note that Sen's critique is essentially a much more forceful version of the argument made in the text—namely, that applying the strict Pareto criterion in the real world would result in too much loyalty to the status quo, regardless of the attractiveness of alternative proposals.
30. In fact, this equivalence between maximizing net benefits and satisfying the potential Pareto criterion, though commonly asserted in the economics literature, turns out not to be precisely correct. Paul Milgrom has pointed out that the equivalence does not hold when people are altruistic, that is, when they care directly about the utility of others. Suppose person A cares about person B's utility in addition to her own. Then a policy that provides positive net benefits, with B receiving the gains and A bearing the losses, might fail the potential Pareto criterion: Any transfers from B to A would lower B's utility and so diminish the gain to A, possibly ruling out a strict Pareto improvement. Although the logic of this argument is airtight, it is also somewhat contrived. In our view, the close link between maximizing net benefits and the Kaldor–Hicks criterion continues to offer a strong normative grounding for benefit–cost analysis. See Paul Milgrom, "Is Sympathy an Economic Value? Philosophy, Economics, and the Contingent Valuation Method," in Jerry A. Hausman, ed., *Contingent Valuation: A Critical Assessment* (Amsterdam: North–Holland, 1993), 417–435.

Chapter 4: *The Efficiency of Markets*

1. Limbaugh, *The Way Things Ought to Be* (New York: Pocket Books, 1992), 156.
2. See Amory B. Lovins and Boyd Cohen, *Climate Capitalism* (New York: Hill and Wang, 2011); and Al Gore and David Blood, "A Manifesto for Sustainable Capitalism," *The Wall Street Journal* (December 14, 2011).
3. An important caveat arises when we define benefits in terms of willingness to pay. As we saw in Chapter 3, willingness to pay depends in part on ability to pay and therefore on income. A more subtle and complete analysis of social welfare would take this into account, perhaps by attaching more weight to the benefits of low-income consumers in order to reflect their lesser ability to pay.

4. More precisely, the cost referred to is total *variable* costs, or the costs of inputs that depend on how much is produced. For example, to produce more steel, a steel factory must consume more coal to run its furnaces, buy more iron ore to make steel, and hire more workers (or pay existing workers for overtime). Those inputs—fuel, materials, and labor—are variable inputs, because the amount needed varies with the amount of output. In contrast, the amount of physical capital—the machines and buildings that make up a factory—is typically *fixed*, at least in the short run. Thus capital costs are generally fixed costs. Because they do not increase with output, they do not affect a firm's decision of how much to produce—only whether or not to produce.
5. The fact that market outcomes are (in general) Pareto efficient does not mean that markets must necessarily always lead to desirable outcomes, if we (as a society) value equity as well as efficiency. Rather, markets always maximize net social welfare, given the distribution of wealth and resources among the members of society. An important and complementary result in economics, called the Second Theorem of Welfare Economics, states that any particular efficient outcome—for example, a perfectly equitable one—can be reached by a market economy from some initial distribution of resources. More practically, equity concerns can be addressed by redistributive policies such as progressive taxation. Although taxes typically diminish economic efficiency, most economists endorse them in some form, in order to ensure some level of distributional equity. Of course, how far society should go in trading efficiency for equity is not something economists agree upon; indeed, it is a question of social policy that is properly left up to the realm of deliberative politics. (See also note 29 in Chapter 3.)
6. Regarding Microsoft, see Richard J. Gilbert and Michael L. Katz, “An Economist's Guide to U.S. v. Microsoft,” *Journal of Economic Perspectives* 15(2): 25–44 (Spring 2001). Regarding diamonds and eyeglasses, respectively, see Eric Goldschein, “The Incredible Story of How De Beers Created and Lost the Most Powerful Monopoly Ever,” *Business Insider* (December 19, 2011), and Ana Swanson, “Meet the Four-Eyed, Eight-Tentacled Monopoly That Is Making Your Glasses So Expensive,” *Forbes* (September 10, 2014).

Chapter 5: Market Failures in the Environmental Realm

1. Garrett Hardin, “The Tragedy of the Commons,” *Science* 162 (3859): 1243–1248 (December 13, 1968).
2. This definition follows that given by William Baumol and Wallace Oates in their classic treatise *The Theory of Environmental Policy*, 2nd ed. (New York: Cambridge University Press, 1988).
3. In a private good setting, each individual consumes a different amount of the good, but they all pay the same unit price; therefore we sum up *quantities consumed at a given price* (i.e., a given marginal willingness to pay) rather than summing up willingness to pay at a given quantity (as we must for public goods). In terms of figure 5.4, we can frame the distinction between public and private goods as follows. For public goods, the social marginal benefit function represents the sum of individuals' willingness to pay at each given quantity of the public good. This implies that we sum the individual *MB* curves vertically. In contrast, for private goods, aggregate demand (loosely equivalent to social marginal benefits) is derived

by summing the individual demand curves *horizontally*, because the price is the same across individuals while the quantity varies.

4. Although Hardin's article is often cited for its memorable metaphor of collective action problems in the environmental arena, the thrust of his article was actually a dire (and, in retrospect, vastly overstated) warning about overpopulation. Thus it was more in line with the "Limits to Growth" style arguments we will encounter in Chapter 11.
5. We choose 550 ppm simply for the purposes of illustration. A number of scientists, environmental organizations, and government officials have advocated stabilizing atmospheric concentrations at 450 ppm, the level associated with a roughly 50 percent chance of keeping *long-term* global temperatures below 2°C above pre-industrial levels. Among many simplifications, this example considers only costs and benefits for the year 2100, because that is what IPCC estimates. Importantly, atmospheric concentrations would continue to rise after that year, resulting in higher long-term levels; reducing emissions would incur costs and provide benefits long before and after that year. We are also eliding the distinction between the atmospheric concentration of carbon dioxide (the chief greenhouse gas) and that of all greenhouse gases taken together (which requires expressing the impact of different greenhouse gases in terms of carbon dioxide equivalence). For reference, atmospheric concentrations of carbon dioxide rose above 400 ppm in 2014; atmospheric concentrations of greenhouse gases as a whole, measured in parts per million carbon dioxide equivalent, are somewhat higher. Cost and temperature estimates are taken from IPCC, *Climate Change 2014: Impacts, Adaptation and Vulnerability* (Cambridge, UK: Cambridge University Press, 2014); damages are based on the estimated damage function used in William Nordhaus, *A Question of Balance: Weighing the Options on Global Warming Policies* (New Haven, CT: Yale University Press, 2008).
6. We have presented an admittedly bleak case. In the real world, social norms or sanctions may arise to encourage contribution. Moreover, repeated play can change the outcome dramatically: If A and B face the situation over and over again, they may find cooperation much easier. Nonetheless, the simple model presented here provides a useful framework for thinking about collective action problems, the tragedy of the commons, and indeed public goods generally.

Chapter 6: Managing Stocks: Natural Resources as Capital Assets

1. For more information on the breakeven prices of these oil resources in the mid-2010s, see Jeff Lewis, "Is Oil Sands Development Still Worth It?" *The Globe and Mail* (October 28, 2014); Jim Burkhard, "Tight Oil Test: US Output at Lower Prices," *IHS Unconventional Energy Blog* (November 20, 2014); and Brian Dumaine, "Why America's Fracking Revolution Won't Be Hurt (Much) by Low Oil Prices," *Fortune* (December 2, 2014).
2. See Tom H. Tietenberg, *Environmental and Natural Resource Economics*, 6th ed. (Boston: Addison-Wesley, 2003), 4. Britain is still mining coal. Jevons also worried about the world running out of paper. Anticipating an eventual shortage, Jevons collected a stash of blank writing paper so large that his children were still using it 50 years after his death. This famous anecdote is related by John Maynard Keynes. See *The Collected Writings of John Maynard Keynes: Volume X, Essays in Biography* (London: Macmillan, 1972), 117.

3. For a discussion of the many factors contributing to the profitable extraction of gas and oil from shale formations, see Zhongmin Wang and Alan Krupnick, "A Retrospective Review of Shale Gas Development in the United States," RFF Discussion Paper 13-12, Washington, DC: Resources for the Future (April 2013). For a comprehensive review of the benefits and costs of shale gas development, see Charles F. Mason, Lucija A. Muehlenbachs, and Sheila M. Olmstead, "The Economics of Shale Gas Development," *Annual Review of Resource Economics* 7. doi:10.1146/annurev-resource-100814-125023 (2015).
4. The two-period nonrenewable natural resources problem we solve here is originally due to James M. Griffin and Henry B. Steele, *Energy Economics and Policy* (New York: Academic Press, 1980) and is modeled on the interpretation in Teitenberg (2003), cited in note 3, above, pp. 89–93.
5. Notice that this problem is constructed so as to exhaust the stock of oil entirely. In reality, it may not be efficient to completely exhaust many nonrenewable resource stocks because the cost of extracting these resources can become very high as stocks dwindle. (Recall that we have assumed that marginal extraction costs were constant in the two-period problem.)
6. Among other conditions, the Hotelling Rule relies on the assumptions that marginal extraction costs and ore grades are constant. If these assumptions are not met, the Hotelling model is somewhat more complicated, although the general no-arbitrage feel of the model is preserved.
7. See James L. Smith, "World Oil: Market or Mayhem?" *Journal of Economic Perspectives* 23(3): 145–164 (2009).
8. China's difficulty in exercising market power for rare earths was also compromised by the incentive to "cheat" on internal production quotas, as domestic smuggling was also spurred by high global prices (similar to OPEC's problems sustaining cooperation between member countries). Background on rare earths and an analysis of the 2010 Chinese embargo is provided in Eugene Gholz, "Rare Earth Elements and National Security," New York: Council on Foreign Relations Energy Report (October 2014).
9. See Allen L. Torell, James D. Libbin, and Michael D. Miller, "The Market Value of Water in the Ogallala Aquifer," *Land Economics* 66(2): 163–175 (1990).

Chapter 7: Stocks That Grow: The Economics of Renewable Resource Management

1. The following two-period problem is adapted from Barry C. Field, *Natural Resource Economics: An Introduction* (Long Grove, IL: Waveland Press, Inc., 2001).
2. Figure 7.4 is adapted from Field (2001), 232.
3. G. Cornelius van Kooten, Clark S. Binkley, and Gregg Delcourt, "Effect of Carbon Taxes and Subsidies on Optimal Forest Rotation Age and Supply of Carbon Services," *American Journal of Agricultural Economics* 77(2): 365–374 (1995).
4. See John Creedy and Anke D. Wurzbacher, "The Economic Value of a Forested Catchment with Timber, Water and Carbon Sequestration Benefits," *Ecological Economics* 38(1): 71–83 (2001).
5. See Jonathan Rubin, Gloria Helfand, and John Loomis, "A Benefit–Cost Analysis of the Northern Spotted Owl," *Journal of Forestry* 89: 25–30 (December 1991).
6. See D. Hagen, J. Vincent, and D. Welle, "Benefits of Preserving Old-Growth Forests and the Spotted Owl," *Contemporary Policy Issues* 10: 13–26 (April 1992).

7. See U.S. Forest Service, Northern Research Station, "Who Owns America's Forests? Forest Ownership Patterns and Family Forest Highlights from the National Woodland Owner Survey," NRS-INF-06-08, Washington, DC: U.S. Department of Agriculture (May 2008).
8. See Food and Agriculture Organization of the United Nations, "Understanding Forest Tenure in Africa: Opportunities and Challenges for Forest Tenure Diversification," Forestry Policy and Institutions Working Paper no. 19, Rome, Italy: FAO (2008); and Food and Agriculture Organization of the United Nations, "Understanding Forest Tenure in South and Southeast Asia," Forestry Policy and Institutions Working Paper no. 14, Rome, Italy: FAO (2006).
9. See Charles Wood and Robert Walker, "Saving the Trees by Helping the Poor: A Look at Small Producers along Brazil's Transamazon Highway," *Resources* 136(Summer): 14–17 (1999).
10. See Robert T. Deacon, "Deforestation and the Rule of Law in a Cross-Section of Countries," *Land Economics* 70: 414–430 (1994).
11. See Robert L. Mendelsohn, "Property Rights and Tropical Deforestation," *Oxford Economic Papers* 46: 750–756 (1994).
12. See Wood and Walker, "Saving the Trees by Helping the Poor."
13. See Zachary Liscow, "Do Property Rights Promote Investment but Cause Deforestation? Quasi-Experimental Evidence from Nicaragua," *Journal of Environmental Economics and Management* 65(2): 241–261 (2013).
14. Millennium Ecosystem Assessment, *Ecosystems and Human Well-Being: Policy Responses* (Washington, DC: Island Press, 2005); Kwaw S. Andam, Paul J. Ferraro, Alexander Pfaff, G. Arturo Sanchez-Azofeifa, and Juan A. Robalino, "Measuring the Effectiveness of Protected Area Networks in Reducing Deforestation," *Proceedings of the National Academy of Sciences* 105(42): 16089–16094 (2008).
15. See Katharine R. E. Sims, Jennifer M. Alix-Garcia, Elizabeth Shapiro-Garza, Leah R. Fine, Volker C. Radeloff, Glen Aronson, Selene Castillo, Carlos Ramirez-Reyes, and Patricia Yañez-Pagans, "Improving Environmental and Social Targeting through Adaptive Management in Mexico's Payments for Hydrological Services Program," *Conservation Biology* 28(5): 1151–1159; and C. Muñoz-Piña, A. Guevara, J. Torres, and J. Braña, "Paying for the Hydrological Services of Mexico's Forests: Analysis, Negotiations and Results," *Ecological Economics* 65(4): 725–736.
16. Jennifer M. Alix-Garcia, Katharine R. E. Sims, and Patricia Yañez-Pagans, "Only One Tree from Each Seed? Environmental Effectiveness and Poverty Alleviation in Mexico's Payments for Ecosystem Services Program," *American Economic Journal: Economic Policy*, forthcoming 2015.
17. Empirical evidence suggests that spillovers from Costa Rica's tropical forest protected areas have been minimal (see Andam et al., "Measuring the Effectiveness of Protected Area Networks in Reducing Deforestation").
18. Kwaw S. Andam, Paul J. Ferraro, Katharine R. E. Sims, Andrew Healy, and Margaret B. Holland, "Protected Areas Reduced Poverty in Costa Rica and Thailand," *Proceedings of the National Academy of Sciences* 107(22): 9996–10001 (2010).
19. The model is due originally to M. B. Schaefer, "Some Considerations of Population Dynamics and Economics in Relation to the Management of Marine Fisheries," *Journal of the Fisheries Research Board of Canada* 14: 669–681 (1957).

20. The steady-state analysis, though a simplification, provides a good basis from which to develop the intuition and general conclusions of the bioeconomic fishing model. Nonetheless, some important aspects of the problem will be lost in our departure from the dynamic context. In particular, the simplification erases the link between today's harvest and tomorrow's stock, as well as the time value of money. We will be more explicit about the costs of the steady-state assumption later in the chapter.
21. This is simply an assumption that the yield–effort function is a multiplicative function of the stock and the level of effort. A common yield–effort function is the constant returns per unit effort function: $Y = qXE$, where Y is yield, X is the fish stock, E is the level of fishing effort, and $q > 0$ is a catchability coefficient.
22. At the efficient level of fishing effort, the slopes of the total benefit and total cost curves are equal, so marginal benefit equals marginal cost. If this is unclear, see our discussion of the relationship between total and marginal cost and benefit curves in chapter 2.
23. For a fascinating discussion of common–property resources, based on a vast body of research on community–level natural resource management, see Elinor Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action* (New York: Cambridge University Press, 1990).
24. For an exposition of the dynamic fishing model, see Chapter 3 in Jon M. Conrad, *Resource Economics*, 2nd ed. (New York: Cambridge University Press, 2010).
25. The following is summarized from F. Berkes, D. Feeny, B. J. McCay, and J. M. Acheson, “The Benefits of the Commons,” *Nature* 340(6229): 91–93 (1989).
26. See Food and Agriculture Organization of the United Nations, “The State of World Fisheries and Aquaculture: Opportunities and Challenges,” Rome: FAO, 2014 (p. 25), available at <http://www.fao.org/3/d1eaa9a1-5a71-4e42-86c0-f2111f07de16/i3720e.pdf>, accessed January 6, 2015.
27. For a wonderful description of the initial abundance of these stocks and their subsequent decimation, see Mark Kurlansky, *Cod: A Biography of the Fish That Changed the World* (New York: Penguin Books, 1997).
28. See Suzanne Iudicello, Michael Weber, and Robert Wieland, *Fish, Markets and Fishermen: The Economics of Overfishing* (Washington, DC: Island Press, 1999); and Robert Buchsbaum, Judith Pederson, and William E. Robinson, eds., *The Decline of Fisheries Resources in New England: Evaluating the Impact of Overfishing, Contamination, and Habitat Degradation* (Cambridge, MA: MIT Sea Grant College Program, 2005).
29. Annual losses in net benefits in the Bering Sea fishery at the time were estimated at \$124 million; see Daniel Huppert, “Managing the Groundfish Fisheries of Alaska: History and Prospects,” *Reviews in Aquatic Sciences* 4(4): 339–373 (1991).
30. John M. Ward and Jon G. Sutinen, “Vessel Entry–Exit Behavior in the Gulf of Mexico Shrimp Fishery,” *American Journal of Agricultural Economics* 76: 916–923 (1994).
31. See Food and Agriculture Organization of the United Nations, “Marine Fisheries and the Law of the Sea: A Decade of Change,” special chapter (rev.) in *The State of Food and Agriculture 1992*, FAO Fisheries Circular no. 853, Rome: FAO, 1993.
32. See Ussif Rashid Sumaila and Daniel Pauly, eds., “Catching More Bait: A Bottom–Up Re–estimation of Global Fisheries Subsidies,” University of British Columbia

- Fisheries Centre Research Report Vol. 14, no. 6, Vancouver, BC, 2006, available at http://www.fisheries.ubc.ca/webfm_send/126, accessed January 6, 2015.
33. See National Research Council, *Dolphins and the Tuna Industry* (Washington, DC: National Academy Press, 1992).
34. See Mario F. Teisl, Brian Roe, and Robert L. Hicks, "Can Eco-Labels Tune a Market? Evidence from Dolphin-Safe Labeling," *Journal of Environmental Economics and Management* 43: 339–359 (2000).

Chapter 8: Principles of Market-Based Environmental Policy

1. Ronald Coase, "The Problem of Social Cost," *Journal of Law and Economics* 3: 1–44 (1969).
2. Coase, of course, realized the importance of transaction costs in the real world. Indeed, much of his (very long) article is devoted not to the so-called Coase Theorem but rather to the proposition that even when transaction costs are sizeable—so that assigning liability does matter—government action is not necessarily preferable to *laissez faire*. For example, Coase argued that taxing a polluting factory created its own problems: The resulting reduction in smoke would attract more firms and residents to the area, raising the subsequent damages from pollution and lowering the factory's productivity. Note how narrow this argument is: It implicitly treats a single factory (rather than the industry) as the object of government policy and relies on the quaint assumption that smoke from a factory affects only a well-defined region in its vicinity. More importantly, the argument misunderstands the nature of the tax remedy. The proper tax is equal to marginal damages *at the efficient point* rather than (say) to marginal damages at the unfettered level of production or under some other arbitrary conditions.
3. Many other such cases are described in Terry Anderson and Gary Libecap, *Environmental Markets: A Property Rights Approach* (New York: Cambridge University Press, 2014), which makes a robust argument in favor of Coasean solutions to environmental problems.
4. This discussion is based on Danièle Perrot-Maitre, and Patsy Davis, "Case Studies of Markets and Innovative Financial Mechanisms for Water Services from Forests," mimeo produced by ForestTrends and the Katoomba Group, May 2001.
5. To be more precise, a tax and subsidy create the same incentives for pollution control *on the margin*. Of course firms would prefer to receive a subsidy rather than pay a tax. Indeed, that points to a potential problem in the long run: A subsidy would create too much incentive for firms to enter the regulated market. In contrast, although a tax might encourage firms to exit the industry, this is not distortionary (at least if the tax is set according to marginal damages), because only the very dirtiest firms with the most expensive abatement options will choose to exit rather than to pay the tax, and those are the firms for which the damages from pollution are greater than the social benefits from production. In addition, a tax raises government revenue, whereas subsidies spend it—a difference that may have implications for efficiency depending on how tax revenue is used, as we discuss later in the chapter.
6. Two sources for further reading on information provision programs and certification are James T. Hamilton, *Regulation through Revelation: The Origin, Politics, and Impacts of the Toxics Release Inventory Program* (New York: Cambridge University

- Press, 2005), on the TRI program; and Benjamin Cashore, Graeme Auld, and Deanna Newsom, *Governing through Markets: Forest Certification and the Emergence of Non-State Authority* (New Haven, CT: Yale University Press, 2004), on forest certification.
7. See Hunt Allcott and Todd Rogers, “The Short-Run and Long-Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation,” *American Economic Review* 104(10): 3003–3037 (2014); and Hunt Allcott, “Social Norms and Energy Conservation” *Journal of Public Economics* 95(9–10): 1082–1095 (2011).
 8. See Paul J. Ferraro and Michael Price, “Using Non-Pecuniary Strategies to Influence Behavior: Evidence from a Large-Scale Field Experiment,” *Review of Economics and Statistics* 95(1): 64–73 (2013).
 9. Figures for 2013 are from International Energy Agency, *World Energy Outlook 2014* (Paris: International Energy Agency, 2014); abatement estimates from International Energy Agency, *Redrawing the Energy–Climate Map* (Paris: International Energy Agency, 2013).
 10. This discussion is based on research by Ian W. H. Parry and Kenneth A. Small, “Does Britain or the United States Have the Right Gasoline Tax?” *American Economic Review* 95(4): 1276–1289 (2005).
 11. You may recall from our discussion in Chapter 5 that some provision of public goods will typically arise even in the absence of government intervention. In our two-person example of the flower garden, the neighbor with the higher value for the garden ends up providing a positive amount of the good. But remember that in real-world cases of interest, where large populations share the public good (e.g., clean air), free riding will dominate, and private provision of the public good will be much smaller than the efficient level.
 12. For more on the EU-ETS, see the EU case study by Environmental Defense Fund and the International Emission Trading Association, part of “The World’s Carbon Markets: A Case Study Guide to Emission Trading,” available at <http://www.ieta.org/worldscarbonmarkets>, as well as information and fact sheets on the website of the European Commission Directorate-General for Climate Action, http://ec.europa.eu/clima/policies/ets/index_en.htm. Information on the Chicago Climate Exchange is taken from fact sheets and datasets at <https://www.theice.com/ccx>.
 13. Of course, privatizing the resource is not the only solution. As we saw in Chapter 7, with the example of beaver hunting in James Bay, local communities can manage natural resources sustainably through social norms and customs of stewardship rather than private property. See also the work of Elinor Ostrom, cited in note 23 of Chapter 7.
 14. See Scott C. Matulich and Michael L. Clark, “North Pacific Halibut and Sablefish IFQ Policy Design: Quantifying the Impacts on Processors,” *Marine Resource Economics* 18: 149–166 (2003); and Clarence G. Pautzke and Chris W. Oliver, “Development of the Individual Fishing Quota Program for Sablefish and Halibut Longline Fisheries off Alaska,” Presentation to the National Research Council’s Committee to Review Individual Fishing Quotas, Anchorage, Alaska, September 4, 1997. Available at <https://alaskafisheries.noaa.gov/ram/ifq/ifqpaper.htm>, accessed February 16, 2015.

15. See Christopher Costello, Stephen D. Gaines, and John Lynham, "Can Catch Shares Prevent Fisheries Collapse?" *Science* 321: 1678–1681 (2008); and Geoffrey Heal and Wolfram Schlenker, "Economics: Sustainable Fisheries," *Nature* 455: 1044–1045 (2008).
16. Unfortunately, this is not necessarily the free lunch it might seem to be. It turns out that environmental regulations may exacerbate the preexisting distortions caused by those other taxes. For example, a carbon tax (or a cap-and-trade program) would raise the price of gasoline. This effectively lowers real wages by reducing purchasing power. The effect is much the same as an increase in the income tax and therefore contributes to distortions in the labor market. Whether or not the efficiency gain from reducing other taxes offsets the hidden cost of the regulation (what economists call the tax interaction effect) is an empirical question—that is, it can go one way or the other depending on the circumstances.
17. You may notice that the true marginal cost curves are simply parallel shifts of the expected (average) marginal cost curve. This might seem at first to be a very special case. As Weitzman (1974) showed rigorously, however, such parallel shifts are a fairly general way of modeling uncertainty, for reasons beyond the scope of this discussion.

Chapter 9: The Case for Market-Based Instruments in the Real World

1. Although minimizing costs is not the same thing as maximizing net benefits, the two are closely related. If the benefits of pollution control depend only on the total amount of abatement (in other words, if the damages from pollution depend only on the total amount of pollution), and two policies achieve the same objective, the one with lower costs must be more efficient than the other (even if the goal itself is not fully efficient). Cost minimization is a necessary (though not sufficient) condition for efficiency. In other words, it is necessary to minimize costs in order to maximize benefits minus costs.
2. Two other things are worth noting about the equal marginal abatement cost condition for cost-effectiveness. First, it is necessary but not sufficient. It is possible to imagine scenarios in which firms' marginal costs are equated but total costs are not minimized—for example, if the capital costs of abatement were very high. Second, we have been careful to state that the equal-marginal-costs principle applies only to all firms that abate pollution. It is possible that some firms with very high marginal abatement costs ought not to control pollution at all (at least from a cost minimization standpoint).
3. Recall that a supply curve in a product market makes sense only in a competitive industry, because it embodies the assumptions that individual producers are price takers and that they seek to maximize profit. The same is true with an aggregate marginal abatement cost curve: It makes sense only in a context in which all regulated firms take the price of abatement (e.g., the tax) as given and seek to minimize compliance costs. Indeed, the aggregate curve essentially has cost-effectiveness built in: It can be thought of as tracing out the cost of abating pollution on the margin, when abatement is allocated cost-effectively between firms. In other words, the cost being measured corresponds to the *least-cost allocation of abatement* between firms. For this reason, it doesn't make much sense to think about the

- aggregate marginal abatement cost curve under uniform standards, because (as we have already seen) uniform standards are not cost-effective in general.
4. In fact, we have seen this result already, in the Coase Theorem, discussed in Chapter 8. Saying that the equilibrium allocation of tradable allowances is unaffected by the initial allocation, as long as transactions costs are small, is just a restatement of Coase's insight that private bargaining in the absence of transaction costs will lead to efficient outcomes regardless of the assignment of property rights.
 5. The cost of adoption—the capital cost of installing or building new control equipment—will remain implicit in our analysis. This cost is what the firm must balance against the potential cost savings from using the new technology to decide whether to adopt it. Note that we can safely assume that the cost of installing a particular technology does not depend on whether a firm faces a tax or a standard. Therefore, it will not affect our comparison between various policy instruments.
 6. We have glossed over two subtle points. First, we have implicitly treated more technology as a good thing. In fact, from the perspective of economics, the question we should ask is not “Which policy leads to more investment in new technologies?” but rather “Which policy leads to more *efficient* investment in new technologies?” Answering this question turns out to be harder than it seems (which helps explain why we tackled the easier question in the text). Nonetheless, there are strong reasons to believe that market-based instruments will be superior to performance standards on efficiency grounds in the face of technological change. The real question is which market-based instrument—an emission tax or a cap-and-trade policy—is preferable. The answer depends in large part on the relative slopes of marginal cost and benefit, as in the case of uncertainty over marginal cost discussed in Chapter 8. Indeed, you might be able to see the intuitive connection between “cost uncertainty” and technological change. Second, we have assumed that the regulator does not anticipate the adoption of the new technology. You can think of this as corresponding to a scenario in which the form and stringency of regulation remain set in stone over a reasonably long duration—long enough for new technologies to arise.
 7. For the story of the longitude prize, see Dava Sobel, *Longitude* (New York: Walker & Company, 2005).
 8. See Richard G. Newell and Nathan E. Wilson, “Technology Prizes for Climate Change Mitigation,” Resources for the Future Discussion Paper 05-33 (2005).
 9. Two nuances are worth noting. First, because we are assuming a fixed market price for fish and have equated net benefits with net revenues, maximizing net benefits is equivalent to minimizing cost. We make the distinction here mainly for the purposes of intuition. Second, it might appear at first that by introducing quotas with a market price, an IFQ policy would reduce the net benefits to fishers. In fact, however, that is true for a particular fisher only to the extent that she must buy more IFQs than she sells. Moreover, from the perspective of society as a whole, the value of IFQ transfers is a wash; what matters is the difference between the total value of the harvest (reflected in its market price) and the real costs of catching it.
 10. Nicholas Z. Muller and Robert Mendelsohn, “Efficient Pollution Regulation: Getting the Prices Right,” *American Economic Review* 99(5): 1714–1739.
 11. Nonuniform mixing and regional transport of air pollution have garnered significant attention since 2008, when the U.S. Federal Court of Appeals ruled that the

- Clean Air Act's emission trading programs did not adequately address this problem, particularly for the transport of air pollution from midwestern power plants to the East Coast. These issues will be discussed in greater detail in Chapter 10, in the section that describes the U.S. sulfur dioxide market.
12. Meredith Fowlie and Nicholas Muller, "Market-Based Emissions Regulation When Damages Vary across Sources: What Are the Gains from Differentiation?" NBER Working Paper 18801 (Cambridge, MA: National Bureau of Economic Research, 2013).
 13. The following example is from R. Scott Farrow, Martin T. Schultz, Pinar Celikkol, and George L. Van Houtven, "Pollution Trading in Water Quality Limited Areas: Use of Benefits Assessment and Cost-Effective Trading Ratios," *Land Economics* 81(2): 191–205 (2005).
 14. The Minnesota River Basin Trading program and the Long Island Sound Nitrogen Credit Exchange are good examples. See Karen Fisher-Vanden and Sheila Olmstead, "Moving Pollution Trading from Air to Water: Potential, Problems and Prognosis," *Journal of Economic Perspectives* 27(1): 147–172 (2013).

Chapter 10: Market-Based Instruments in Practice

1. If you are interested in reading more, some comprehensive studies of market-based policies in the real world include Charles E. Kolstad and Jody Freeman, eds., *Moving to Markets in Environmental Regulation: Lessons from Twenty Years of Experience* (New York: Oxford University Press, 2007); Robert N. Stavins, "Experience with Market-Based Environmental Policy Instruments," in Karl-Göran Mäler and Jeffrey R. Vincent, eds., *Handbook of Environmental Economics*, Vol. 1 (Amsterdam: Elsevier Science B.V., 2003), 355–435; Thomas Sterner, *Policy Instruments for Environmental and Natural Resource Management* (Washington, DC: Resources for the Future, 2003), 363; and Theodore Panayotou, *Instruments of Change: Motivating and Financing Sustainable Development* (London: Earthscan Publications, Ltd., 1998).
2. Congress anticipated this problem in the 1977 Clean Air Act Amendments, which declared that significant capital investments at existing plants, designed to extend their lifetimes, would trigger the stringent federal regulation that applied to new sources. The so-called New Source Review process entailed a costly survey of planned investments that might trigger those regulatory standards. However, enforcing this provision has bedeviled subsequent administrations and stirred up fierce opposition from electric power plants. The line between routine maintenance and "significant" upgrades has been difficult to establish. As a result, electric utilities have been reluctant to make investments that would improve the operating efficiency of their plants. At the same time, the EPA has been largely unsuccessful in forcing utilities that made past investments to install new control equipment.
3. Many of the details about the allowance trading program created by the 1990 Clean Air Act are taken from the definitive overview provided by A. Denny Ellerman, Paul J. Joskow, Richard Schmalensee, Juan-Pablo Montero, and Elizabeth M. Bailey, *Markets for Clean Air: The U.S. Acid Rain Program* (New York: Cambridge University Press, 2000).
4. See Richard Schmalensee and Robert N. Stavins, "The SO₂ Allowance Trading System: The Ironic History of a Grand Policy Experiment," *Journal of Economic Perspectives* 27: 103–122 (2013).

5. See Dallas Burtraw, Alan Krupnick, Erin Mansur, David Austin, and Deirdre Farrell, "Costs and Benefits of Reducing Air Pollutants Related to Acid Rain," *Contemporary Economic Policy* 16(4): 379–400 (October 1998).
6. See Nathaniel O. Keohane, "Cost Savings from Allowance Trading in the 1990 Clean Air Act," in Charles E. Kolstad and Jody Freeman, eds., *Moving to Markets in Environmental Regulation: Lessons from Twenty Years of Experience* (New York: Oxford University Press, 2007), 194–229.
7. See Nathaniel O. Keohane, "Environmental Policy and the Choice of Abatement Technique: Evidence from Coal-Fired Power Plants," Yale University mimeo available at www.som.yale.edu/faculty/nok4/files/papers/scrubbers.pdf (on adoption decisions); and David Popp, "Pollution Control Innovations and the Clean Air Act of 1990," *Journal of Policy Analysis and Management* 22 (4): 641–660 (Fall 2003) (on patents and technological innovation).
8. In fact, a small fraction—roughly 3 percent—of each year's allowances were auctioned off by the EPA. This provision was included in the trading program to overcome concerns that giving away all of the allowances would make it possible for the incumbent firms to hoard their allowances, driving up the price for firms that wanted to enter. But in a sign of the political power of the regulated industry, the allowance revenue was credited to the electric utilities that had given up their allowances to be auctioned rather than being retained by the government.
9. See Nicholas Muller and Robert Mendelsohn, "Efficient Pollution Regulation: Getting the Prices Right," *American Economic Review* 99(5): 1714–1739 (2009).
10. The geographic distribution of pollution under the allowance trading program is also discussed by Ellerman and his colleagues in Chapter 5 of *Markets for Clean Air* (see note 3).
11. Recall from our discussion of fisheries in Chapter 7 that fishing to MSY is economic but not biological overfishing; thus the TACs would have to be lower than their current levels to induce the efficient level of fishing effort.
12. At the program's inception, IFQs represented a right to an absolute amount of fish rather than a percentage of the TAC. But this policy led to total shares that exceeded the TAC in several fisheries soon after the program was implemented. The government had to buy back almost 16 metric tons of shares, at significant cost. Issuing quotas as a percentage of the TAC avoids causing this problem in the future.
13. Cindy Chu, "Thirty Years Later: The Global Growth of ITQs and Their Influence on Stock Status in Marine Fisheries," *Fish and Fisheries* 10: 217–230 (2009).
14. For a recent overview of carbon pricing systems around the world, see World Bank, *State and Trends of Carbon Pricing 2014* (Washington, DC: The World Bank, 2014). Detailed country-level case studies of emission trading systems, prepared jointly by the International Emission Trading Association and the Environmental Defense Fund, can be found online at <http://www.ieta.org/worldscarbonmarkets>.
15. See John H. Annala, "New Zealand's ITQ System: Have the First Eight Years Been a Success or a Failure?" *Reviews in Fish Biology and Fisheries* 6: 43–62 (1996).
16. The excluded thirty stocks covered areas fished lightly for only a few species, not suspected to be at serious risk of overfishing. See John H. Annala, *Report from the Fishery Assessment Plenary, May 1994: Stock Assessments and Yield Estimates* (Wellington, NZ: MAF Fisheries Greta Point, 1994), quoted in Annala, "New Zealand's ITQ System," 47.

17. All data on market performance, including quota and lease prices and market activity, are taken from Richard G. Newell, James N. Sanchirico, and Suzi Kerr, "Fishing Quota Markets," *Journal of Environmental Economics and Management* 49(3): 437–462 (2005).
18. Note that this equity concern is distinct from a potential efficiency concern. If some market participants grow large enough that they can affect the market price of quotas, then efficiency will suffer, just as in any market where firms have market power. Recall our discussion of market failures at the end of Chapter 5.
19. See James Sanchirico and Richard Newell, "Catching Market Efficiencies," *Resources* 150: 8–11 (2003).
20. See Clement and Associates, *New Zealand Commercial Fisheries: The Guide to the Quota Management System* (Tauranga, NZ: Clement & Associates, 1997); and P. Major, "Individual Transferable Quotas and Quota Management Systems: A Perspective from the New Zealand Experience," in K. L. Gimbel, ed., *Limiting Access to Marine Fisheries: Keeping the Focus on Conservation* (Washington, DC: Center for Marine Conservation and World Wildlife Fund, 1994), 98–106.
21. These fees covered more than 80 percent of the program's cost in 1994–1995. See Suzanne Iudicello, Michael Weber, and Robert Wieland, *Fish, Markets and Fisheries: The Economics of Overfishing* (Washington, DC: Island Press, 1999).
22. See Iudicello et al. *Fish, Markets and Fishermen*, 105.
23. See Christopher M. Dewees, "Effects of Individual Quota Systems on New Zealand and British Columbia Fisheries," *Ecological Applications* 8(1): S133–S138 (1998).
24. See R. W. Mayer, W. B. DeOreo, E. M. Opitz, J. C. Kiefer, W. Y. Davis, B. Dziegielewski, and J. O. Nelson, *Residential End Uses of Water* (Denver, CO: American Water Works Association Research Foundation, 1998).
25. See Mary E. Renwick and Richard D. Green, "Do Residential Water Demand Side Management Policies Measure Up? An Analysis of Eight California Water Agencies," *Journal of Environmental Economics* 40(1): 37–55 (2000).
26. See Ellen M. Pint, "Household Responses to Increased Water Rates," *Land Economics* 75(2): 246–266 (1999).
27. More precisely, price elasticity measures the percentage change in demand that results from a 1 percent increase in price.
28. See Christopher Timmins, "Demand-Side Technology Standards under Inefficient Pricing Regimes: Are They Effective Water Conservation Tools in the Long Run?" *Environmental and Resource Economics* 26: 107–124 (2003).
29. See Erin T. Mansur and Sheila M. Olmstead, "The Value of Scarce Water: Measuring the Inefficiency of Municipal Regulations," *Journal of Urban Economics* 71: 332–346 (2012).
30. This section draws on A. Denny Ellerman, Frank J. Convery, and Christian de Perthuis, *Pricing Carbon: The European Union Emissions Trading Scheme* (Cambridge, UK: Cambridge University Press, 2010).
31. Raphael Calel and Antoine Dechezleprêtre, "Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market," Grantham Research Institute on Climate Change and the Environment Working Paper No. 75, London School of Economics (revised February 2013).

32. The information in this section is summarized from Karen A. Fisher-Vanden and Sheila M. Olmstead, "Moving Pollution Trading from Air to Water: Potential, Problems and Prognosis," *Journal of Economic Perspectives* 27(1): 147–172 (2013).
33. See George Van Houtven, Ross Loomis, Justin Baker, Robert Beach, and Sara Casey, "Nutrient Credit Trading for the Chesapeake Bay: An Economic Study," Report prepared for the Chesapeake Bay Commission, May 2012, RTI International, Research Triangle Park, NC.
34. See Robert Repetto, Roger C. Dower, Robin Jenkins, and Jacqueline Geoghegan, *Green Fees: How a Tax Shift Can Work for the Environment and the Economy* (Washington, DC: World Resources Institute, 1992).
35. Skumatz Economic Research Associates, Inc., "Pay as You Throw (PAYT) in the U.S.: 2006 Update and Analyses," Final report to the EPA Office of Solid Waste, December 30, 2006, Superior, CO, available at www.epa.gov/osw/conserve/tools/payt/pdf/sera06.pdf, accessed February 21, 2015.
36. See Don Fullerton and Thomas C. Kinnaman, "Household Responses to Pricing Garbage by the Bag," *American Economic Review* 86(4): 971–984 (1996).
37. See Robin R. Jenkins, Salvador A. Martinez, Karen Palmer, and Michael J. Podolsky, "The Determinants of Household Recycling: A Material-Specific Analysis of Recycling Program Features and Unit Pricing," *Journal of Environmental Economics and Management* 45: 294–318 (2003); and Ju-Chin Huang, John M. Halstead, and Shanna B. Saunders, "Managing Municipal Solid Waste with Unit-Based Pricing: Policy Effects and Responsiveness to Pricing," *Land Economics* 87(4): 645–660 (2011).
38. The Brazilian example is summarized from Kenneth M. Chomitz, "Transferable Development Rights and Forest Protection: An Exploratory Analysis," *International Regional Science Review* 27(3): 348–373 (2004).
39. See Virginia McConnell and Margaret Walls, "U.S. Experience with Transferable Development Rights," *Review of Environmental Economics and Policy* 3(2): 288–303 (2009); and Virginia McConnell, Margaret Walls, and Elizabeth Kopits, "Zoning, Transferable Development Rights, and the Density of Development," *Journal of Urban Economics* 59: 440–457 (2009).
40. Figures for 1993–2000 from National Research Council, *Compensating for Wetland Losses under the Clean Water Act* (Washington, DC: National Academy Press, 2001). Figures on wetland mitigation banks in 2001 and 2005 taken from Jessica Wilkinson and Jared Thompson, *2005 Status Report on Compensatory Mitigation in the United States* (Washington, DC: Environmental Law Institute, 2006). Credit prices from Richard R. Rogoski, "Liquid Assets: New Breed of Bankers Deal with Wetlands," *Triangle Business Journal*, September 8, 2006.
41. The discussion of conservation banking draws on a report titled "Mitigation Banking as an Endangered Species Conservation Tool" (Washington, DC: Environmental Defense Fund, 1999), downloaded from www.environmentaldefense.org/documents/146_mb.PDF#search=%22ESA%20endangered%20species%20banking%22.
42. See Christian Langpap and Joe Kerkvliet, "Endangered Species Conservation on Private Land: Assessing the Effectiveness of Habitat Conservation Plans," *Journal of Environmental Economics and Management* 64: 1–15 (2012).

43. This phenomenon is known colloquially as “shoot, shovel, and shut up.” For example, among the owners of more than 400 forest plots in the U.S. Southeast, the average age of timber harvest (the average rotation, discussed in Chapter 7) falls from nearly 70 years if there are no endangered red-cockaded woodpeckers nearby, to nearly half that length if there are 25 colonies within 25 miles of the logging site. Because the woodpeckers prefer old-growth pine, it appears as if landowners seek to avoid the costs of federal regulation under the ESA by discouraging woodpeckers from settling in their trees. See Dean Lueck and Jeffrey Michael, “Preemptive Habitat Destruction Under the Endangered Species Act,” *Journal of Law and Economics* 46: 27–60 (2003). Broader studies of the impacts of species listing under the ESA have also suggested that listing can be harmful to species recovery, unless substantial funds are invested in recovery efforts. See Paul J. Ferraro, Craig McIntosh, and Monica Ospina, “The Effectiveness of the U.S. Endangered Species Act: An Econometric Analysis Using Matching Methods,” *Journal of Environmental Economics and Management* 54: 245–261 (2007).
44. See Kelly Chinnners Reiss, Erica Hernandez, and Mark T. Brown, *An Evaluation of the Effectiveness of Mitigation Banking in Florida: Ecological Success and Compliance with Permit Conditions*. Florida Department of Environmental Protection #WM881. Gainesville, FL (2007), available at http://www.dep.state.fl.us/water/wetlands/docs/mitigation/Final_Report.pdf, accessed February 22, 2015; and Rebecca L. Kihlslinger, “Success of Wetlands Mitigation Projects,” *National Wetlands Newsletter* 30(2): 14–16 (2008).

Chapter 11: Sustainability and Economic Growth

1. See Joseph Stiglitz, “The Ethical Economist,” a review of *The Moral Consequences of Economic Growth*, by Benjamin M. Friedman, *Foreign Affairs* (November/December): 128–134 (2005).
2. See Donella H. Meadows, D. L. Meadows, J. Randers, and W. W. Behrens, *The Limits to Growth* (New York: Universe Books, 1972); Donella H. Meadows et al., *Beyond the Limits* (Post Mills, VT: Chelsea Green Publishing Company, 1992).
3. See Julian Simon, *The Ultimate Resource* (Princeton, NJ: Princeton University Press, 1981); Julian Simon, *The Ultimate Resource II* (Princeton, NJ: Princeton University Press, 1996).
4. See William D. Nordhaus, “Lethal Model 2: The Limits to Growth Revisited,” *Brookings Papers on Economic Activity* (2): 1–59 (1992).
5. Economist Simon Kuznets established such a hill-shaped relationship between economic growth and income inequality (initially increasing, then decreasing) in the 1950s, which explains the origin of the name of this hypothesis.
6. For a summary of these issues, see Richard T. Carson, “The Environmental Kuznets Curve: Seeking Empirical Regularity and Theoretical Structure,” *Review of Environmental Economics and Policy* 4(1): 3–23 (2010). For evidence of pollution reductions in high-income countries, see John List and Mitch Kuncie, “Environmental Protection and Economic Growth: What Do the Residuals Tell Us?” *Land Economics* 76(2): 267–282 (2000); and Michael Greenstone, “Did the Clean Air Act Cause the Remarkable Decline in Sulfur Dioxide Concentrations?” *Journal of Environmental Economics and Management* 47(3): 585–611 (2004).

7. See, for example, World Bank, *World Development Report 1992* (New York: Oxford University Press, 1992); Gene M. Grossman and Alan B. Krueger, "Economic Growth and the Environment," *Quarterly Journal of Economics* 110(2): 353–377 (1995).
8. See William T. Harbaugh, Arik Levinson, and David Molloy Wilson, "Reexamining the Empirical Evidence for an Environmental Kuznets Curve," *Review of Economics and Statistics* 84(3): 541–551 (2002); and Cynthia C.-Y. Lin and Zachary D. Liscow, "Endogeneity in the Environmental Kuznets Curve: An Instrumental Variables Approach," *American Journal of Agricultural Economics* 95(2): 268–274 (2013).
9. See Charles D. Kolstad, "Interpreting Estimated Environmental Kuznets Curves for Greenhouse Gases," *Journal of Environment and Development* 15(1): 42–49 (2006).
10. See Joseph E. Aldy, "An Environmental Kuznets Curve Analysis of U.S. State-Level Carbon Dioxide Emissions," *Journal of Environment and Development* 14(1): 48–72 (2005).
11. See Nordhaus, "Lethal Model 2"; and Jeffrey A. Krautkraemer, "Economics of Scarcity: The State of the Debate," in David Simpson, Michael A. Toman, and Robert U. Ayres, eds., *Scarcity and Growth Revisited: Natural Resources and the Environment in the New Millennium* (Washington, DC: RFF Press, 2005), 54–77.
12. For a summary of these trends, see Nordhaus, "Lethal Model 2."
13. It is also possible that declining ore quality is contributing to some price decreases.
14. See Martin L. Weitzman, "Pricing the Limits to Growth from Minerals Depletion," *Quarterly Journal of Economics* 114(2): 691–706 (1999).
15. See World Commission on Environment and Development, *Our Common Future*, Report of the United Nations World Commission on Environment and Development (Oxford: Oxford University Press, 1987).
16. Robert M. Solow, "Sustainability: An Economist's Perspective," J. Seward Johnson Lecture to the Marine Policy Center, June 14, Woods Hole Oceanographic Institution, Woods Hole, Massachusetts. Reprinted in Robert N. Stavins, ed., *Economics of the Environment: Selected Readings*, 6th ed. (New York: W.W. Norton and Company, 2012), 543–550.
17. This idea is due to Robert N. Stavins, A. F. Wagner, and G. Wagner, "Interpreting Sustainability in Economic Terms: Dynamic Efficiency plus Intergenerational Equity," *Economic Letters* 79: 339–343 (2003).
18. Thomas C. Schelling, "The Cost of Combating Global Warming," *Foreign Affairs* 76(6): 8–4 (1997).
19. Capital depreciation is simply the loss in value of an economy's capital equipment due to use and the passage of time. For example, this year's U.S. GDP would include the total value of tractors produced by labor and property in the United States this year, net of the decrease in value of all tractors produced by labor and property in the United States in this and previous years.
20. See Arthur C. Pigou, *The Economics of Welfare* (London: Macmillan and Co., 1932), part 1, Chapter 3, available at www.econlib.org/library/NPDBooks/Pigou/pgEW3.html.
21. Graham Davis, Interview for "Measuring Our Worth," *Living on Earth*, National Public Radio (April 9, 2004).
22. See William Nordhaus and James Tobin, "Is Growth Obsolete," Cowles Foundation Working Paper 398, New Haven, CT: Yale University, reprinted from Milton

- Moss, ed., *The Measurement of Economic and Social Performance: Studies in Income and Wealth*, Vol. 38 (Cambridge, MA: National Bureau of Economic Research, 1973).
23. See Robert Repetto, William Magrath, Michael Wells, Christine Beer, and Fabrizio Rossini, *Wasting Assets: Natural Resources in the National Income Accounts* (Washington, DC: World Resources Institute, 1989).
 24. See William D. Nordhaus and Edward C. Kockelenberg, eds., *Nature's Numbers: Expanding the National Economic Accounts to Include the Environment* (Washington, DC: National Academy Press, 1999).
 25. Nicholas Z. Muller, Robert Mendelsohn, and William Nordhaus, "Environmental Accounting for Pollution in the United States Economy," *American Economic Review* 101(5): 1649–1675.
 26. See Martin L. Weitzman and Karl-Gustaf Löfgren, "On the Welfare Significance of Green Accounting as Taught by Parable," *Journal of Environmental Economics and Management* 32: 139–153 (1992).
 27. This phrase is due to Robert Solow, "An Almost Practical Step toward Sustainability," Lecture on the occasion of the 40th anniversary, Resources for the Future, Washington, DC (1992).
 28. The following is summarized from Kenneth Arrow, Partha Dasgupta, Lawrence Goulder, Gretchen Daily, Paul Erlich, Geoffrey Heal, Simon Levin, Karl-Göran Mäler, Stephen Schneider, David Starrett, and Brian Walker, "Are We Consuming Too Much?" *Journal of Economic Perspectives* 18(3): 147–172, especially table 2 and pp. 163–166.
 29. See A. B. Jaffe, S. R. Peterson, P. R. Portney, and R. N. Stavins, "Environmental Regulation and the Competitiveness of U.S. Manufacturing: What Does the Evidence Tell Us?" *Journal of Economic Literature* 33: 132–63 (March 1995); and Arik Levinson, "Environmental Regulations and Industry Location: International and Domestic Evidence," in J. N. Bhagwati and R. E. Hudec, eds., *Fair Trade and Harmonization: Prerequisites for Free Trade?* Vol. 1 (Cambridge, MA: MIT Press, 1996), 429–457.
 30. See Smita B. Brunnermeier and Arik Levinson, "Examining the Evidence on Environmental Regulations and Industry Location," *Journal of Environment and Development* 13(1): 6–41 (2004).
 31. See Josh Ederington, Arik Levinson, and Jenny Minier, "Footloose and Pollution-Free," *Review of Economics and Statistics* 87(1): 92–99 (2005).
 32. See Arik Levinson, "Environmental Regulatory Competition: A Status Report and Some New Evidence," *National Tax Journal* 56(1): 91–106 (2003).
 33. Carolyn Fisher, "Does Trade Help or Hinder the Conservation of Natural Resources?" *Review of Environmental Economics and Policy* 4(1): 103–121 (2010).
 34. M. Scott Taylor, "Buffalo Hunt: International Trade and the Virtual Extinction of the North American Bison," *American Economic Review* 101(7): 1–36 (2011).
 35. For a succinct description of this point of view, see the first issue of the journal *Ecological Economics*, published in February 1989.

Chapter 12: Conclusion

1. Hal Quinn, "EPA's New Greenhouse Gas Regulations Are Overzealous," *USA Today*, June 2, 2014.

2. Union of Concerned Scientists, "The Clean Power Plan: A Climate Game Changer," available at <http://www.ucsusa.org/our-work/global-warming/reduce-emissions/what-is-the-clean-power-plan#.VOthlMbqPtZ>, accessed February 23, 2015.
3. U.S. Environmental Protection Agency, "Regulatory Impact Analysis for the Proposed Carbon Pollution Guidelines for Existing Power Plants and Emission Standards for Modified and Reconstructed Power Plants," EPA-452/R-14-002, Research Triangle Park, NC: EPA Office of Air Quality Planning and Standards, June 2014, available at <http://www2.epa.gov/sites/production/files/2014-06/documents/20140602ria-clean-power-plan.pdf>, accessed February 23, 2015.
4. For example, Tom Zeller, writing in *Bloomberg View* (<http://www.bloombergvie.com/articles/2014-06-01/obama-s-carbon-rules-can-boost-the-economy>), cites an article in the *Van Nuys Valley News* from September 10, 1970, quoting Ford Motor Company officials who stated that the Clean Air Act could "cut off automobile production in 1975, lead to huge price increases for cars even if production were not stopped, do 'irreparable damage' to the American economy—and still lead to only small improvements in the quality of the air." Claims that environmental regulation actually generates net benefits without counting the benefits to the environment—that it *reduces* firms' production costs, on average—are similarly deceptive.
5. See Nicholas Z. Muller and Robert Mendelsohn, "Efficient Pollution Regulation: Getting the Prices Right," *American Economic Review* 99(5): 1714–1739 (2009).
6. See Randolph M. Lyon and Scott Farrow, "An Economic Analysis of Clean Water Act Issues," *Water Resources Research* 31: 213–223 (1995).
7. See Peter Gordon and Harry W. Richardson, "Prove It: The Costs and Benefits of Sprawl," *Brookings Review* (Fall): 23–25 (1998).

Further Reading

Chapter 2

- Brennan, Timothy J. 1999. "Discounting the Future: Economics and Ethics," in Wallace E. Oates, ed., *The RFF Reader in Environmental and Resource Management*, Resources for the Future, Washington, DC, 35–41.
- Goulder, Lawrence H., and Robert N. Stavins. 2002. "An Eye on the Future," *Nature* 419 (October): 673–674.
- Mankiw, Gregory N. 2015. "Ten Principles of Economics," in *Principles of Economics*, 7th ed., Cengage Learning, Stamford, CT, pp. 3–18.
- Nordhaus, William D. 1999. "Discounting and Public Policies that Affect the Distant Future," in Paul R. Portney and John P. Weyant, eds., *Discounting and Intergenerational Equity*, Resources for the Future, Washington, DC, 145–162.
- Oates, Wallace E. 2010. "An Economic Perspective on Environmental Policy and Resource Management: An Introduction," in Wallace E. Oates, ed., *The RFF Reader in Environmental and Resource Policy*, 2nd ed., Resources for the Future, Washington, DC.
- Portney, Paul, and John Weyant. 1999. *Discounting and Intergenerational Equity*, Resources for the Future, Washington, DC.
- Weitzman, Martin L. 1998. "Why the Far-Distant Future Should Be Discounted at Its Lowest Possible Rate," *Journal of Environmental Economics and Management* 36(3): 201–208.

Chapter 3

- Arrow, Kenneth J., Maureen L. Cropper, George C. Eads, Robert W. Hahn, Lester B. Lave, Roger G. Noll, Paul R. Portney, Milton Russell, Richard Schmalensee, V. Kerry Smith, and Robert N. Stavins. 1996. "Is There a Role for Benefit–Cost Analysis in Environmental, Health and Safety Regulation?" *Science* 272: 221–222.
- Burtraw, Dallas, Alan J. Krupnick, Erin Mansur, David Austin, and Deirdre Farrell. 1998. "The Costs and Benefits of Reducing Air Pollutants Related to Acid Rain," *Contemporary Economic Policy* 16: 379–400.
- Freeman, A. Myrick III. 2003. *The Measurement of Environmental and Resource Values*, 2nd ed., Resources for the Future, Washington, DC.

- Kelman, Steven. 1981. "Cost-Benefit Analysis: An Ethical Critique," with replies by James Delong, Robert Solow, Gerard Butters, John Calfee, and Pauline Ippolito, *AEI Journal on Government and Social Regulation* (January/February): 33–40. Reprinted in Robert N. Stavins, ed. 2012. *Economics of the Environment: Selected Readings*, 6th ed., Norton, New York, 254–269.
- Metrick, Andrew, and Martin L. Weitzman. 1998. "Conflict and Choices in Biodiversity Preservation," *Journal of Economic Perspectives* 12(3): 21–34.
- Morgenstern, Richard D., ed. 1997. *Economic Analyses at EPA: Assessing Regulatory Impact*, Resources for the Future, Washington, DC.
- National Center for Environmental Economics, U.S. Environmental Protection Agency. 2010. *Guidelines for Preparing Economic Analyses*. EPA 240-R-10-001, U.S. Environmental Protection Agency, Washington, DC. Available at <http://yosemite.epa.gov/ee/epa/eerm.nsf/vwAN/EE-0568-50.pdf/%24file/EE-0568-50.pdf>.
- Pasurka, Carl. 2008. "Perspectives on Pollution Abatement and Competitiveness: Theory, Data, and Analyses," *Review of Environmental Economics and Policy* 2: 194–218.
- Pizer, William, Matthew Adler, Joseph Aldy, David Anthoff, Maureen Cropper, Kenneth Gillingham, Michael Greenstone, Brian Murray, Richard Newell, Richard Richels, Arden Rowell, Stephanie Waldhoff, and Jonathan Wiener. 2014. "Using and Improving the Social Cost of Carbon," *Science* 346(6214): 1189–1190.
- Sen, Amartya K. 1970. *Collective Choice and Welfare*, Holden-Day, San Francisco.
- Toman, Michael. 1998. "Why Not to Calculate the Value of the World's Ecosystem Services and Natural Capital," *Ecological Economics* 25: 57–60.

Chapter 4

- Mankiw, N. Gregory. 2015. *Principles of Microeconomics*, 7th ed., Cengage Learning, Stamford, CT.
- Pindyck, Robert S., and Daniel L. Rubinfeld. 2013. *Microeconomics*, 8th ed., Pearson Prentice-Hall, Upper Saddle River, NJ.

Chapter 5

- Baumol, William, and Wallace Oates. 1988. *The Theory of Environmental Policy*, 2nd ed., Cambridge University Press, New York.
- Hardin, Garrett. 1968. "The Tragedy of the Commons," *Science* 162(3859): 1243–1248.
- Stavins, Robert N. 2011. "The Problem of the Commons: Still Unsettled After 100 Years," *American Economic Review* 101: 81–108.
- Stavins, Robert N., ed. 2012. *Economics of the Environment: Selected Readings*, 6th ed., Norton, New York.
- Tietenberg, Tom H., and Lynne Lewis. 2015. *Environmental and Natural Resource Economics*, 10th ed., Pearson Prentice-Hall, Upper Saddle River, NJ.

Chapter 6

- Heal, Geoffrey. 1998. *Valuing the Future: Economic Theory and Sustainability*, Columbia University Press, New York.
- Hotelling, Harold. 1931. "The Economics of Exhaustible Resources," *Journal of Political Economy* 39(2): 137–175.
- Jevons, Stanley. 1866. *The Coal Question*, Macmillan, London.
- Krautkraemer, Jeffrey A. 2005. "Economics of Scarcity: The State of the Debate," in

- R. David Simpson, Michael A. Toman, and Robert U. Ayres, eds., *Scarcity and Growth Revisited: Natural Resources and the Environment in the New Millennium*, Resources for the Future, Washington, DC, 54–77.
- Livernois, John R. 2009. “On the Empirical Significance of the Hotelling Rule,” *Review of Environmental Economics and Policy* 3(1): 22–41.
- Maugeri, Leonardo. 2009. “Understanding Oil Price Behavior through an Analysis of a Crisis,” *Review of Environmental Economics and Policy* 3: 147–166.
- Smith, James L. 2009. “World Oil: Market or Mayhem?” *Journal of Economic Perspectives* 23(3): 145–164.

Chapter 7

- Andam, Kwaw S., Paul J. Ferraro, Alexander Pfaff, G. Arturo Sanchez-Azofeifa, and Juan A. Robalino. 2008. “Measuring the Effectiveness of Protected Area Networks in Reducing Deforestation,” *Proceedings of the National Academy of Sciences* 105(42): 16089–16094.
- Berkes, F., D. Feeny, B. J. McCay, and J. M. Acheson. 1989. “The Benefits of the Commons,” *Nature* 340(6229): 91–93.
- Clark, Colin W. 1973. “Profit Maximization and the Extinction of Animal Species,” *Journal of Political Economy* 81(4): 950–961.
- Conrad, Jon M. 1999. *Resource Economics*, Cambridge University Press, New York.
- Costello, Christopher, S. Gaines, and J. Lynham. 2008. “Can Catch Shares Prevent Fisheries Collapse?” *Science* 321: 1678–1681.
- Field, Barry. 2001. *Natural Resource Economics: An Introduction*, McGraw-Hill, New York.
- Gordon, H. Scott. 1954. “The Economic Theory of a Common Property Resource: The Fishery,” *Journal of Political Economy* 62(2): 124–142.
- Iudicello, Suzanne, Michael Weber, and Robert Wieland. 1999. *Fish, Markets and Fishermen: The Economics of Overfishing*, Island Press, Washington, DC.
- Kurlansky, Mark. 1997. *Cod: A Biography of the Fish That Changed the World*, Penguin Books, New York.
- Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*, Cambridge University Press, New York.
- Stavins, Robert N. 2011. “The Problem of the Commons: Still Unsettled After 100 Years,” *American Economic Review* 101: 81–108.
- Wood, Charles, and Robert Walker. 1999. “Saving the Trees by Helping the Poor: A Look at Small Producers along Brazil’s Transamazon Highway,” *Resources* 136 (Summer): 14–17.

Chapter 8

- Anderson, Terry L., and Gary D. Libecap. 2014. *Environmental Markets: A Property Rights Approach*, Cambridge University Press, New York.
- Benbear, Lori S., and Cary Coglianese. 2013. “Flexible Approaches to Environmental Regulation,” in *The Oxford Handbook of U.S. Environmental Policy*, Oxford University Press, New York, 582–604.
- Coase, Ronald. 1969. “The Problem of Social Cost,” *Journal of Law and Economics* 3: 1–44.
- Dales, J. H. 2002. *Pollution, Property, and Prices: An Essay in Policy-Making and Economics*, Edward Elgar, Cheltenham, UK.

- Fung, Archon, Mary Graham, and David Weil. 2007. *Full Disclosure: The Perils and Promise of Transparency*, Cambridge University Press, New York.
- Goulder, Lawrence H., and Ian W. H. Parry. 2008. "Instrument Choice in Environmental Policy," *Review of Environmental Economics and Policy* 2(2): 152–174.
- Parry, Ian W. H., and Kenneth A. Small. 2005. "Does Britain or the United States Have the Right Gasoline Tax?" *American Economic Review* 95(4): 1276–1289.
- Sandel, Michael J. 1997. "It's Immoral to Buy the Right to Pollute," *New York Times*, December 15, with replies by Shavell, Stavins, Gaines, Maskin, and Leifman. Reprinted in Robert N. Stavins, ed., *Economics of the Environment: Selected Readings*, 6th ed., W.W. Norton and Company, New York, 344–348.
- Tietenberg, Tom H. 1990. "Economic Instruments for Environmental Regulation," *Oxford Review of Economic Policy* 6(1): 17–33.
- Weitzman, Martin L. 1974. "Prices v. Quantities," *Review of Economic Studies* 41: 477–491.

Chapter 9

- Aldy, Joseph E., Alan J. Krupnick, Richard G. Newell, Ian W. H. Parry, and William A. Pizer. 2010. "Designing Climate Mitigation Policy," *Journal of Economic Literature* 48(4): 903–934.
- Anderson, Terry L., and Gary D. Libecap. 2014. *Environmental Markets: A Property Rights Approach*, Cambridge University Press, New York.
- Farrow, R. Scott, Martin T. Schultz, Pinar Celikkol, and George L. Van Houtven. 2005. "Pollution Trading in Water Quality Limited Areas: Use of Benefits Assessment and Cost-Effective Trading Ratios," *Land Economics* 81(2): 191–205.
- Jaffe, Adams B., Richard G. Newell, and Robert N. Stavins. 2003. "Technological Change and the Environment," in Karl-Göran Mäler and Jeffrey Vincent, eds., *Handbook of Environmental Economics*, Vol. I, Elsevier Science, Amsterdam.
- Muller, Nicholas Z., and Robert Mendelsohn. 2009. "Efficient Pollution Regulation: Getting the Prices Right," *American Economic Review* 99(5): 1714–1739.
- Newell, Richard G., and Nathan E. Wilson. 2005. "Technology Prizes for Climate Change Mitigation," RFF Discussion Paper 05-33, Resources for the Future, Washington, DC.
- Tietenberg, Tom H. 2006. *Emissions Trading: Principles and Practice*, 2nd ed., Resources for the Future, Washington, DC.

Chapter 10

- Burtraw, Dallas, Alan Krupnick, Erin Mansur, David Austin, and Deirdre Farrell. 1998. "Costs and Benefits of Reducing Air Pollutants Related to Acid Rain," *Contemporary Economic Policy* 16(4): 379–400.
- Ellerman, A. Denny, Frank J. Convery, and Christian de Perthuis. 2010. *Pricing Carbon: The European Union Emissions Trading Scheme*, Cambridge University Press, New York.
- Ellerman, A. Denny, Paul J. Joskow, Richard Schmalensee, Juan-Pablo Montero, and Elizabeth M. Bailey. 2000. *Markets for Clean Air: The U.S. Acid Rain Program*, Cambridge University Press, New York.
- Fisher-Vanden, Karen A., and Sheila M. Olmstead. 2013. "Moving Pollution Trading from Air to Water: Potential, Problems and Prognosis," *Journal of Economic Perspectives* 27(1): 147–172.

- Freeman, Jody, and Charles D. Kolstad. 2007. *Moving to Markets in Environmental Regulation: Lessons from Twenty Years of Experience*, Oxford University Press, New York.
- Huang, Ju-Chin, John M. Halstead, and Shanna B. Saunders. 2011. "Managing Municipal Solid Waste with Unit-Based Pricing: Policy Effects and Responsiveness to Pricing," *Land Economics* 87(4): 645–660.
- Iudicello, Suzanne, Michael Weber, and Robert Wieland. 1999. *Fish, Markets and Fishermen: The Economics of Overfishing*, Island Press, Washington, DC.
- Keohane, Nathaniel O. 2007. "Cost Savings from Allowance Trading in the 1990 Clean Air Act," in Charles E. Kolstad and Jody Freeman, eds., *Moving to Markets in Environmental Regulation: Lessons from Twenty Years of Experience*, Oxford University Press, New York, 194–229.
- Langpap, Christian, and Joe Kerkvliet. 2012. "Endangered Species Conservation on Private Land: Assessing the Effectiveness of Habitat Conservation Plans," *Journal of Environmental Economics and Management* 64: 1–15.
- Mansur, Erin T., and Sheila M. Olmstead. 2012. "The Value of Scarce Water: Measuring the Inefficiency of Municipal Regulations," *Journal of Urban Economics* 71: 332–346.
- McConnell, Virginia, and Margaret Walls. 2009. "U.S. Experience with Transferable Development Rights," *Review of Environmental Economics and Policy* 3(2): 288–303.
- Newell, Richard G., James N. Sanchirico, and Suzi Kerr. 2005. "Fishing Quota Markets," *Journal of Environmental Economics and Management* 49(3): 437–462.
- Popp, David. 2003. "Pollution Control Innovations and the Clean Air Act of 1990," *Journal of Policy Analysis and Management* 22(4): 641–660.
- Sanchirico, James, and Richard Newell. 2003. "Catching Market Efficiencies," *Resources* 150: 8–11.
- Schmalensee, Richard, and Robert N. Stavins. 2013. "The SO₂ Allowance Trading System: The Ironic History of a Grand Policy Experiment," *Journal of Economic Perspectives* 27: 103–122.
- Stavins, Robert N. 2003. "Experience with Market-Based Environmental Policy Instruments," in Karl-Göran Mäler and Jeffrey R. Vincent, eds., *Handbook of Environmental Economics*, Vol. 1, Elsevier Science B.V., Amsterdam, 355–435.
- Stern, Thomas. 2003. *Policy Instruments for Environmental and Natural Resource Management*, Resources for the Future, Washington, DC., 363.
- World Bank Group. 2014. *State and Trends of Carbon Pricing 2014*, World Bank, Washington, DC. Available at <http://documents.worldbank.org/curated/en/2014/05/19572833/state-trends-carbon-pricing-2014>.

Chapter 11

- Arrow, Kenneth, Partha Dasgupta, Lawrence Goulder, Gretchen Daily, Paul Erlich, Geoffrey Heal, Simon Levin, Karl-Göran Mäler, Stephen Schneider, David Starrett, and Brian Walker. 2004. "Are We Consuming Too Much?" *Journal of Economic Perspectives* 18(3): 147–172.
- Brunnermeier, Smita B., and Arik Levinson. 2004. "Examining the Evidence on Environmental Regulations and Industry Location," *Journal of Environment and Development* 13(1): 6–41.
- Carson, Richard T. 2010. "The Environmental Kuznets Curve: Seeking Empirical Regularity and Theoretical Structure," *Review of Environmental Economics and Policy* 4(1): 3–23.

- Ederington, Josh, Arik Levinson, and Jenny Minier. 2005. "Footloose and Pollution-Free," *Review of Economics and Statistics* 87(1): 92–99.
- Fischer, Carolyn. 2010. "Does Trade Help or Hinder the Conservation of Natural Resources?" *Review of Environmental Economics and Policy* 4(1): 103–121.
- Frankel, Jeffrey A., and Andrew K. Rose. 2005. "Is Trade Good or Bad for the Environment? Sorting Out the Causality," *Review of Economics and Statistics* 87(1): 95–91.
- Friedman, Benjamin M. 2005. *The Moral Consequences of Economic Growth*, Knopf, New York.
- Jaffe, A. B., S. R. Peterson, P. R. Portney, and R. N. Stavins. 1995. "Environmental Regulation and the Competitiveness of U.S. Manufacturing: What Does the Evidence Tell Us?" *Journal of Economic Literature* 33(March): 132–163.
- Krautkraemer, Jeffrey A. 2005. "Economics of Scarcity: The State of the Debate," in David Simpson, Michael A. Toman, and Robert U. Ayres, eds., *Scarcity and Growth Revisited: Natural Resources and the Environment in the New Millennium*, RFF Press, Washington, DC, 54–77.
- Nordhaus, William D. 1992. "Lethal Model 2: The Limits to Growth Revisited," *Brookings Papers on Economic Activity* 1992(2): 1–59.
- Nordhaus, William D., and Edward C. Kokkelenberg, eds. 1999. *Nature's Numbers: Expanding the National Economic Accounts to Include the Environment*, National Academy Press, Washington, DC.
- Solow, Robert M. 1991. "Sustainability: An Economist's Perspective," J. Seward Johnson Lecture to the Marine Policy Center, June 14, Woods Hole Oceanographic Institution, Woods Hole, MA. Reprinted in Robert N. Stavins, ed., *Economics of the Environment: Selected Readings*, 6th ed., W.W. Norton and Company, New York, 543–550.
- Solow, Robert. 1992. "An Almost Practical Step toward Sustainability," Lecture on the occasion of the 40th anniversary, Resources for the Future, Washington, DC.
- Stiglitz, Joseph. 2005. "The Ethical Economist," a review of *The Moral Consequences of Economic Growth*, by Benjamin M. Friedman, *Foreign Affairs* (November/December): 128–134.
- Weitzman, Martin L. 1999. "Pricing the Limits to Growth from Minerals Depletion," *Quarterly Journal of Economics* 114(2): 691–706.

Index

Pages numbers followed by “f” and “t” indicate figures and tables.

- Abatement, defined, 13
- Abatement benefit functions, 16, 17f
- Abatement cost functions, 14–15, 17f
- Ability to pay, willingness to pay and, 47
- Accounting, green, 232, 245–249
- Acid rain, 200, 203
- Adoption, innovation vs., 183
- Adsorbable organic halides (AOX), 23
- Africa, 250
- Aggregate marginal abatement costs, 176
- Aggregates, 175–176
- Alcoa, 223
- Aleutian Islands fisheries, 135
- Allocations, 145, 219–220. *See also*
 - Allowance trading
- Allowance trading, 145. *See also*
 - Cap-and-trade systems
- Aluminum recycling, 223
- Ansari X-Prize, 185
- AOX. *See* Adsorbable organic halides
- Aracruz Celulose, 23
- Arctic National Wildlife Refuge, 31
- Arrow, Kenneth, 59
- Arsenic, 57
- Assumptions in cost models, 39
- Asymmetric information, 146–147
- Backloading, 220
- Bargaining, private, 140–143
- Beaver hunting, 134
- Benefit-cost analysis
 - critiques of, 57–65
 - EPA study of lead in gasoline, 58
 - overview of, 7, 55–57
 - uncertainty and, 65–68
- Benefits. *See also* Marginal benefits; Values
 - defining, 44–48
 - measuring, 49–55
 - overview of, 6
 - of policies to mitigate climate change, 4–5
 - types of values and, 48
- Bering Sea fisheries, 135
- Biodiversity, as public good, 85
- Bioeconomic models, 114, 128–130, 134
- Biological rotation, 115–116, 116f
- Bottom-up cost models, 38–39
- Boundaries, 194–196
- Brazil, 125–126, 225
- Brundtland Commission, 239
- Bush, George H.W., 200–201
- Bush, George W., 202
- Bycatch, 137, 214
- Byrd, Robert, 61–62
- California condor, 24, 25f
- CalTrans, 227
- Canals, 63

- Cap-and-trade systems. *See also*
 Allowance trading
 advantages of, 178
 cost effectiveness and, 177–178
 emissions taxes vs., 162–166, 163f,
 178–179
 hot spot problem and, 190–191
 markets for greenhouse gas emission
 trading, 157–158, 208–210,
 209f
 missing demand curve and, 152–156,
 154f, 156f
 for nitrogen oxides, 192
 raising revenues and, 160–162
 taxes vs., 182
 for U.S. sulfur dioxide market,
 200–207, 204t
 water quality trading, 221–223
- Capital, interchangeability of, 240
- Capital assets, 99–100, 115
- Capitalism, 70
- Carbon emission abatement. *See also*
 Greenhouse gas emissions
 benefits of, 4–6
 Clean Power Plan and, 254–255
 costs of, 38–39, 39f
 standards for new power plants,
 66–67, 67f
- Carbon sequestration, 122–123
- Cartels, 110
- Catalytic converters, 197
- Catch shares, 158–160, 187–188, 188f,
 207–214
- CCX. *See* Chicago Climate Exchange
- Certification programs, 146–147
- CFCs. *See* Chlorofluorocarbons
- Chesapeake Bay, 221–223
- Chicago Climate Exchange (CCX),
 157–158
- Chile, 210
- China, 110–111, 210, 250
- Chlorofluorocarbons (CFCs), 28–29
- Clean Air Act, 44, 55, 197, 200, 255
- Clean Air Act Amendments of 1977,
 144, 204
- Clean Air Act Amendments of 1990,
 61, 201
- Clean Development Mechanism,
 208, 218
- Clean Power Plan, 254
- Clean Water Act, 54, 221, 223, 255
- Climate change, 2–6. *See also* Carbon
 emission abatement; Greenhouse
 gas emissions; Sulfur dioxide
 abatement
- Club goods, 86
- Coal industry, 61–62
The Coal Question (Jevons), 101
- Coase, Ronald, 139, 140, 167
- Coase Theorem, 133–134, 140–143,
 159–160
- Co-benefits, 4
- Collective action problems, 94–97, 97f
- Columbia River, 2
- Command-and-control. *See also*
 Performance standards; Technology
 standards
 advantages of, 197
 cost effectiveness and, 172–173
 fisheries and, 159
 market-based policies vs., 196–198
 monitoring and enforcement and,
 195
 municipal water pricing and, 215
 overview of, 139, 144
- Common property, open access vs.,
 132–134
- Companies, role of, 256–257
- Completeness, market efficiency and,
 77–78
- Compliance, 206–207, 213–214, 217
- Concave functions, 16
- Condors, 24, 25f
- Congestion, 153
- Conservation banking, 226–228
- Consolidation, 211–212
- Consumers, role of, 256–257
- Consumer surplus, 74–75, 76–77
- Contingent valuation (CV), 52–55,
 122–123, 124
- Copenhagen Accord, 5
- Costanza, Robert, 56
- Cost-benefit analysis. *See* Benefit-cost
 analysis

- Cost effectiveness
 command-and-control approaches and, 172–173
 emission taxes and, 174–176, 175f, 176f
 emission trading and, 177–178
 heterogeneity and, 179
 overview of, 170–172
- Cost reduction
 command-and-control approaches and, 172–173
 cost effectiveness and, 170–172
 emission trading and, 177–178
 overview of, 169–170
- Costs. *See also* Benefit-cost analysis;
 Marginal costs
 of compliance in private sector, 36–37
 of government sector compliance, 36, 37–40
 market-based instruments for reduction of, 169–179, 170f
 measuring, 35–44
 of mitigating greenhouse gas emissions, 38–39, 39f
 overview of, 6
 of policies to mitigate climate change, 5
 social welfare, 36, 37
 of transitional effects, 36
- CV. *See* Contingent valuation
- Dales, J.H., 159–160
- Damage assessment, 54
- Deadweight losses, 84–85, 164–166
- Dead zones, 222
- De Beers, 77
- Decentralization, 71, 146
- Decision-making, criteria for, 57–59
- Deepwater Horizon oil spill, 54
- Deforestation, 123–128. *See also* Forest resources
- Demand
 benefits to consumers and, 74–75, 75f
 markets and, 71–72, 72f
 for pollution control, 152–156
- Demand curves, 71–72, 72f, 152–156
- Depletion externality, 137
- Desalination, 240–241
- Development rights, 225
- Diminishing marginal returns, 93
- Dioxin, 23
- Directional drilling, 102
- Director, Aaron, 141
- Discounting, 31–33, 59–60
- Discount rate, 33t, 41, 42–43, 43f
- Disease, 4. *See also* Specific diseases
- Distortionary taxes, 152, 161–162
- Dolphins, 137, 214
- Double dividends, 162
- Drinking water, 240–241
- Driving, externalities of, 153
- Drought, municipal water pricing and, 214–217
- Dukakis, Michael, 201
- Dynamic efficiency, 7, 30–34, 243–244
- Dynamic incentives, 168
- ECF methods. *See* Elemental chlorine free methods
- Ecolabeling programs, 146–147
- Ecological economics, 252
- Economic efficiency. *See also* Efficiency
 dynamic efficiency and, 30–34
 environmental policy and, 20–22
 equimarginal rule and, 26–30
 overview of, 6, 12–13, 168
 sulfur dioxide abatement and, 14–20, 17f
 technical efficiency vs., 16, 19
- Economic growth
 conclusions from theory and experience and, 249–250
 green accounting and, 245–249
 gross domestic product and, 250
 limits to, 232–238
 overview of, 231–232, 253
 sustainability and, 238–245, 252
 trade, growth, environment and, 251–252
- Economic overfishing, 131
- Economic scarcity, 100–103. *See also* Scarcity
- Economics of Welfare* (Pigou), 140
- Economies of scale, 21
- EDF. *See* Environmental Defense Fund

- Education, perception of value and, 46
- Efficiency. *See also* Economic efficiency;
Market efficiency
benefit-cost analysis and, 55–57
competitive market equilibrium and,
70–74
of competitive markets, 74–78
defined, 12–13, 14
overview of, 69–70, 78–79
- Ehrlich, Paul, 236
- EKC. *See* Environmental Kuznets curve
- Elemental chlorine free (ECF) methods, 23
- Emission rates, performance standards
and, 144
- Emissions reduction, 155–156, 156f. *See
also* Carbon emission abatement;
Greenhouse gas emissions; Sulfur
dioxide abatement
- Emissions taxes
cap-and-trade policies vs., 162–166,
163f, 178–179
cost effectiveness of, 174–176,
175f, 176f
hot spot problem and, 190–191
as market-based policies, 144–145
- Emissions trading. *See* Cap-and-trade
systems
- Emissions Trading System. *See* European
Union Emissions Trading System
- Employment. *See* Jobs
- Endangered Species Act (ESA), 124, 227
- Endangered species preservation,
226–228
- Endogeneity, 251
- Energy efficiency paradox, 15–16, 38, 148
- Enforcement
individual fishing quotas and, 213–214
market-based policies and, 194–196
municipal water supplies and, 217
sulfur dioxide abatement and, 206–207
- Engineering costs, 37
- Environmental Defense Fund (EDF),
200, 204, 228
- Environmental economics, overview of,
1–2, 254–256
- Environmental Kuznets curve (EKC), 235
- Environmental policies. *See* Policies
- Equilibrium, 72–74, 72f, 149–152, 212
- Equimarginal rule, 26–30, 34
- Equity, intergenerational and
intragenerational, 244–245
- Equivalence, determining, 228–229
- European Union Allowances (EUA), 218
- European Union Emissions Trading
System (EU-ETS), 157–158, 209,
217–220
- Exchanges, markets and, 71
- Exclusion, tragedy of the commons
and, 160
- Expected marginal costs, 163–164, 163f
- Externalities. *See also* Marginal user costs
Coase Theorem and, 139, 140–143
externalities and, 81–85
fisheries and, 137
market failures and, 81–85, 83f
Pigouvian taxes and, 151–152
research and development and, 184
- Exxon Valdez* oil spill, 46, 54
- Faustmann, Martin, 122
- Faustmann Rule, 118–122, 121f
- Federal Payments for Hydrological
Services, 127
- Federal Water Pollution Control Act, 195
- Firms, role of, 256–257
- Fisheries
bioeconomic model of, 128–130
catch shares and, 158–160
cost-effectiveness and, 184–188, 190
efficient management of vs. open
access, 130–132, 131f
incentives for technological change
and, 189–190
logistic growth and, 128, 129f
market-based instruments for in New
Zealand, 207–214
overview of, 137–138
property rights for halibut in Gulf of
Alaska, 159–160
real-world, economics and, 135
subsidies, externalities and, 136–137
substitution, sustainability and, 241
use of steady-state approach and,
134–135

- Flexibility, 173, 228
- Forest resources
 Faustmann Rule and, 118–122, 121f
 forest growth, biological rotation and, 115–116
 nontimber benefits, efficient rotation and, 122–123, 123f, 124
 overview of, 137–138
 overview of economics of, 114–115
 public goods, property rights, deforestation and, 123–128
 Wicksell Rule and, 116–118
- Forest Stewardship Council, 146
- Fossil fuel consumption, perverse incentives and, 151
- Fracking, 102
- Free markets, 3, 78
- Free-riding problem, 90
- FRONAFOR. *See* Federal Payments for Hydrological Services
- Future value, discounting and, 31–33
- Gasoline, 58, 153
- GATT. *See* Global Agreement on Tariffs and Trade
- GDP. *See* Gross domestic product
- Genuine wealth, 250
- Global Agreement on Tariffs and Trade (GATT), 251
- Global climate change. *See* Climate change
- Global ecosystem services, valuation of, 56
- GNP. *See* Gross national product
- Gore, Al, 69
- Governments, role of, 256–257
- Grand Banks fisheries, 135
- Green accounting, 232, 245–249
- Greenhouse gas emissions, 208–210, 209f. *See also* European Union Emissions Trading System
- Green net national product (green NNP), 246, 247, 248, 249–250
- Gross domestic product (GDP), 245, 250
- Gross national product (GNP), 245, 247
- Gross value, net benefit vs., 11–12
- Groundwater as open-access resource, 111–112
- Growth. *See* Economic growth; Logistic growth
- Gulf of Mexico fisheries, 135
- Gulf Stream, reversal of, 3
- Habitat Conservation Plans (HCP), 227
- Habitat management, 224–225
- Habitat value, 122, 123f
- Halibut, Pacific, 159–160, 189
- Hardin, Garrett, 91
- Harrison, John, 185
- HCP. *See* Habitat Conservation Plans
- Hedonic property method, 50–51
- Hedonic wage method, 51–52
- Heterogeneity, 179, 190, 251
- Hicks, John, 63–64
- High grading, 213
- Home energy reports, 148
- Hotelling, Harold, 109
- Hotelling Rule, 108–110, 111, 113
- Hot spot problem, 190–194
- Hunter River Salinity Trading program, 221
- IAM. *See* Integrated assessment models
- IFQ markets. *See* Individual fishing quota markets
- Illegal dumping, 224
- Implicit pricing, 50
- Incentives. *See also* Market-based instruments
 causes of climate change and, 3
 protected areas and, 126
 as root cause of environmental problems, 98
 for technological change, 180–183, 181f, 185, 189
 wrong pricing and, 256
- Income, growth, environment and, 235
- India, 151, 250
- Indirect use values, 48, 54
- Individual fishing quota (IFQ) markets, 158–160, 187–188, 188f, 207–214
- Indonesia, 151

- Information-based approaches, 146–147
- Innovation
 adoption vs., 183
 European Union Emissions Trading System and, 218–219
 policies aimed at promoting, 183–184
 prizes and, 185
- Input standards, 195
- Integrated assessment models (IAM), 41–42
- Interest rates, 108–109
- Intergenerational equity, 244–245
- Internalization of damages, 84, 145
- Invisible hand, 69
- James Bay, Quebec, 134
- Jevons, Stanley, 101, 233
- Jobs, environmental regulations and, 39, 43–44
- Kaldor, Nicholas, 63–64
- Kaldor-Hicks criterion, 64
- Kelvin, Lord, 242
- Krupp, Fred, 200
- Kuznets, Simon, 246
- Kuznets curve, 235
- Kyoto Protocol, 157, 208, 217–218
- Land as scarce resource, 124
- Landfills, 223–224
- Landing taxes, 160
- Land management, 224–225
- Land prices, 119
- Land use, forest resources and, 118–119
- Lead, in gasoline, 58
- Limbaugh, Rush, 69
- The Limits to Growth* (Meadows), 233–234
- Lindbergh, Charles, 185
- Linkages, 195, 209
- Logistic growth, fisheries and, 128, 129f
- Longitude Act, 185
- Lovins, Amory, 69
- Luxottica, 77
- Macroeconomics, overview of, 8–9, 231–232
- MAI. *See* Mean annual increment
- Malthus, Thomas, 233
- Maori people, 212–213
- Marginal benefit functions, 24, 27f
- Marginal benefits
 demand curves as measure of, 75, 75f
 dynamic efficiency and, 34
 overview of, 22–24
 public goods and, 87–88, 88f
 total benefits and, 30
- Marginal cost curve, 75–76, 75f
- Marginal cost functions, 22–24, 24f, 171–172, 180
- Marginal costs
 cap-and-trade policies vs. emissions taxes and, 162–166, 178–179
 cost reduction and, 169–172, 170f, 186–187, 187f
 dynamic efficiency and, 34
 overview of, 22–24, 24f
 total costs and, 30
- Marginal damages (MD), 82–83, 83f
- Marginal extraction costs, 100
- Marginal net benefit costs, 106, 106f
- Marginal returns, diminishing, 93
- Marginal user costs, 104, 107–108, 112, 236
- Marginal willingness to pay, 49–53, 56, 71, 74–75, 90
- Marine Mammal Protection Act, 137
- Marine Stewardship Council, 146
- Market-based instruments
 advantages of, 198
 Coase Theorem and, 140–143
 command-and-control vs., 196–198
 conservation banking, 226–228
 cost reduction and, 169–179
 creating property rights, 158–160
 emissions trading system of E. U., 157–158, 217–220
 emission taxes and allowance trading, 144–145, 174–178, 175f, 176f
 for fishing in New Zealand, 207–214
 getting prices right and, 147–152, 149f
 for greenhouse gas emissions, 208–210, 209f
 greenhouse gas emissions trading, 157–158

- for habitat and land management, 224–225
- information-based approaches to, 145–147
- for managing natural resources, 184–189, 187f, 188f
- missing demand curve and, 152–156, 154f, 156f
- monitoring and enforcement and, 194–196
- for municipal water pricing, 214–217
- nudges and, 147, 148
- optimal gasoline tax, 153
- overview of, 139, 144–145, 167, 168–169, 198, 199–200, 228–230
- Perrier, Vitell and, 143
- pollution hot spots and, 190–194
- property rights for halibut in Gulf of Alaska, 159–160
- raising revenues and, 160–162
- setting prices vs. setting quantities and, 162–166
- technological change promotion and, 179–184, 181f, 185
- tradeable development rights, 225
- for U.S. sulfur dioxide market, 200–207, 204t
- for waste management, 223–224
- water quality trading, 221–223
- wetland mitigation banking, 225–226
- Market efficiency, overview of, 69–70, 76–78
- Market equilibrium. *See* Equilibrium
- Market failures. *See also* Externalities; Public goods; Tragedy of the commons
 - importance of correcting, 257–258
 - market-based policies to overcome, 147–160, 149f
 - overview of, 7, 78, 80–81, 97–98, 147
 - policies to address, 7–8
- Market power, 110–111
- Markets
 - demand, benefits to consumers and, 74–75, 75f
 - demand and supply and, 71–72, 72f
 - equilibrium and, 72–74
 - nonrenewable resources and, 110–111
 - overview of, 7, 70–71
 - supply, cost to producers and, 75–76, 75f
- Market stability reserve, 220
- Markey, Ed, 208
- Maximum abatement, 18
- Maximum contaminant level (MCL) standards, 55
- Maximum sustainable yield (MSY), 129–130, 129f, 207
- McCain, John, 208
- McKelvey Diagrams, 100, 101f, 108, 236
- MCL standards. *See* Maximum contaminant level standards
- MD. *See* Marginal damages
- Mean annual increment (MAI), 115–118
- Mexico, 127, 210
- Microeconomics, 8–9, 231
- Microsoft, 77
- Mileage taxes, 153
- Millikan, Robert, 243
- Mitigation banking, 225–226, 228–229
- Mixing, 190. *See also* Heterogeneity; Uniform mixing assumption
- Monetary valuation as cheapening worth of environment, 60–61
- Monitoring, market-based policies and, 194–196
- Monopolies, 77, 110
- Monte Carlo analysis, 65, 68
- Morality, valuation and, 60–61
- Morgenstern, Richard, 55
- MSY. *See* Maximum sustainable yield
- Municipal water pricing, 107–108, 214–217
- Myopic state efficiency, 103–104, 104f
- NAFTA. *See* North American Free Trade Agreement
- Natural resources. *See also* Nonrenewable resources; Renewable resources; *Specific resources*
 - economic scarcity of, 234–238
 - market-based instruments for management of, 184–189, 187f, 188f
 - rate of extraction of, 252
 - tragedy of the commons and, 93

- NDP. *See* Net domestic product
- Negative externalities, defined, 80
- Nestlé Waters Group, 143
- Net benefits, 11–12, 64
- Net domestic product (NDP), 246
- Net national product (NNP), 246, 247–248, 249–250
- New Zealand, 158–160, 187–188, 188f, 207–214
- Nitrogen oxides (NO_x), 192
- NNP. *See* Net national product
- No-arbitrage conditions, 109, 116
- Nonexcludability, 85–86, 87f, 93, 98
- Nonrenewable resources
 - decreased scarcity of, 102
 - economic scarcity and, 100–103, 101f, 102f
 - efficient extraction and, 105–106, 106f
 - Hotelling Rule and, 108–110
 - impact of scarcity of, 238
 - marginal user cost and, 107–108
 - market power and, 110–111
 - overview of, 7, 113
 - overview of economics of, 99
 - property rights and, 111–113
 - two-period model of extraction of, 103–105, 104f
- Nonrivalry, 85–86, 87f
- Nonuniform mixing, 228–229
- Nonuse values, 48, 54
- Nordhaus, William, 59
- Normative decisions, discount rate as, 33
- North American Free Trade Agreement (NAFTA), 251
- Northern spotted owl, 65, 122–123, 124
- North Sea, 241–242, 248
- Northwest Economic Adjustment Initiative, 65
- Northwest Forest Plan, 65
- Norway, 241–242, 248
- Norwegian Central Bank, 242
- No Surprises Rule, 227
- NO_x. *See* Nitrogen oxides
- Nudge* (Thaler and Sunstein), 148
- Nudges, 147, 148
- Nutrient contamination, 221–223
- Obama, Barack, 208
- Offset credits, 208
- Ogallala Aquifer, 112
- Ohio River Basin, 192, 193
- Oil Pollution Act of 1990, 54
- Oil shocks, 110
- Oil spills, 46, 54
- OPEC. *See* Organization of the Petroleum Exporting Countries
- Open access, 130–132, 131f, 132–134
- Open-access resources, 93, 112
- Opower example, 148
- Opportunity costs, 15, 35–36, 119–120
- Optimal aging problem, 116
- Organization of the Petroleum Exporting Countries (OPEC), 110
- Orteig Prize, 185
- Owl, northern spotted, 65, 122–123, 124
- Ozone depletion, 28–29
- Pacific halibut, 159–160, 189
- Pareto, Vilfredo, 63
- Pareto efficiency, 63, 69, 244
- Parry, Ian, 153
- Partial equilibrium analyses, 39
- “Pay as You Throw,” 223–224
- Pennsylvania Nutrient Credit trading program, 221–223
- Perception, value and, 45–46
- Performance standards, 144, 173
- Perrier, 143
- Pervasiveness of markets, 3
- Perverse incentives, 151
- Petroleum Fund, 242
- Pigou, Arthur, 140, 246–247
- Pigouvian taxes, 140, 148–152, 160
- Planned economies, markets vs., 70
- PMC. *See* Private marginal costs
- Poaching, 126
- Policies. *See also* Market-based instruments; Nudges
 - to address market failure, 7–8
 - benefit–cost analysis and losers from, 61–63
 - criteria for comparing, 63–65
 - dynamic efficiency and, 30–34

- economic definitions of sustainability and, 243–245
- economic efficiency and, 20–22
- factors relevant to design of, 8
- implications of economics for, 254–256
- for promotion of innovation, 183–184
- Pollution, climate change and, 4
- Pollution, Property, and Prices* (Dales), 160
- Pollution havens hypothesis, 251
- Positive externalities, 82, 90–91
- Potential Pareto criterion, 64
- Poverty, protected areas and, 127
- Powder River Basin, 61–62, 204
- Power plants, 200–201. *See also* Sulfur dioxide abatement
- Preference, 49–55
- Prescriptive regulation. *See* Command-and-control
- Present value, 31–33, 33t, 34, 104
- Price-dumping, 213
- Pricing
- economic definitions of sustainability and, 243–244
 - incentives for wrong, 256
 - market-based policies to overcome market failure and, 147–152, 149f
 - municipal water supplies and, 215–216
 - setting quantities vs., 162–166
- Prisoner's Dilemma, 96–97, 97f
- Private bargaining, 140–143
- Private compliance costs, 36–37
- Private marginal costs (PMC), 83, 83f
- Prizes, for technological change, 185
- PROARBOL. *See* Federal Payments for Hydrological Services
- “The Problems of Social Cost” (Coase), 140
- Producer surplus, 76–77
- Property rights
- creating, 158–160
 - fisheries and, 158–160
 - forest resources and, 115, 123–128
 - for halibut in Gulf of Alaska, 159
 - nonrenewable resources and, 111–113
 - poorly defined, 113
 - water quality trading and, 223
- Protected areas, forest resources and, 126–127
- Public goods
- carbon emission abatement as, 4
 - defined, 80
 - as externalities, 90–91
 - free-riding problem and, 90
 - market failures and, 85–91, 87f, 88f, 89f
 - private provision of, 87–89, 88f, 89f
 - property rights, deforestation and, 123–128
 - public goods as, 90–91
- Quantities, setting pricing vs., 162–166
- Quotas. *See* Individual fishing quota markets
- Railroads, 63
- Rare earth elements, 110–111
- Rebound effects, 15
- Red-cockaded woodpecker, 20–21
- Regional Greenhouse Gas Initiative (RGGI), 209
- Regulatory Impact Analyses (RIA), 57, 58, 67
- Renewable resources, 7, 238. *See also* Fisheries; Forest resources
- Rents, 130–131, 236, 241–242. *See also* Marginal user costs
- Research and development, 184, 185
- Reserves-to-production ratio, 235–236, 237t
- Reserve-to-use ratios, 100
- Resource depletion, 248–249
- Resource rents. *See* Marginal user costs; Rents
- Revealed preference approaches to valuation, 49–52
- Revenues, market-based policies and, 160–162
- RGGI. *See* Regional Greenhouse Gas Initiative
- Ricardo, David, 233
- Right-to-know laws, overview of, 146–147
- Rockefeller, John D., 77
- Rotation, overview of, 115–116, 116f

- Safe Drinking Water Act, 55, 57
- Salmon, 2
- Scale, 41. *See also* Economies of scale
- Scarcity, 7, 101–103, 214–217. *See also*
Economic scarcity
- Scarcity rent. *See* Marginal user costs
- SCC. *See* Social cost of carbon
- Schaefer logistic model, fisheries and,
128, 129f
- Schnelling, Tom, 245
- Scientific Certification Systems, 146
- Scrubbers. *See* Sulfur dioxide abatement
- Seattle Stomp, 224
- Seawater, 240–241
- Seismic imaging, 102
- Shale revolution, 102
- Sierra Club, 256–257
- Simon, Julian, 234, 236
- Site value, 119–120
- Small, Kenneth, 153
- Smith, Adam, 69, 140
- Snake River, 2
- Social cost of carbon (SCC), 41–43, 43f
- Socialism, 70
- Social marginal costs, 83, 83f
- Social welfare costs
benefit–cost analysis and, 61–63
of carbon, 41–43, 43f
overview of, 36, 37–40
- Solow, Robert, 60, 110, 239–240
- South Africa, 210
- South Korea, 210
- SpaceShip One*, 185
- Species habitat value, 122, 123f
- Spillover effects, 127
- Spotted owl, northern, 65, 122–123, 124
- Standard Oil, 77
- Stated preference approaches to
valuation, 49, 52–55
- State-owned enterprises, 70
- Static efficiency rule, 26–27, 27f, 103
- Static reserve indices, 100
- Steady-state models, 129
- Stigler, George, 141
- Stiglitz, Joseph, 231
- Stock pollutants, 31
- Subsidies, 136–137, 145, 151
- Substitution, 237, 240–241
- Sulfur dioxide abatement
benefits of, 16–18, 17f
costs of, 14–16, 17f
economic efficiency and, 13,
18–20, 19f
market-based instruments for in
U.S., 200–207, 204t
overview of, 13–14
- Sunstein, Richard, 148
- Supply, 71–72, 72f, 75–76, 75f
- Supply curves, 72, 72f
- Surplus. *See* Consumer surplus;
Producer surplus
- Sustainability
defined, 232
economic growth and, 252
environmental policy and, 243–245
investment of natural resource rents
and, 241–242
overview of, 238–240
problems with economic definition
of, 242–243
substitution, technological change
and, 240–241
- Sustainable development, 246
- Symmetry, market efficiency and, 77
- TAC. *See* Total allowable catch
- Taxes, 149–152, 153, 215–216. *See also*
Specific taxes
- TCF methods. *See* Totally chlorine free
methods
- TDR. *See* Tradeable development rights
- Technical efficiency, 16, 19
- Technological change
cap-and-trade instruments and, 205
market-based instruments promoting,
179–184, 181f, 185, 189
scarcity and, 237
substitution, sustainability and,
240–241
- Technology standards, 144, 172–173,
183, 195
- Thaler, Richard, 148

- Thatcher, Margaret, 242
- TMDL program. *See* Total Maximum Daily Load program
- Tollbooths, 160
- Top-down cost models, 38–39
- Total allowable catch (TAC), 207–214
- Total benefits, marginal benefits and, 30
- Total costs, marginal costs and, 30
- Totally chlorine free (TCF) methods, 23
- Total Maximum Daily Load (TMDL) program, 221–222
- Toxics Release Inventory, 146
- Tradeable development rights (TDR), 225
- Trading programs. *See* Cap-and-trade systems
- Trading ratios, 192, 229
- Tragedy of the commons
 - as collective action problem, 94–97, 97f
 - defined, 80
 - as general model, 93–94
 - individual fishing quotas and, 159–160
 - market failures and, 91–97
 - overview of, 91–93
- Transaction costs, 142
- Transatlantic flights, 185
- Travel cost method, 49–50
- True costs, defined, 36
- Two-period model of nonrenewable resource extraction, 103–106, 106f
- Uncertainty, 65–68, 67f
- Uncompensated, externalities as, 81
- Uniform mixing assumption, 190
- Unintentional, externalities as, 81
- United Kingdom, 241–242
- United Nations Statistical Commission, 248
- Upper Ohio River Basin, 192, 193
- Use value, defined, 48
- Value of a statistical life (VSL), 51–52
- Values, 48, 56. *See also* Benefits
- Vittell, 143
- Volume-based waste disposal charge, 223–224
- VSL. *See* Value of a statistical life
- Waste management, 223–224
- Wasting Assets* (Repetto et al.), 248
- Water pollution, 143, 191–194
- Water pricing, 107–108, 214–217
- Water quality trading, 221–223
- Watershed protection, 122–123
- Waxman, Henry, 208
- Weak sustainability, 252
- Weitzman Rule, 165
- West Virginia, coal industry and, 61–62
- Wetland mitigation banking, 225–226, 228–229
- Wicksell Rule, 116–118, 117f
- Widodo, Joko, 151
- Willingness to accept (WTA), 46–47, 216
- Willingness to pay (WTP)
 - ability to pay and, 47
 - defined, 45–46
 - gross value vs. net benefit and, 11–12
 - overview of, 56, 216
 - willingness to accept vs., 46–47
- Woodpecker, red-cockaded, 20–21
- World Trade Organization (WTO), 251
- WTA. *See* Willingness to accept
- WTO. *See* World Trade Organization
- WTP. *See* Willingness to pay
- Zero pollution, 18–19
- Zoning regulations, 225

Island Press | Board of Directors

Katie Dolan

(Chair)

Environmental Writer

Merloyd Ludington Lawrence

(Secretary)

Merloyd Lawrence, Inc.
and Perseus Books

Pamela B. Murphy

(Vice-Chair)

William H. Meadows

(Treasurer)

Counselor and Past President
The Wilderness Society

Decker Anstrom

Board of Directors
Discovery Communications

Lisa A. Hook

President and CEO
Neustar Inc.

Stephen Badger

Board Member
Mars, Inc.

Mary James

Executive Committee Member
Prime Group, LLC

Terry Gamble Boyer

Author

Charles C. Savitt

President
Island Press

Paula A. Daniels

Founder
LA Food Policy Council

Alison Sant

Cofounder and Partner
Studio for Urban Projects

Melissa Shackleton Dann

Managing Director
Endurance Consulting

Ron Sims

Former Deputy Secretary
US Department of Housing
and Urban Development

Margot Paul Ernst

Anthony Everett

Alison Greenberg

Programme Officer
International Union
for the Conservation
of Nature

Sarah Slusser

Principal
emPower Partners, LLC

Deborah Wiley

Chair
Wiley Foundation, Inc.

About Island Press

Since 1984, the nonprofit organization Island Press has been stimulating, shaping, and communicating ideas that are essential for solving environmental problems worldwide. With more than 1,000 titles in print and some 30 new releases each year, we are the nation's leading publisher on environmental issues. We identify innovative thinkers and emerging trends in the environmental field. We work with world-renowned experts and authors to develop cross-disciplinary solutions to environmental challenges.

Island Press designs and executes educational campaigns in conjunction with our authors to communicate their critical messages in print, in person, and online using the latest technologies, innovative programs, and the media. Our goal is to reach targeted audiences—scientists, policymakers, environmental advocates, urban planners, the media, and concerned citizens—with information that can be used to create the framework for long-term ecological health and human well-being.

Island Press gratefully acknowledges major support of our work by The Agua Fund, The Andrew W. Mellon Foundation, The Bobolink Foundation, The Curtis and Edith Munson Foundation, Forrest C. and Frances H. Lattner Foundation, The JPB Foundation, The Kresge Foundation, The Oram Foundation, Inc., The Overbrook Foundation, The S.D. Bechtel, Jr. Foundation, The Summit Charitable Foundation, Inc., and many other generous supporters.

The opinions expressed in this book are those of the author(s) and do not necessarily reflect the views of our supporters.

Here is a thoroughly revised edition of an instant classic in environmental economics, with updated information on current environmental policy and real-world examples of market-based instruments. Key topics include cost-benefit analysis and economic efficiency, design and implementation of market-based policies, and economic growth and sustainability.

PRAISE:

“Two of the world’s leading young environmental economists offer a concise, balanced, and highly accessible treatment of one of the most dynamic fields of contemporary scholarship. Here is a book that will serve both as a superb text for an introductory course and an entertaining and effective introduction for a solo reader who wishes to learn about this growing field of inquiry.”

— **ROBERT N. STAVINS**, Albert Pratt Professor of Business and Government, John F. Kennedy School of Government, Harvard University and Director, Harvard Environmental Economics Program

“Succinct and well-written . . . Covers the material with admirable clarity, pertinent examples, and readily available analytics.”

— **BARRY C. FIELD**, author of *Environmental Economics: An Introduction* and *Natural Resource Economics: An Introduction*

“. . . a highly accessible and concise introduction to mainstream thinking in economics of the environment and resource management.”

— ***Environmental Conservation***

“. . . an excellent argument for why an economist deserves a seat at the environmental policy table.”

— ***Choice***

Markets and the Environment, Second Edition is part of the Foundations of Contemporary Environmental Studies Series produced by Island Press in collaboration with the Yale School of Forestry and Environmental Studies.

Nathaniel O. Keohane is Vice President at Environmental Defense Fund, where he leads EDF’s Global Climate program. **Sheila M. Olmstead** is Associate Professor of Public Affairs at the University of Texas at Austin.



Washington | Covelo | London
www.islandpress.org

All Island Press books are printed on recycled, acid-free paper.

Cover design: Amy Stirnkorb

Cover image: Joel Nakamura ©VEER

