Robert S. Pindyck, Daniel L. Rubinfeld (2017). Microeconomics, Global Edition (9th ed.). Pearson. ISBN: 978-1-292-213316.

1. Choose an output level represented by an isoquant in Figure 7.6 (a). Then find the point of tangency of that isoquant with an isocost line.
2. From the chosen isocost line, determine the minimum cost of producing the output level that has been selected.
3. Graph the output-cost combination in Figure 7.6 (b).

Suppose we begin with an output of 100 units. The point of tangency of the 100-unit isoquant with an isocost line is given by point *A* in Figure 7.6 (a). Because *A* lies on the $1000 isocost line, we know that the minimum cost of producing an output of 100 units in the long run is $1000. We graph this combination of 100 units of output and $1000 cost as point *D* in Figure 7.6 (b). Point *D* thus represents the $1000 cost of producing 100 units of output. Similarly, point *E* represents the $2000 cost of producing 200 units which corresponds to point *B* on the expansion path. Finally, point *F* represents the $3000 cost of 300 units corresponding to point *C*. Repeating these steps for every level of output gives the *long-run total cost curve* in Figure 7.6 (b)—i.e., the minimum long-run cost of producing each level of output.

In this particular example, the long-run total cost curve is a straight line. Why? Because there are constant returns to scale in production: As inputs increase proportionately, so do outputs. As we will see in the next section, the shape of the expansion path provides information about how costs change with the scale of the firm's operation.

EXAMPLE 7.5 REDUCING THE USE OF ENERGY

B

Policymakers around the world have been concerned with finding ways to reduce the use of energy. In part, this reflects environmental concerns—most energy consumption uses fossil fuels and thus contributes to the emission of greenhouse gases and global warming. But energy, whether in the form of oil, natural gas, coal or nuclear, is also expensive, so if companies can find ways to reduce their energy use, they can lower their costs.

There are essentially two ways that companies can reduce the amount of energy they use. The first is to substitute other factors of production for energy. For example, some machines might be more costly but also use less energy, so if energy prices rise, firms could respond by buying and using those energy-efficient machines, effectively substituting capital for energy. This is exactly what has happened as energy prices rose in recent years: firms bought and installed expensive but more energy-efficient heating and cooling systems, industrial processing equipment, trucks, cars, and other vehicles.

The second way to reduce energy use is through technological change. As time passes, research and development lead to innovations that make it possible to produce the same output using fewer inputs—less labor, less capital, and less energy. Thus even if the relative prices of energy and capital stay the same, firms will use less energy (and less capital) to produce the same output. Advances in robotics during the past two decades are an example of this; cars and trucks are now produced with less capital and energy (as well as less labor).

These two ways of reducing energy use are illustrated in Figures 7.7 (a) and (b), which show how capital and energy are combined to produce output.[8] The isoquants in each figure represent the various combinations of capital and energy that can be used to generate the same level of output. The figures illustrate how reductions in energy use can be achieved in two ways. First, firms can substitute more capital for energy, perhaps in response to a government subsidy for investment in energy-saving equipment and/or an increase in the cost of electricity. This is shown as a movement along isoquant q_1 from point *A* to point *B* in Figure 7.7(a), with capital increasing

[8]This example was inspired by Kenneth Gillingham, Richard G. Newell, and Karen Palmer, "Energy Efficiency Economics and Policy," *Annual Review of Resource Economics*, 2009, Vol. 1: 597–619.

Still another piece of good news is that from 2005 onward, the use of coal to generate electricity in the U.S. declined, replaced largely by natural gas, as Figure 18.9 illustrates. Why is this good news? Because CO_2 emissions from burning natural gas are less than half the emissions from burning coal. Thus the shift from coal to natural gas is an important step in reducing CO_2 emissions and global warming.

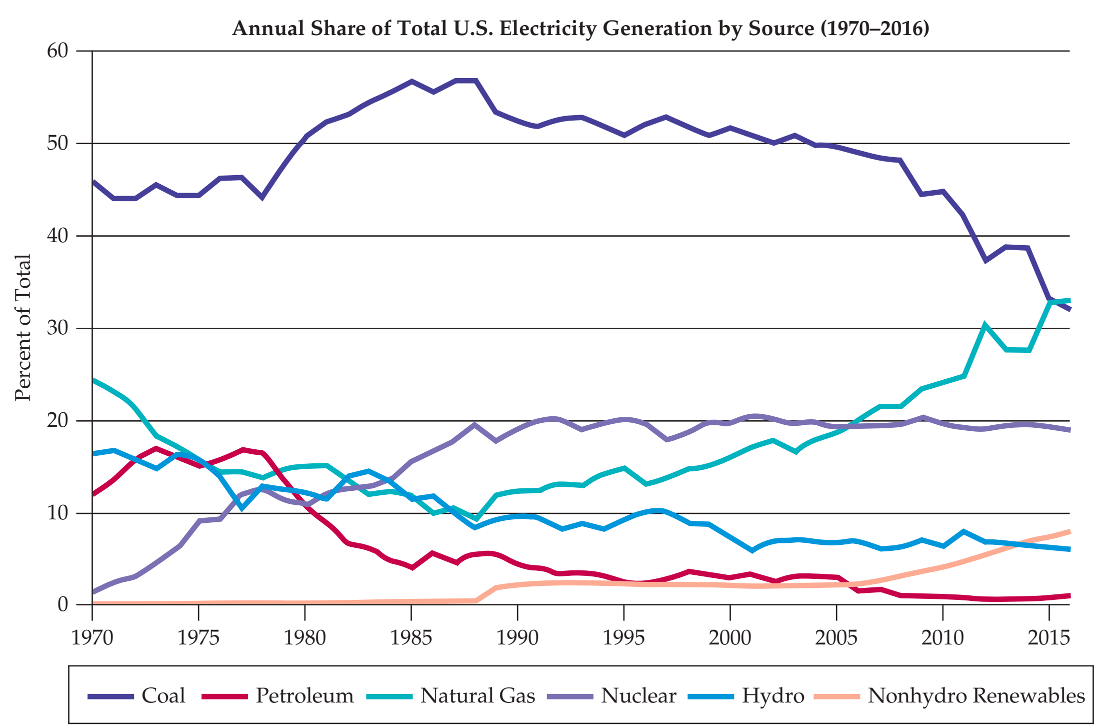

B

FIGURE 18.9
FUEL MIX FOR U.S. ELECTRICITY GENERATION

Until recently, coal had been the primary fuel for generating electricity, accounting for over 50% of electricity generation during the 1980s and 1990s. But starting around 2000, coal has been increasingly displaced by natural gas.

Recycling

To the extent that the disposal of waste products involves little or no private cost to either consumers or producers, society will dispose of too much waste material. The overutilization of virgin materials and the underutilization of recycled materials will result in a market failure that may require government intervention. Fortunately, given the appropriate incentive to recycle products, this market failure can be corrected.[11]

To see how recycling incentives can work, consider a typical household's decision with respect to the disposal of glass containers. In many communities, households are charged a fixed annual fee for trash disposal. As a result, these

[11]Even without market intervention, some recycling will occur if the price of virgin material is sufficiently high. For example, recall from Chapter 2 that when the price of copper is high, there is more recycling of scrap copper.

Sometimes, however, the damage to society comes not directly from the emissions flow, but rather from the *accumulated stock* of the pollutant. A good example is global warming. Global warming is thought to result from the accumulation of carbon dioxide and other greenhouse gasses (GHGs) in the atmosphere. (As the GHG concentration grows, more sunlight is absorbed into the atmosphere rather than being reflected away, causing an increase in average temperatures.) GHG emissions do not cause the kind of immediate harm that sulfur dioxide emissions cause. Rather, it is the *stock of accumulated GHGs in the atmosphere* that ultimately causes harm. Furthermore, the *dissipation rate* for accumulated GHGs is very low: Once the GHG concentration in the atmosphere has increased substantially, it will remain high for many years, even if further GHG emissions were reduced to zero. That is why there is concern about reducing GHG emissions now rather than waiting for concentrations to build up (and temperatures to start rising) fifty or more years from now.

stock externality
Accumulated result of action by a producer or consumer which, though not accounted for in the market price, affects other producers or consumers.

Stock externalities (like flow externalities) can also be positive. An example is the stock of "knowledge" that accumulates as a result of investments in R&D. Over time, R&D leads to new ideas, new products, more efficient production techniques, and other innovations that benefit society as a whole, and not just those who undertake the R&D. Because of this positive externality, there is a strong argument for the government to subsidize R&D. Keep in mind, however, that it is the *stock* of knowledge and innovations that benefits society, and not the flow of R&D that creates the stock.

Recall from §15.1 that a firm's capital is measured as a stock, while the investment that creates the capital is a flow. The firm's output is also measured as a flow.

Recall from §15.2 that the present discounted value (PDV) of a series of expected future cash flows is the sum of those cash flows discounted by the appropriate interest rate. Moreover, we observe in §15.4 that, according to the net present value (NPV) rule, a firm should invest if the PDV of the expected future cash flow from an investment is greater than the cost.

We examined the distinction between a stock and a flow in Chapter 15. As we explained in Section 15.1 (page 574), the capital that a firm owns is measured as a *stock*, i.e., as a quantity of plant and equipment that the firm owns. The firm can increase its stock of capital by purchasing additional plant and equipment, i.e., by generating a *flow* of investment expenditures. (Recall that inputs of labor and raw materials are also measured as *flows*, as is the firm's output.) We saw that this distinction is important, because it helps the firm decide whether to invest in a new factory, equipment, or other capital. By comparing the *present discounted value (PDV)* of the additional profits likely to result from the investment to the cost of the investment, i.e., by calculating the investment's *net present value (NPV)*, the firm can decide whether or not the investment is economically justified.

The same net present value concept applies when we want to analyze how the government should respond to a stock externality—though with an additional complication. For the case of pollution, we must determine how any ongoing level of emissions leads to a buildup of the stock of pollutant, and we must then determine the economic damage likely to result from that higher stock. We will then be able to compare the present value of the ongoing costs of reducing emissions each year to the present value of the economic benefits resulting from a reduced future stock of the pollutant.

Stock Buildup and Its Impact

Let's focus on pollution to see how the stock of a pollutant changes over time. With ongoing emissions, the stock will accumulate, but some fraction of the stock, δ, will dissipate each year. Thus, assuming the stock starts at zero, in the first year, the stock of pollutant (S) will be just the amount of that year's emissions (E):

A

$$S_1 = E_1$$

In the second year, the stock of pollutant will equal the emissions that year plus the nondissipated stock from the first year—

$$S_2 = E_2 + (1 - \delta)S_1$$

—and so on. In general, the stock in any year t is given by the emissions generated that year plus the nondissipated stock from the previous year:

$$S_t = E_t + (1 - \delta)S_{t-1}$$

If emissions are at a constant annual rate E, then after N years, the stock of pollutant will be[14]:

$$S_N = E[1 + (1 - \delta) + (1 - \delta)^2 + \cdots + (1 - \delta)^{N-1}]$$

As N becomes infinitely large, the stock will approach the long-run equilibrium level E/δ.

The impact of pollution results from the accumulating stock. Initially, when the stock is small, the economic impact is small; but the impact grows as the stock grows. With global warming, for example, higher temperatures result from higher concentrations of GHGs: thus the concern that if GHG emissions continue at current rates, the atmospheric stock of GHGs will eventually become large enough to cause substantial temperature increases—which, in turn, could have adverse effects on weather patterns, agriculture, and living conditions. Depending on the cost of reducing GHG emissions and the future benefits of averting these temperature increases, it may make sense for governments to adopt policies that would reduce emissions now, rather than waiting for the atmospheric stock of GHGs to become much larger.

A

NUMERICAL EXAMPLE We can make this concept more concrete with a simple example. Suppose that, absent government intervention, 100 units of a pollutant will be emitted into the atmosphere every year for the next 100 years; the rate at which the stock dissipates, δ, is 2 percent per year, and the stock of pollutant is initially zero. Table 18.1 shows how the stock builds up over time. Note that after 100 years, the stock will reach a level of 4,337 units.

TABLE 18.1 BUILDUP IN THE STOCK OF POLLUTANT

YEAR	E	S_t	DAMAGE ($ BILLION)	COST OF $E = 0$ ($ BILLION)	NET BENEFIT ($ BILLION)
2010	100	100	0.100	1.5	−1.400
2011	100	198	0.198	1.5	−1.302
2012	100	296	0.296	1.5	−1.204
…	…	…	…	…	…
2110	100	4,337	4.337	1.5	2.837
…	…	…	…	…	…
∞	100	5,000	5.000	1.5	3.500

[14] To see this, note that after 1 year, the stock of pollutant is $S_1 = E$, in the second year the stock is $S_2 = E + (1 - \delta)S_1 = E + (1 - \delta)E$, in the third year, the stock is $S_3 = E + (1 - \delta)S_2 = E + (1 - \delta)E + (1 - \delta)^2E$, and so on. As N becomes infinitely large, the stock approaches E/δ.

(If this level of emissions continued forever, the stock will eventually approach $E/\delta = 100/.02 = 5{,}000$ units.)

Suppose that the stock of pollutant creates economic damage (in terms of health costs, reduced productivity, etc.) equal to $1 million per unit. Thus, if the total stock of pollutant were, say, 1000 units, the resulting economic damage for that year would be $1 billion. And suppose that the annual cost of reducing emissions is $15 million per unit of reduction. Thus, to reduce emissions from 100 units per year to zero would cost 100 × $15 million = $1.5 billion per year. Would it make sense, in this case, to reduce emissions to zero starting immediately?

To answer this question, we must compare the present value of the annual cost of $1.5 billion with the present value of the annual benefit resulting from a reduced stock of pollutant. Of course, if emissions were reduced to zero starting immediately, the stock of pollutant would likewise be equal to zero over the entire 100 years. Thus, the benefit of the policy would be the savings of social cost associated with a growing stock of pollutant. Table 18.1 shows the annual cost of reducing emissions from 100 units to zero, the annual benefit from averting damage, and the annual *net* benefit (the annual benefit net of the cost of eliminating emissions). As you would expect, the annual net benefit is negative in the early years because the stock of pollutant is low; the net benefit becomes positive only later, after the stock of pollutant has grown.

A

To determine whether a policy of zero emissions makes sense, we must calculate the NPV of the policy, which in this case is the present discounted value of the annual net benefits shown in Table 18.1. Denoting the discount rate by R, the NPV is:

$$\text{NPV} = (-1.5 + .1) + \frac{(-1.5 + .198)}{1 + R} + \frac{(-1.5 + .296)}{(1 + R)^2} + \cdots + \frac{(-1.5 + 4.337)}{(1 + R)^{99}}$$

Recall from §15.1 that the NPV of an investment declines as the discount rate becomes larger. Figure 15.3 shows the *NPV* for an electric motor factory; note the similarity to our environmental policy problem.

Is this NPV positive or negative? The answer depends on the discount rate, R. Table 18.2 shows the NPV as a function of the discount rate. (The middle row of Table 18.2, in which the dissipation rate δ is 2 percent, corresponds to Table 18.1. Table 18.2 also shows NPVs for dissipation rates of 1 percent and 4 percent.) For discount rates of 4 percent or less, the NPV is clearly positive, but if the discount rate is large, the NPV will be negative.

Table 18.2 also shows how the NPV of a "zero emissions" policy depends on the dissipation rate, δ. If δ is lower, the accumulated stock of pollutant will reach higher levels and cause more economic damage, so the future benefits of reducing emissions will be greater. Note from Table 18.2 that for any given

TABLE 18.2 NPV OF "ZERO EMISSIONS" POLICY

		DISCOUNT RATE, R				
		.01	.02	.04	.06	.08
DISSIPATION RATE, δ	.01	108.81	54.07	12.20	−0.03	−4.08
	.02	65.93	31.20	4.49	−3.25	−5.69
	.04	15.48	3.26	−5.70	−7.82	−8.11

Note: Entries in table are NPVs in $billions. Entries for $\delta = .02$ correspond to net benefit numbers in Table 18.1.

discount rate, the NPV of eliminating emissions is much larger if $\delta = .01$ and much smaller if $\delta = .04$. As we will see, one of the reasons why there is so much concern over global warming is the fact that the stock of GHGs dissipates very slowly; δ is only about .005.

Formulating environmental policy in the presence of stock externalities therefore introduces an additional complicating factor: What discount rate should be used? Because the costs and benefits of a policy apply to society as a whole, the discount rate should likewise reflect the opportunity cost to society of receiving an economic benefit in the future rather than today. This opportunity cost, which should be used to calculate NPVs for government projects, is called the **social rate of discount**. But as we will see in Example 18.4, there is little agreement among economists as to the appropriate number to use for the social rate of discount.

social rate of discount Opportunity cost to society as a whole of receiving an economic benefit in the future rather than the present.

In principle, the social rate of discount depends on three factors: (1) the expected rate of real economic growth; (2) the extent of risk aversion for society as a whole; and (3) the "rate of pure time preference" for society as a whole. With rapid economic growth, future generations will have higher incomes than current generations, and if their marginal utility of income is decreasing (i.e., they are risk-averse), their utility from an extra dollar of income will be lower than the utility to someone living today; that's why future benefits provide less utility and should thus be discounted. In addition, even if we expected no economic growth, people may simply prefer to receive a benefit today than in the future (the rate of pure time preference). Depending on one's beliefs about future real economic growth, the extent of risk aversion for society as a whole, and the rate of pure time preference, one could conclude that the social rate of discount should be as high as 6 percent—or as low as 1 percent. And herein lies the difficulty. With a discount rate of 6 percent, it is hard to justify almost any government policy that imposes costs today but yields benefits only 50 or 100 years in the future (e.g., a policy to deal with global warming). Not so, however, if the discount rate is only 1 or 2 percent.[15] Thus for problems involving long time horizons, the policy debate often boils down to a debate over the correct discount rate.

EXAMPLE 18.4 GLOBAL WARMING

Emissions of carbon dioxide and other greenhouse gases have increased dramatically over the past century as economic growth has been accompanied by the greater use of fossil fuels, which has in turn led to an increase in atmospheric concentrations of GHGs. Even if worldwide GHG emissions were to be stabilized at current levels, atmospheric GHG concentrations would continue to grow throughout the next century. By trapping sunlight, these higher GHG concentrations are likely to cause a significant increase in global mean temperatures in 50 years or so and could have severe environmental consequences—flooding of low-lying areas as the polar ice caps melt and sea levels rise, more extreme weather patterns, disruption of ecosystems, and reduced agricultural output. GHG emissions could be reduced from their current levels—governments, for example,

[15]For example, with a discount rate of 6 percent, \$100 received 100 years from now is worth only \$0.29 today. With a discount rate of 1 percent, that same \$100 is worth \$36.97 today, i.e., 127 times as much.

could impose stiff taxes on the use of gasoline and other fossil fuels—but this solution would be costly. The problem is that the costs of reducing GHG emissions would occur today, but the benefits from reduced emissions would be realized only in some 50 or more years. Should the world's industrialized countries agree to adopt policies to dramatically reduce GHG emissions, or is the present discounted value of the likely benefits of such policies simply too small?

Many climate scientists and economists have studied the probable build-up of GHG concentrations and resulting increases in global temperatures if no steps are taken to reduce emissions. Although there is considerable uncertainty over the economic impact of higher temperatures, the consensus view is that the impact could be significant, so that there would be a future benefit from reducing emissions today.[16] The cost of reducing emissions (or preventing them from growing above current levels) can be assessed as well, although here too there is uncertainty over the specific numbers.

Table 18.3 shows GHG emissions and average global temperature change at ten-year intervals for two scenarios, starting in 2020. The first is a "business as usual" scenario in which GHG emissions are projected to more than double over the next century so that the average GHG concentration rises considerably, and by 2120 the average temperature is 4 degrees Celsius above its current level. The resulting damage each year from this temperature increase is estimated to be 1.5 percent of world GDP per degree Celsius of temperature increase. World GDP is assumed to grow at 2.5 percent per year in real terms from its 2016 value of $74 trillion, reaching $965 trillion in 2120. Thus the annual damage from global warming reaches about $(.015)(4)(965) = \$57.9$ trillion in 2120.

The second scenario shown in Table 18.3 is one in which the GHG concentration is stabilized at 550 ppm

TABLE 18.3 REDUCING GHG EMISSIONS

	"BUSINESS AS USUAL"				EMISSIONS REDUCED BY 1% PER YEAR					
YEAR	E_t	S_t	ΔT_t	DAMAGE	E_t	S_t	ΔT_t	DAMAGE	COST	NET BENEFIT
2020	55	460	0°	0	45	460	0°	0	0.82	−0.82
2030	62	490	0.4°	0.63	41	485	0.4°	0.63	1.05	−1.05
2040	73	520	0.8°	1.61	37	510	0.8°	1.61	1.34	−1.34
2050	85	550	1.2°	3.08	33	530	1.2°	3.08	1.71	−1.71
2060	90	580	1.6°	5.26	30	550	1.6°	5.26	2.19	−2.19
2070	95	610	2°	8.42	27	550	2°	8.42	2.81	−2.81
2080	100	640	2.4°	12.94	25	550	2°	10.78	3.59	−1.44
2090	105	670	2.8°	19.32	22	550	2°	13.80	4.60	0.92
2100	110	700	3.2°	28.27	20	550	2°	17.67	5.89	4.71
2110	115	730	3.6°	40.71	18	550	2°	22.61	7.54	10.55
2120	120	760	4°	57.90	16	550	2°	28.95	9.65	19.30

Notes: E_t is measured in gigatonnes (billions of metric tons) of CO_2 equivalent (CO_2e), S_t is measured in parts per million (ppm) of atmospheric CO_2e, the change in temperature ΔT is measured in degrees Celsius, and costs, damages, and net benefits are measured in trillions of 2007 dollars. Cost of reducing emissions is estimated to be 1 percent of GDP each year. World GDP is projected to grow at 2.5% in real terms from a level of $74 trillion in 2016. Damage from warming is estimated to be 1.5% of GDP per year for every 1°C of temperature increase. Under BAU, temperature is predicted to rise by 0.04° per year.

[16]For a consensus view, see the 2007 *Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge University Press or online at **http://www.ipcc.ch**.

so that the temperature increase is limited to only 2 degrees Celsius, which is reached in 2070. To achieve this, GHG emissions be reduced by 1 percent per year starting in 2010. The annual cost of this emissions reduction policy is estimated to be 1 percent of world GDP.[17] (Because world GDP is assumed to increase each year, so too does the cost of this policy.) Also shown in the table is the annual net benefit from the policy, which equals the damage under the "business as usual" scenario minus the (smaller) damage when emissions are reduced minus the cost of reducing emissions.

A

Does this emissions-reduction policy make sense? To answer that question, we must calculate the present value of the flow of net benefits, which depends critically on the discount rate. A review conducted in the United Kingdom recommends a social rate of discount of 1.3 percent. With that discount rate, the NPV of the policy is $11.41 trillion, which shows that the emissions-reduction policy is clearly economical. However, if the discount rate is 2 percent, the NPV drops to −$12.19 trillion, and with a discount rate of 3 percent, the NPV is −$23.68 trillion.

We have examined a particular policy—and a rather stringent one at that—to reduce GHG emissions. Whether that policy or any other policy to restrict GHG emissions makes economic sense clearly depends on the rate used to discount future costs and benefits. Be warned, however, that economists disagree about what rate to use, and as a result, they disagree about what should be done about global warming.[18]

18.4 Externalities and Property Rights

We have seen how government regulation can deal with the inefficiencies that arise from externalities. Emissions fees and transferable emissions permits work because they change a firm's incentives, forcing it to take into account the external costs that it imposes. But government regulation is not the only way to deal with externalities. In this section we show that in some circumstances, inefficiencies can be eliminated through private bargaining among the affected parties or by a legal system in which parties can sue to recover the damages they suffer.

Property Rights

Property rights are the legal rules that describe what people or firms may do with their property. If you have property rights to land, for example, you may build on it or sell it and are protected from interference by others.

property rights Legal rules stating what people or firms may do with their property.

To see why property rights are important, let's return to our example of the firm that dumps effluent into the river. We assumed both that it had a property right to use the river to dispose of its waste and that the fishermen did not have

[17]This policy is the one recommended by the Stern Review, commissioned by the U.K. Government, and available online at **http://www.hm-treasury.gov.uk/stern_review_report.htm**. The cost estimate of 1 percent of GDP is from the Stern Review, and is probably too optimistic. The estimate of the damage from higher temperatures (1.3 percent of GDP for each 1 degree Celsius increase) is an amalgam of estimates from the Stern Review and the IPCC Report.

[18]This disagreement over the discount rate and its crucial role in assessing policies to reduce GHG emissions is spelled out quite nicely in Martin Weitzman, "The Stern Review of the Economics of Climate Change," *Journal of Economic Literature* (September 2007). Also, there are many uncertainties about the size of possible future temperature increases and their social and economic impact. Those uncertainties can have implications for policy but have been ignored in this example. See, for example, R. S. Pindyck, "Uncertainty in Environmental Economics," *Journal of Environmental Economics and Policy* (Winter 2007), R. S. Pindyck, "Uncertain Outcomes and Climate Change Policy," *Journal of Environmental Economics and Management*, 2012. A

19.5 Behavioral Economics and Public Policy

When economists design public policies, they often assume that the consumers and firms affected by those policies are fully rational and, more importantly, fully informed. An example is the use of a tax to reduce or eliminate a negative externality, such as pollution. Go back to Figure 18.1 in the previous chapter. Part (b) of that figure describes an industry where the firms emit a pollutant roughly proportionally to their output. As a result there is a marginal external cost (MEC), which in the figure rises linearly with industry output. Thus the marginal social cost (MSC) of industry output is greater than the marginal private cost (MC), which defines the industry supply curve. If unregulated, industry output, Q_1, is greater than the socially optimal output, Q^*. As explained in Chapter 18, the public policy solution might be a tax on output that makes the marginal private cost equal to the marginal social cost, so that output is reduced to the socially optimal level Q^*.

But perhaps there is another way to think about—and respond to—this pollution problem. Suppose the pollutant at issue is carbon dioxide, CO_2, which creates an external cost because it contributes to global warming and climate change. (Recall our analysis of global warming in Example 18.4.) The typical response suggested by policy analysts is a *carbon tax*, which would raise the marginal private cost of burning fossil fuels and thereby reduce CO_2 emissions. This assumes, however, that firms and consumers are fully informed about their own private costs of burning fossil fuels. But this assumption might be incorrect.

It might well be that if they were properly informed, consumers and firms, on their own and without the incentives of a tax, would reduce their use of fossil fuels. Why might this be the case? Consider the option of replacing incandescent light bulbs with LED bulbs, which are far more energy efficient. (A 12-watt LED bulb, for example, will produce the same amount of light as a 100-watt incandescent bulb.) While LED bulbs are more expensive than incandescent bulbs, the savings in electricity usage are so large that the LED bulbs will usually pay for themselves in a year or two and thus save money. So, with almost no effort and little or no cost, consumers (and firms) can reduce energy consumption (and thus reduce CO_2 emissions by electric power plants). They just need to be properly informed in order to be incentivized to change their light bulbs. Likewise, energy consumption can be reduced if consumers are better informed about the cost savings of better insulation, "smart" thermostats, and so on.

This is where behavioral economics comes into play in the design of public policy. If the policy objective is to reduce energy use, we need to understand how *people's behavior* affects the decisions they make regarding energy. If consumers were fully rational and utility-maximizing, they would make the effort on their own to learn about the cost savings from LED light bulbs and then act accordingly by switching to LED bulbs. But that's asking too much of consumers. Instead, one element of public policy would be to educate consumers about LED bulbs, perhaps through paid advertisements or even classes on "home economics and finance" in schools and colleges.

The behavioral approach to public policy is illustrated in Figure 19.6, which generalizes Figure 18.1.[23] In this figure, an industry emits a pollutant so there is a

[23] A very similar and more detailed example of a behavioral approach to public policy, but based on a positive externality (vaccinations), is in Brigitte C. Madrian, "Applying Insights from Behavioral Economics to Policy Design," *Annual Review of Economics* 6 (2014): 663–88. Figure 19.6 is based on Figure 1 from that article. See also Allison Demeritt and Karla Hoff, "Small Miracles – Behavioral Insights to Improve Development Policy," Policy Research Working Paper 7197, World Bank, 2015.

A

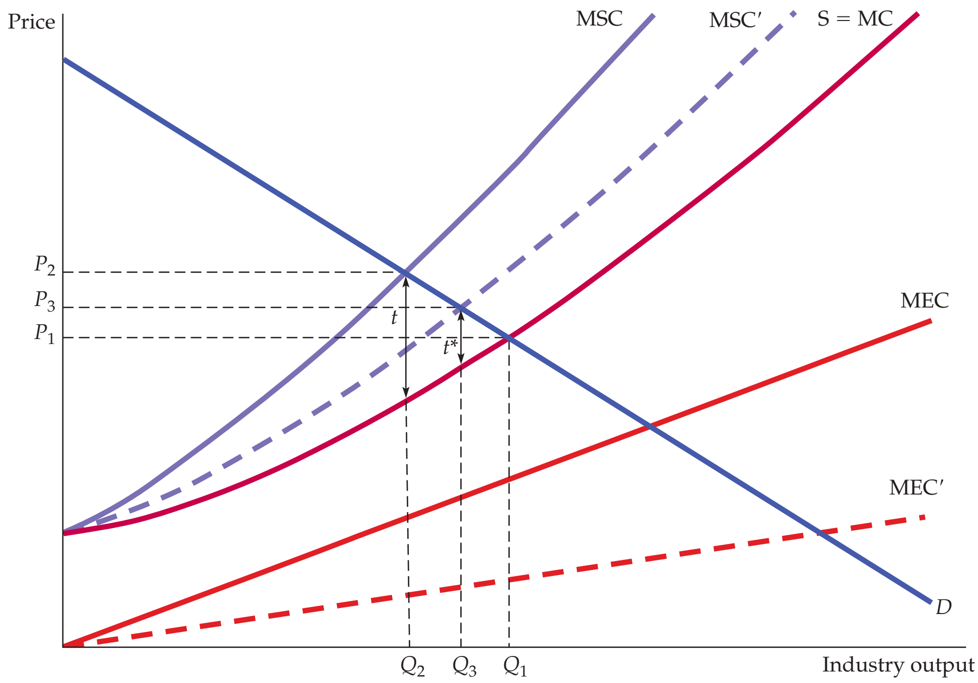

FIGURE 19.6
EXTERNAL COST—A BEHAVIORAL ANALYSIS

With pollution emissions, the marginal social cost of industry output is greater than the marginal private cost and industry output is greater than the socially optimal level.

marginal external cost of industry output, given by the curve MEC. The marginal social cost of industry output, MSC, is the sum of the marginal private cost (MC) and the marginal external cost. Industry output is therefore too large: Q_1 instead of the socially optimal output Q_2. One solution to this problem would be to impose a tax, t, which would equate the marginal social and private costs. But suppose that with a bit of education, consumers and firms would realize that they can save money by reducing their emissions of the pollutant. That would lower the marginal external cost curve (from MEC to MEC′ in the figure) and likewise lower the marginal social cost curve (to MSC′). Output might still be too large (in the figure, Q_3 instead of Q_2), but a much smaller tax (t^* instead of t) would be needed to correct the problem.

In the case of CO_2 emissions and climate change, there are other ways that consumers and firms might be induced to reduce their energy consumption. One example would be *moral persuasion*. If consumers were convinced that they have a moral obligation to conserve energy (even if doing so was inconvenient or costly), they might indeed reduce their energy consumption. You might say that if they do so, they are not maximizing utility subject to a budget constraint, along the lines of Chapters 3 and 4. But that's short-sighted. Moral obligations can certainly enter people's utility functions, so that reducing energy consumption would indeed be utility-maximizing.

Other policy options might alter (and hopefully improve) the environment in which individuals make choices. For example, one option might be to change the default option available to consumers. In the light bulb example, local ordinances might require stores to make incandescent bulbs available only upon request. Another option might be for stores to be encouraged or even required to feature LED bulbs prominently but to make incandescent bulbs less visible. Finally, stores might be required to simplify consumer choices by displaying signage that clearly describes the advantages of LED bulbs over incandescent bulbs. All of these policy options change the environment in which consumers make choices. Giving consumers a "nudge" in the direction of choices that they would likely make if fully informed can help them to maximize utility.[24]

Summing Up

Where does this leave us? Should we dispense with the traditional consumer theory discussed in Chapters 3 and 4? Not at all. In fact, the basic theory that we learned up to now works quite well in many situations. It helps us to understand and evaluate the characteristics of consumer demand and to predict the impact on demand of changes in prices or incomes. Although it does not explain all consumer decisions, it sheds light on many of them. The developing field of behavioral economics tries to explain and to elaborate on those situations that are not well explained by the basic consumer model.

If you continue to study economics, you will notice many cases in which economic models are not a perfect reflection of reality. Economists have to carefully decide, on a case-by-case basis, what features of the real world to include and what simplifying assumptions to make so that models are neither too complicated to study nor too simple to be useful.

---

[24]See Cass R. Sunstein and Richard H. Thaler, *Nudge: Improving Decisions about Health, Wealth, and Happiness*, Yale University Press (2008).

SUMMARY

1. Individual behavior sometimes seems unpredictable, even irrational, and contrary to the assumptions that underlie the basic model of consumer choice. The study of behavioral economics enriches consumer theory by accounting for reference points, endowment effects, anchoring, fairness considerations, and deviations from the laws of probability.
2. **Reference points**—the points from which individuals make consumption decisions—can strongly affect the way individuals approach economic decisions. An example of a reference point is the **endowment effect**—the effect of owning a good on the value that individuals place on the good.
3. Other psychological perspectives serve to explain the impact of reference points. They include loss aversion (the tendency of individuals to prefer avoiding losses over acquiring gains), **framing** (a tendency to rely on the context in which a choice is described), and **salience** (the perceived importance of one or more of a good's features).
4. Models of consumer theory can be modified to take into account the effects of fairness on consumer behavior. To illustrate, concerns about the fairness of the pricing of a good can be seen as a shift in the demand curve to one that is very elastic at relatively high prices.
5. Individuals often resort to **rules of thumb**—mental shortcuts—in making decisions. However, rules of thumb can lead to biases in decision making, as illustrated by the **law of small numbers,** whereby individuals tend to overstate the probability that certain events will occur when faced with little information from recent memory.
6. **Overconfidence** (overestimating one's prospects or abilities) is a common bias in human decision-making. It can take the form of over-precision, whereby individuals have an unrealistic belief that they can accurately predict outcomes.
7. A **bubble** is an increase in the price of a good that is not based on the fundamentals of demand. Bubbles are often the result of an irrational belief that the price of a good will keep going up. A bubble can result from an **informational cascade,** when one's action depends on the actions of others, which in turn were based on the actions of still others, and so on.
8. An understanding of behavioral economics can help economists design appropriate public policies, as illustrated by the design of tax policies that reduce or eliminate externalities.

EXAMPLE 4.5 THE LONG-RUN DEMAND FOR GASOLINE

Among industrialized countries, the United States is unique in that the price of gasoline is relatively low. The reason is simple: Europe, Japan, and other countries have stiff taxes on gasoline, so that gas prices are typically double or triple that in the United States, which imposes very low taxes on gasoline. Many economists have argued that the United States should substantially increase its tax on gasoline, because doing so would lower gasoline consumption and thereby reduce dependence on imported oil and reduce the greenhouse gas emissions that contribute to global warming (in addition to providing much-needed revenue to the government). Politicians have resisted, however, because they fear that a tax increase would anger voters.

Putting the politics of a gas tax aside, would higher gasoline prices indeed reduce gasoline consumption, or are drivers so wedded to big gas-guzzling cars that higher prices would make little difference? What matters here is the *long-run* demand for gasoline, because we can't expect drivers to immediately scrap their old cars and buy new ones following a price increase. One way to get at the long-run demand curve is by looking at per-capital consumption of gasoline in different countries which historically have had very different prices (because they imposed different gasoline taxes). Figure 4.13 does just that. It plots the per-capita consumption of gasoline on the vertical axis and the price in dollars per gallon for 10 countries on the horizontal axis.[6] (Each circle represents the population of the corresponding country.)

Note that the United States has had by far the lowest gasoline prices and also the highest per-capita gasoline consumption. Australia is roughly in the middle in terms of prices, and likewise in terms of consumption. Most of the European countries, on the other hand, have much higher prices and correspondingly lower per capita consumption levels. The long-run elasticity of demand for gasoline turns out to be about −1.4.

Now we come back to our question: Would higher gasoline prices reduce gasoline consumption? Figure 4.13 provides a clear answer: Most definitely.

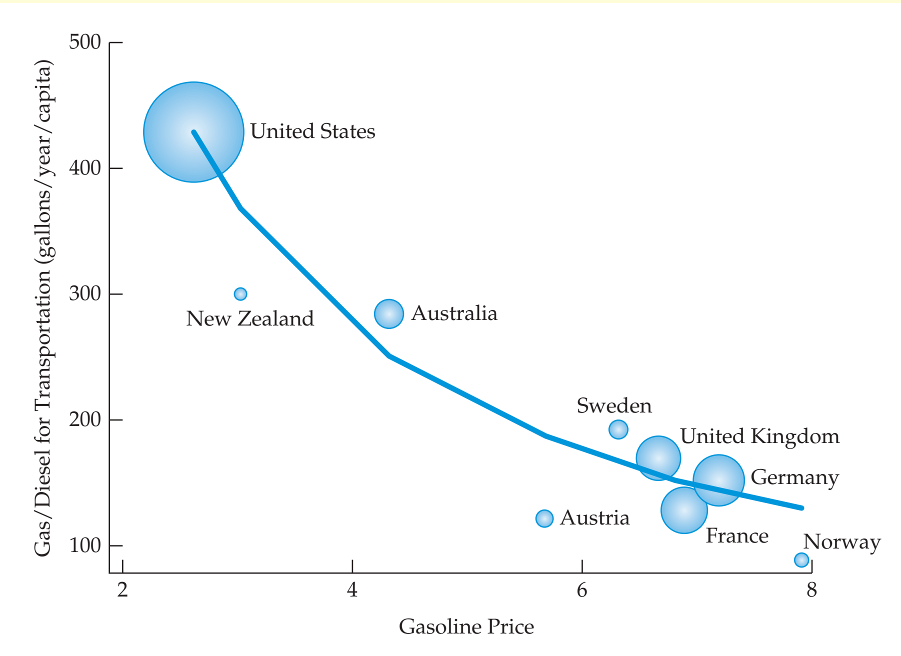

FIGURE 4.13 GASOLINE PRICES AND PER CAPITA CONSUMPTION IN 10 COUNTRIES

The graph plots per capita consumption of gasoline versus the price per gallon (converted to U.S. dollars) for 10 countries over the period 2008 to 2010. Each circle represents the population of the corresponding country.

[6]Our thanks to Chris Knittel for providing us with the data for this figure. The figure controls for income differences and is based on Figure 1 in Christopher Knittel, "Reducing Petroleum Consumption from Transportation," *Journal of Economic Perspectives*, 2012. All underlying data are available from **www.worldbank.org**.