



UNIVERSITY OF TORONTO PRESS

Advancing Knowledge

Chapter Title: Environmental Economics

Book Title: Elements of Environmental Management

Book Author(s): Werner Antweiler

Published by: University of Toronto Press

Stable URL: <https://www.jstor.org/stable/10.3138/j.ctt7zwbx7.9>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



JSTOR

University of Toronto Press is collaborating with JSTOR to digitize, preserve and extend access to Elements of Environmental Management

Chapter 3

Environmental Economics

Protecting the environment is embraced widely, but businesses and society often differ about the scope of protection. How much environmental protection is too little, and how much protection is too much? There is often no scientific answer to that question. Ultimately, it comes down to a trade-off where the costs of environmental protection (or pollution abatement) are weighed against the benefits of a cleaner (less polluted) environment. Often an important consideration is whether environmental harm is reversible or irreversible. With reversible harm, economic policy does not necessarily have to 'get it right' on the first try; but when harm is irreversible, policies that are insufficiently stringent will lock in the damage permanently. Environmental economics is a useful tool for designing environmental policies. In essence, environmental economics rests on the assumption that there are market failures for environmental goods (such as air and water), and that policy interventions (ideally using market and price mechanisms) can be used to correct these market failures. This chapter explores the fundamental concepts of environmental economics with a view towards how economic policies affect businesses and their choices.

While covering some of the basic elements of environmental economics, it is important to note that this chapter does not cover a sibling area: resource economics. A discussion of the economics of renewable and non-renewable resources is beyond the scope of this book.

The material in this chapter leans to some extent on the coverage of environmental economics in a number of popular introductory texts in this area. There are four texts that are highly recommended as further reading. First, Field and Olewiler (2011) is a widely used introductory text on environmental economics with a focus on Canadian applications. Goodstein (2010) is the U.S. counterpart, written with a focus on recent topical issues. Some of the examples in this chapter were adapted from this source. A book with a wider scope (it includes an excellent coverage of resource economics) is Perman et al. (2003), which aims at both undergraduate and graduate readers by including advanced material in chapter appendices. Last but not least, McKittrick (2011) is an advanced text that focuses on the economic analysis of environmental policy, and is aimed at graduate students who are comfortable with mathematical modelling and econometric tools. All four books are excellent sources to gain a deeper understanding of environmental economics. Considering the wide interest in emissions trading markets, another useful source is Tietenberg (2006). This is one of the few rig-

orous treatments of the topic that includes a thorough discussion of transaction costs, spatial and temporal dimensions, market power and enforcement.


3.1 Basic economic concepts

3.1.1 *Pollutants and externalities*

The starting point for economic analysis of the environment is the notion of a **negative externality**, a spillover from an economic activity that is costly to a third party. Pollution is the classic example of a negative externality. It is a cost that is not borne by the buyer and seller but by the public at large. Because negative externalities are not captured by the price of the good that is being sold and bought, they constitute a type of **market failure** that is caused by a lack of **property rights**. The extreme case of a lack of property rights was described as the **tragedy of the commons** by Garrett Hardin in 1968. It describes a situation where multiple parties, acting only in their own self-interest, destroy an open access shared resource even though this is not in their long-term interest. Overfishing that leads to fish stock collapse (see discussion in section 11.3) is an example of such a tragedy of the commons. Environmental economics has an effective recipe for dealing with such negative externalities and property right deficiencies. The general solution to the externality problem is to find a way to internalize the externality, which involves compensating the affected parties for the pollution they are exposed to, or to create property rights over the affected environment or resource. A large part of the analysis in environmental economics is concerned with finding the most efficient way to internalize negative externalities. Often this involves allocating and managing property rights effectively.

Before exploring specific tools in the toolbox of environmental economists, it is useful to step back and first look at the nature of the negative externalities that we are dealing with. Figure 3.1 provides a simple schematic of the economic-environmental system. Industrial activity generates a flow of pollution. Such pollution flow generates some immediate (short-term) environmental damage due to the exposure of humans and the ecosystem at large. Nature is able to absorb many pollutants over time and render them harmless; this ability is described as nature's **assimilative capacity**. This process can be relatively quick (especially when dispersion and dilution are involved), but often it is rather slow. When nature is not able to cope with long-lived pollutants, there is a build-up of a pollution stock, with damage to humans and the

Fig. 3.1 Pollution Sources and Sinks



ecosystem. This means that the harm from these pollutants may continue well after the emission flow has been discontinued.

The environmental consequences of pollution flows and pollution stocks are transformed through the ecosystem itself. A pollutant's health hazard evolves over time as the pollutant may undergo chemical or biological transformations. For example, nitrous oxide emitted from soil fertilizers breaks down when it reaches the stratosphere and forms ozone-destroying nitrogen oxides. Biological processes can also increase the health risk from pollutants. Two important processes are **bioaccumulation** and **biomagnification**.

Bioaccumulation describes the build-up of toxic pollutants in the tissue of living organisms due to ambient exposure to the pollutant. The effect is more pronounced in organisms that are unable to metabolize or excrete the pollutant. Contaminated food supply may increase the pollutant dosage that an organism is exposed to by orders of magnitude from the ambient level. While the ambient pollution concentration may be harmless, the accumulated long-term dosage may be harmful or even lethal to the organism.

Bioaccumulation is often accompanied by biomagnification, where pollutants are passed on from lower organisms in the foodweb (a lower trophic level) to higher organisms in the foodweb (a higher trophic level). For example, prior to 1972 the insecticide DDT was used in the United States to combat mosquitoes, which transmit diseases such as malaria. Rain washed DDT into rivers and lakes and into the fish stock. A number of predatory bird species ate contaminated fish, and the DDT's thinning effect on the birds' egg shells led to their premature

breaking, thus decimating these bird species.¹ DDT has also been linked with direct health effects on humans, including miscarriage, premature birth, difficulty breast-feeding, and developmental neurotoxicity.

Bioaccumulation (within a trophic level) and biomagnification (across trophic levels) require that the toxin is (i) long-lived (it does not decay rapidly), (ii) mobile (typically airborne or waterborne), (iii) soluble in fats, and (iv) biologically active (has a measurable effect on cell metabolism).

The effect of pollutants has both a time and a space dimension. The time dimension is reflected in the process of bioaccumulation and biomagnification. Temporal clustering is also important with respect to the **dose-response** relationship of a pollutant—i.e., the change in effect of an organism caused by differing levels of exposure. The ambient concentration (the dose) of a pollutant often has a non-linear effect on health (the response). Therefore, short-term exposure to a high dose may have more detrimental effects than long-term exposure to a low dose. This non-linear relationship has important consequences for the design of environmental policy. Emission standards need to take into account the **temporal clustering** of exposure and set separate limits for short-term exposure and long-term exposure. For example, Canada's recommended ambient air quality standard for carbon monoxide limits the maximum acceptable level to 30 ppm over a one-hour averaging period and to 13 ppm over an eight-hour averaging period.

The spatial dimension of pollutants matters as well. One needs to distinguish between **local** pollutants (with adverse effects limited to a few kilometres from a point source), **regional** pollutants such as sulphur dioxide that travel over tens or even a few hundreds of kilometres, and **global** pollutants such as the greenhouse gas carbon dioxide. Technically, local and regional pollutants are referred to as non-uniformly mixed pollutants, while global pollutants are referred to as uniformly mixed pollutants. The **spatial clustering** of pollutants creates what is known as **hot spots**, areas with very high ambient pollution concentrations. In the 1960s it was quite common to solve the hot-spot problem by building higher smoke stacks and following the mantra 'the solution to pollution is dilution.' While spreading out the pollutant more equally over a larger area reduces ambient concentrations, it is not a particularly effective method because adverse weather conditions can still lead to extreme local concentration levels.

The spatial dimension matters in yet another way. Pollution can cross jurisdictional borders. Such **transboundary pollution** is very common and limits the ability of a single jurisdiction to tackle environmental problems effectively. The need for policy coordination is often

¹ The 1962 book *Silent Spring* by biologist Rachel Carson, which describes the ill effects of DDT, is often considered one of the defining events in launching the environmental movement in the United States, and perhaps even worldwide.

paramount but limited by the constitutional allocation of responsibilities or the inability to conclude effective cross-jurisdictional or international treaties. These problems are examined later in the chapter on environmental law.

3.1.2 *Public goods and property rights*

Private goods differ in two key dimensions from **public goods**, which are goods enjoyed in common. Public goods are described by two features: **non-excludability** and **non-rivalry**. Excludability requires that it is possible to prevent people who have not paid for the good from using it. A rival good cannot be consumed by two consumers simultaneously. Conversely, a non-rival good is one which can be supplied at zero marginal cost to a second and further consumers, and thus can be consumed by an unlimited number of consumers. The two dimensions form a classification matrix with four different types of goods:

	Excludable	Non-excludable
Rivalrous	Private goods food, clothing, cars, toys	Common goods fish, forests
Non-rivalrous (simultaneous consumption)	Club goods cable television, cinemas	Public goods lighthouses, air, national defence

While public goods are characterized by non-excludability and non-rivalry, common goods and club goods are cousins with only one of these characteristics. Common goods such as fish and forests are non-excludable (or at least not easily excludable) but rivalrous: a fish I eat you can't eat; a tree I log you can't log. Club goods such as cable television are easily excludable, but to the extent that households already have a cable connection, adding another subscriber comes at practically zero marginal cost to the cable company. Put another way, if I watch a television show, any number of people can watch the same show at the same time.

The environment is a public good and as such encounters the problem of **free riding**. Consumers can take advantage of public goods without contributing sufficiently to their creation, and with respect to the environment, the environment can be used freely without contributing to keeping the environment clean. Free riding is a consequence of poorly defined **property rights**. A solution is therefore to create property rights. Property rights can create excludability, or rivalry, or both, depending on the particular nature of the property rights.

Restoring property rights is the crucial step in fixing a problem with a negative externality. The concept was pioneered by Nobel Prize laureate Ronald Coase in the 1960s and is known as the **Coase theorem**.

Coase theorem: When trade in an externality is possible and when there are no transaction costs and free riding, bargaining will lead to an efficient outcome regardless of the *initial* allocation of property rights.

As far as the environment is concerned, the Coase theorem suggests that governments should create property rights over its use. Regardless of who owns the property rights, the externality will be internalized. Who owns the property rights is important only in regard to the distribution of income from these property rights.

If the government endowed the polluter with a right to pollute, then the pollution victim would bargain to pay the polluter to reduce pollution. If, on the other hand, the government endowed a city or province with a right to clean air, then the polluter would bargain to pay the city or province for permission to pollute. Of course, in the first case revenue would flow to the polluter, and in the latter case to the city or province. The outcome will be economically efficient either way, but the two solutions will be perceived as differently fair.

The Coase theorem acknowledges one major caveat: the presence of **transaction costs**. The efficient outcome may not be feasible when the parties find it difficult to negotiate and transact. For example, the victims of pollution may find it difficult to coordinate among themselves, and firms would find it difficult to contract with thousands of individuals rather than a single other party. Where such transaction barriers exist, governments may need to play the role of the contracting party and exercise the property right on behalf of a large number of people.

There is a second caveat: **incomplete information**. When property rights are allocated to one party, that party has an incentive to withhold crucial information to strengthen its bargaining position. In the absence of complete information, the efficient outcome may not be possible to achieve. Furthermore, when the negotiation between parties occurs over multiple rounds, reputation effects can arise that distort the outcome.

In practice, the Coase theorem points us in the right direction (create property rights) without telling us the precise destination (whom to allocate the property rights) or the shortest route there (with the least transaction costs, and with the most complete information). Nevertheless, property rights can be used quite effectively in many instances. For example, individual transferable quotas (ITQs) are used to manage rights over fish stock, and tradeable emission permits (TEPs) are

used to manage emissions into air. In both cases, though, initial allocation matters in practice. While property rights can be very efficient, big questions remain about fair allocation of such rights.

3.1.3 *Efficiency and fairness*

In the discussion about the allocation of property rights and the Coase theorem, it was seen that outcomes involve both efficiency and fairness. An outcome that is economically efficient may be viewed as unfair, while an economically inefficient outcome may be regarded as fair. How does economics reconcile these notions?

Economic efficiency is a straight-forward optimality concept involving a state without scope for further improvement. In environmental economics, the notion of Pareto efficiency is crucial:

Pareto efficiency describes an outcome in which it is impossible to make one person (or contracting party) better off without making anyone else (or the other contracting party) worse off.

Finding economic efficiencies often involves allocating property rights and letting the different parties bargain over the outcome, which may make both parties better off. Box 3.1 illustrates the situation of water rights in California and the potential for a Pareto-improvement.

Efficiency does not imply fairness. Both polluter or victims can pay to reduce pollution and achieve an efficient outcome. Yet most of us would probably instinctively object to giving polluters the right to pollute rather than citizens the right to clean air. In economics, fairness is a notion that applies to the distribution of income and wealth, and thus fairness is closely linked to distributive equality. Good social outcomes are 'fair.' Suppose two equally productive people R and P started out with an income of \$50,000, so that the total is \$100,000. Now double the total income and give all of the extra income to person R and none to P. This is Pareto-efficient (nobody is made off worse), but probably not very fair. The notion of fairness goes beyond economics and is ultimately political in nature: what is fair—or not—is defined by societal norms, which may change over time, and may differ across cultures. Much of the following analysis is therefore focused more on efficiency and less on notions of fairness. The problem with many environmental policies, though, is that many policies that are economically efficient also create winners and losers—and this becomes quickly a matter for political battles.

Fairness is also concerned when it comes to the location choice of polluting firms. There have been concerns about **environmental jus-**

Box 3.1: Water Rights in California

California is legendary for its water wars (which came to Hollywood fame in the 1974 movie *Chinatown*). The unending thirst of a growing Los Angeles led to the building of aqueducts and the draining of lakes, resulting in battles between the city and rural communities. California is also a state of water extremes, with years of strong rainfall and other years of drought. The perceived shortages of water in California are at least in part an example of mismatched supply and demand.

When California was settled, the settling farmers were endowed with most of the water rights. Today, they use more than three-quarters of the available water at a price that is set by the state and the federal government as low as 1 cent per cubic metre. By comparison, residents of Los Angeles pay 9.25 cents per cubic metre, almost ten times as much as farmers pay. This is clearly inefficient and

explains why irrigation is overused to produce agricultural output. This can also be viewed as a huge subsidy to farming.

Is there a Pareto-improvement? Imagine there was a single market price for water that equalized supply and demand for water in California. If farmers were allowed to sell their water entitlements to urban dwellers, quite a few would probably find it profitable to sell water to city dwellers or commercial users rather than use it to produce crops. Reallocating water through a market mechanism could boost California's economy through increased efficiency.

There are, of course, other inefficiencies in California's management of water. As in many other jurisdictions, water use is often not metered and is essentially provided at a flat rate, thus offering little monetary incentive to conserve water.

tice when polluting firms are concentrated in poor neighbourhoods or neighbourhoods with larger shares of demographic minorities. While there may be some self-selection involved (people choosing to live close to the plants they work for, rather than plants choosing to move into neighbourhoods with particular demographics), studies have found unfavourable correlations between poverty and the presence of hazardous waste sites.

In the real world there are large transaction costs between the many individuals in society and the few polluting firms. Efficiency can be achieved under a **polluter pays principle**. Requiring a polluter to pay for the right to pollute is likely to be more efficient than the other way around because it reduces transaction costs and free riding. What is more important is that endowing firms with a right to pollute would lower their overall cost, lead to more firm entry into the industry, and create yet more pollution.

What is an efficient solution to pollution abatement? Firms all have different capabilities to reduce emissions. For some it will be very expensive to reduce emissions on a per-unit basis, and for others it will be much less expensive. Who should reduce emissions more: the 'dirty' or the 'clean' firm? To rank the ability of individual firms to reduce their emissions, we need to know each firm's per-unit **marginal abatement cost (MAC)**. Marginal abatement cost is usually expressed in dollars per physical unit (e.g., metric tonne of emissions). Firms usually have different MACs at different levels of output or emissions. As a firm's share of abated pollution rises, the cost of the extra abatement continues to rise. Reducing 80% of emissions may be relatively cheap, but moving from 80% to 90% abatement efficiency may become more costly, and going from 90% to 99% abatement efficiency may involve yet more costly abatement technology. As a result, firms usually have an upward-sloping MAC curve with respect to the abatement level (or abatement efficiency).

Efficient pollution abatement: the least-cost pollution abatement regime requires that the marginal cost of abatement is equalized across all firms undertaking pollution abatement.

A least-cost solution will in general not involve equal abatement effort by all polluters: some polluters will have higher per-unit abatement costs than others. When MACs differ across firms, cost efficiency implies that relatively low-cost abaters will undertake most of the total abatement effort, but not all of it. On the other hand, the relatively high-cost abaters will carry out less abatement effort. This may lead, counterintuitively, to the 'dirty' firms abating less than the 'clean' firms. But as we are only interested in an overall reduction of emissions, it does not matter who reduces the emissions as long as we achieve the desired emission reduction at the lowest overall cost. But how do we get the low-cost abaters to do most of the pollution abatement? What are the best policy instruments to achieve the efficient outcome?

3.2 Policy instruments

The economic answer to the problem of negative externalities is quite simple: put a price on the negative externality. The fundamental problem with negative environmental externalities is that there is a gap between the private cost and the social cost of the polluting activity. The social cost is borne by society, but not the producer. Therefore, the efficient solution is to close the gap. Arthur Pigou discovered the solution to this problem in 1920:

Pigouvian principle: pollution should be priced at its marginal external cost.

The marginal external cost is the cost resulting from the production of one additional unit accruing to a different party than either the producer or consumer of the product. A **Pigouvian tax** is intended to cover the social cost (negative externality) of an economic activity and restore the efficient outcome. It is important to understand that a Pigouvian tax is efficient only when it is applied to the pollutant (or the externality) directly on a per-unit basis. For example, a tax on automobiles to combat carbon emissions is not as effective as a tax on carbon emissions (or gasoline) because cars differ in fuel efficiency and the mileage they are driven.

When we look at individual firms, a Pigouvian tax will increase the firm's pollution abatement effort. The firm's marginal abatement cost will rise as it steps up this effort. This means that the MAC curve will be upward sloping in abatement effort, or downward sloping in actual emissions. Reducing emissions by 10% will be less costly per unit of output than reducing emissions by 80% per unit of output. In other words, there are low-hanging fruit of pollution abatement (e.g., switching fuel sources) that are relatively cheap, and there are costly capital investments in pollution abatement equipment (e.g., installing a scrubber, perhaps connected to an electrostatic precipitator).

Armed with data on the abatement cost of firms and thus the 'supply' of pollution, what do we know about the 'demand' for pollution? It may seem counterintuitive to think that anyone would 'demand' pollution, but in fact we all do. Any time we get behind the steering wheel in a conventional automobile and drive, we emit pollutants such as carbon monoxide and carbon dioxide. By deciding to drive this particular car we 'demand' the pollution that goes along with this activity. Of course, if this pollution costs us extra money (in the form of carbon taxes or higher gasoline taxes), we may decide to drive less or drive a more fuel-efficient vehicle. Our demand for pollution is leading to environmental damage, and this damage tends to increase on a per-unit level as emissions increase. This is another way of saying that there is a non-linear dose-response function. Technically speaking, marginal damage is upward sloping in emissions. Neoclassical economic analysis tells us that the efficient outcome is where pollution supply meets pollution demand, and this happens when **marginal abatement cost equals marginal damage**. This clean theoretical result encounters an empirical dilemma, unfortunately.

While marginal abatement cost is observable (at least by the individual firm, if not the regulator), marginal damage is often very difficult to pinpoint. In practice, the regulator steps in to reflect society's belief

about the marginal environmental damage—and this may be on target or off target. In the discussion about the different policy instruments we need to keep an eye on the effect of **uncertainty** about marginal damage. What happens when the regulator gets it wrong? Furthermore, the regulator may not always be able to observe correctly each firm's marginal abatement cost. This is a situation of **information asymmetry**: one side (the firm) knows more than the other side (the regulator).

Policy instruments fall into three broad categories:

- **Information-based** policy instruments are voluntary in nature. They induce firms to make public information about their emissions in the hope that pressure from stakeholders, along with knowledge sharing among firms within an industry, will lead to emission reductions without direct government intervention. The *Toxics Release Inventory* in the United States and the *National Pollutant Release Inventory* in Canada are tools that inform the public. Governments promote cooperation among firms through voluntary challenge programs. In the United States, the 33/50 program targeted 17 priority chemicals and set as its goal 33% and 50% emission reduction by 1992 and 1995, respectively. In Canada, the ARET program targeted a wider set of pollutants for virtual elimination between 1995 and 2000. Consumers armed with more information alone are not necessarily putting a lot of pressure on firms: the empirical evidence for **green consumerism** is relatively weak. Another popular information-based approach involves product labelling or certification whereby firms signal environmental performance to consumers. Consumers who value such environmental credentials may be willing to pay a **green price premium** for such products.
- **Mandated** (or direct) policy instruments are command-and-control measures: they give firms a precise set of rules that they are required to follow. Such rules come in two flavours: an **emission standard** regulates the environmental performance of production facilities (such as a maximum level of emission concentrations), while a **technology standard** regulates the production process by mandating specific production methods or pollution abatement equipment.
- **Incentive-based** policy instruments work indirectly through price signals and the market mechanism. The simplest form is a Pigouvian tax in the form of an **emission tax** that sets a price for the pollutant. The **carbon tax** in British Columbia (\$30/tonne since 2012) is an example of such a tax. An alternative is the use of **tradeable emission permits** whereby the regulator auctions or allocates emission rights to firms. Once in possession of these permits, firms can trade freely in them but at the end of a reference period must possess sufficient permits to cover all their emissions. Governments can also use a variety of other taxes to influence environmental outcomes, albeit less directly. As particular intermediate goods and ser-

vices are linked to emissions, **input taxes** can target them. A fuel tax can target simultaneously a variety of pollutants associated with the burning of such fuels. Instead of using taxes, governments can also pay firms directly to reduce emissions. Governments often prefer subsidies over taxes for political reasons: everybody loves getting money from the government, but nobody likes paying money to the government.

3.2.1 *Instrument choice criteria*

There is a large number of possible environmental policy instruments. How should a regulator choose among them? Which instruments are preferred by firms? Regulators and regulated need to consider a number of criteria. Below there are grouped together into four main categories:

1. **Environmental effectiveness**

- **Directedness** is the degree to which the policy is aimed at the environmental problem. First-best approaches target the environmental problem directly, while second-best approaches take aim through a secondary mechanism. For example, reducing automotive emissions through a fuel tax is a direct approach, while taxing vehicles (rather than the amount of driving) is an indirect approach.
- **Short-term effectiveness** is achieved by going after the biggest pollution sources first. These are often the 'low-hanging fruits' of policy intervention, such as shutting down old and highly-polluting facilities before designing policies that encourage further emission reductions across all facilities.
- **Long-run effectiveness** requires that the policy's potency does not diminish over time, for example through substitution effects. If firms or households can avoid the policy by switching to other (unregulated) products or processes with other adverse environmental effects, the policy will become increasingly ineffective.

2. **Economic efficiency**

- **Static efficiency** is achieved when the emission reduction target is obtained at the lowest possible abatement cost. Incentive-based policies are more efficient than mandated policies.
- **Dynamic efficiency** is achieved when the policy creates persistent and predictable incentives to innovate into pollution abatement, emission-reducing products or emission-reducing manufacturing processes.

- **Flexibility** is necessary to ensure that a policy instrument can be adapted quickly and effectively to new situations and circumstances. Firms have a high degree of autonomy in choosing appropriate technologies to achieve the desired environmental outcomes.
- **Robustness** ensures that the policy remains economically efficient even in the presence of high uncertainty (or incomplete information) about abatement costs or environmental damage.
- **Simplicity** is related to informational costs. A policy that is complex and has numerous transaction costs is less desirable than a policy that is simple and does not require much information gathering, monitoring, or enforcement action by regulators.

3. Fairness and participation

- **Fairness** is the notion that the environmental policy does not favour one societal group over another, or put undue burden on a small group in society.
- **Involvement** that is wide and deep across society ensures that stakeholders take ownership of the policy, especially if the policy targets modification of behaviour. Involvement helps entrench a policy against opportunistic abandonment of the policy.

4. Political considerations

- **Political feasibility** requires that an environmental policy is supported by the electorate and elected representatives. Widely unpopular policies may be difficult to adopt even if they are environmentally sound.
- **Budget implications** are the effects of the policy on the government budget. Taxes raise revenue, while subsidies incur costs. Revenue-neutral policies are often preferred.
- **Ancillary benefits** arise when a policy has beneficial secondary effects. Such a policy is said to have a double dividend. For example, tax revenue from the environmental policy can be recycled to offset inefficiencies in other parts of the economy, the tax system, or to pursue other environmental objectives.

This is a long catalogue of criteria, and while the economic efficiency theme stands out among them, it is quite apparent that no single instrument is clearly superior along all the dimensions relevant to policy choice. Significant trade-offs arise in the choice of instruments. In addition to considering economic efficiency and fairness (distributional equity), regulators often find that political feasibility imposes the most important constraints on deploying particular policy instruments.

No single environmental policy instrument is superior along all dimensions; hybrid instruments, or a combination of instruments, are often able to strike a better balance among economic efficiency, distributional fairness, and political feasibility.

In the presence of conflicting choice criteria, **hybrid regimes** may sometimes become preferable over pure instruments (Goulder and Parry, 2008). For example, one can make strong case for a regime of tradeable emission permits with upper and lower price ceilings—i.e., a hybrid system of emission quotas and emission taxes. For a given environmental problem, multiple market failures may be involved. In this case, multiple instruments may need to be deployed simultaneously on efficiency grounds. There is also the possibility of potential interactions among environmental policy instruments. If hybrid regimes are not designed carefully or emerge through overlapping jurisdiction, interactions among these policy instruments may reduce their overall efficacy.

Yet another problem is that environmental policy instruments may need to be adjusted to allow for existing market distortions. Specifically, when an industry is imperfectly competitive, an emission tax may lead to an overly strong output reduction. This implies that emission taxes ought to be somewhat lower in oligopolistic than in competitive markets.

When regulators deploy environmental policy instruments, they affect a firm's choice of (i) input mix (especially different types of energy); (ii) use of abatement technology; (iii) output level; and (iv) innovation activity (with respect to redesigning products or processes). When we look at the individual firm, an environmental policy may affect one or more of these at the same time, and this may not always lead to effects working in the same direction.

3.2.2 Technology and emission standards

Governments often feel at greatest ease setting a **technology standard**, as this can be accomplished through the process of granting a development permit for the construction of a new plant. This is a very simple once-and-for-all minimum standard that the firm has to adhere to, and that can be monitored and enforced easily through an inspection.

In general, technology standards are specific requirements regarding the production process and production technology. They can involve a mandate for an end-of-pipe treatment (such as installing a scrubber or other pollution abatement device) or the use of particular inputs (e.g., natural gas instead of coal).

Although technology standards are appealing because of their relative ease of implementation and monitoring, heterogeneity among firms makes it very difficult for the regulator to possess all necessary information to maximize the cost-effectiveness. Technical standards suffer from the **information asymmetry** problem, where the firm knows more about its own abilities than the regulator. Technical standards can easily evolve into a hodgepodge of plant-specific rules that vary over time and across jurisdictions. But even when technical standards are set as a single mandate that is equal across firms, they will not always result in a cost-effective outcome because different firms will face different marginal abatement costs due to differences in their production process, plant age, inputs, or other factors.

An **emission standard** is aimed at a firm's output rather than inputs or production technology. Such standards are typically expressed as **emission intensity** (pollutant per unit of output). Just like technology standards, emission standards are subject to the information asymmetry problem. In practice, emission standards are often sufficiently modest that they can be met by reducing emissions through input substitution and relatively benign end-of-pipe treatment. Emission standards such as fuel efficiency limits for automobiles—in particular, the Corporate Average Fuel Economy (CAFE) standard in the United States—often do not sufficiently encourage output reduction (e.g., number of automobiles on our roads) or output use (e.g., kilometres driven). On the other hand, emission standards are also relatively easy to enforce and monitor, and thus remain popular with regulators.

Despite their relative economic inefficiency, technology and emission standards also remain popular with regulated firms. The reason is their predictability and certainty. Technology standards are often once-and-for-all requirements that can be fulfilled during the construction stage of a new plant. Emission standards often evolve gradually, and are sometimes easily subverted. For example, the CAFE standard for automobiles was eroded by classifying the popular sport-utility vehicles as light trucks rather than cars, which were subject to different standards.


3.2.3 *Environmental taxes*

Among incentive-based policies, environmental taxes (also referred to as 'green taxes') stand out as providing the most immediate and stable price signal to firms and consumers. Unlike mandated policies, incentive-based policies are usually able to achieve economic efficiency because they let firms figure out how to respond optimally to the price signal. It is entirely up to the firms to find the cheapest way to reduce emissions. An individual firm will increase pollution abatement until the marginal abatement cost equals the price of the pollutant as defined

by the environmental tax or through the price in a tradeable permit market. The regulator does not need to know each firm's marginal abatement cost curve, and thus information asymmetry is not a problem. As long as the regulator knows that firms will equate the emission price to marginal abatement cost, the cost-efficient outcome will emerge. The regulator simply needs to figure out the correct price for the negative externality.

How should environmental taxes be introduced? The answer is: gradually. Like any regime shift, businesses require time to adapt. It is therefore highly advisable that government phase in an environmental tax over a number of years and raise the tax once a year until it has the desired effect. Such a gradual introduction avoids the risk of **overshooting** when regulators have imprecise knowledge about the market equilibrium that corresponds to the desired emission reduction.

Fig. 3.2 Cap-and-Trade versus Green Tax




Environmental taxes and tradeable emission permits are equally efficient in a static sense (comparing their effect in a single period). They can achieve the same outcome at the same abatement cost. For any environmental tax there is an equivalent cap-and-trade emission target. Figure 3.2 illustrates how the two policy regimes differ. The two diagrams show consumer demand as a downward sloping curve, which mirrors a normal demand curve with output and price. The only difference is that the horizontal axis shows emissions rather than output. Emissions are simply a by-product of output. Fluctuating demand will move the demand curve lower or higher.

With an emission tax (right panel in figure 3.2), the regulator sets a tax rate (horizontal line) and may thus hit or miss the emission target as demand shifts up or down. This is a case of quantity uncertainty. With a cap-and-trade regime (left panel in figure 3.2), the regulator sets

the emission target (vertical line) and the market discovers the price through the trading of emission permits. As demand shifts up or down, the market price fluctuates. This is a case of price uncertainty.

Fig. 3.3 Effect of Green Tax on Firm



The effect of an environmental tax on an individual firm is illustrated in figure 3.3. The economic logic for a cap-and-trade system is identical; it is different for a subsidy, however. Panel A in figure 3.3 shows a double diagram where the upper part depicts a conventional supply-and-demand diagram. Demand is downward-sloping: higher prices reduce demand. Supply is upward-sloping: higher prices will increase supply. Given supply and demand, the market equilibrium will emerge where the supply curve and demand curve intersect, with an optimal price P and quantity Q . In the lower half of the diagram, output Q translates into a corresponding amount of emissions E as governed by the currently used technology without abatement.

Panel B in figure 3.3 demonstrates the effects of an environmental tax, which shifts up the supply curve by making output more expensive for any given amount of output. The distance between the two supply curves is the tax rate. The tax-induced upward shift of the supply

curve leads to a new market equilibrium, with a higher price $P' > P$ and a lower output level $Q' < Q$. There are now two effects to consider. First, the reduction in output will reduce emissions from E to E'' before abatement technology is deployed. This is the **scale effect**. Second, the environmental tax induces the firm to deploy abatement technology until its marginal abatement cost equals the environmental tax. This amounts to a rotation of the emission technology line in the lower half of the diagram. The blue line rotates into the red line. At the new and lower output level Q' , the new emission technology line translates output Q' into the new final emission level E' . This further reduction in emissions from E'' to E' is the **technique effect**. With an environmental tax and cap-and-trade regime, scale effect and technique effect work in tandem and complement each other. Which of the two effects is larger or smaller depends on the steepness of the firm's supply curve and the rotation of the emission technology curve. A flat supply curve magnifies the scale effect, which is usually not welcome. Lost output also means lost revenue for a firm. What is welcome, however, is a large rotation of the emission technology curve, which reflects low abatement costs and magnifies the technique effect.

Emission taxes generate government revenue, and this may make emission taxes politically unpalatable. In principle, governments can take the revenue from an emission tax and redistribute it in the form of lower taxes elsewhere. This has been described as a (revenue-neutral) **green tax reform**. British Columbia's carbon tax, introduced in 2008, returned the revenue to its citizens in the form of lower personal and corporate income taxes, as well as a one-time \$100 per head 'dividend payment.' Governments can also funnel the revenue into other environmental projects, from public transit to 'green subsidies.' This dual environmental benefit of green taxes is often described as the **double dividend**. Despite the possibility of revenue neutrality or a double dividend, many politicians consider proposing (let alone introducing) an environmental tax 'toxic' for their electoral prospects. Of course, this raises the question whether there is a method to construct an environmental tax that generates very little revenue while preserving the emission price. There are hybrid emission standards plus emission taxes that can accomplish this. Consider an environmental tax on firms only, where firms' MAC is ranked on an **abatement ladder**: the best firm with the lowest MAC firm is at the lowest rung, and the worst firm with the highest MAC is on the highest rung. The best firm probably uses state-of-the-art technology while the others are catching up technologically. The regulator can set an efficient emission tax based on a firm's excess emissions over an emission standard. If the emission standard is based on the best firm's emission intensity, the best firm will pay no tax, and firms that are performing worse pay increasingly more. If, instead, the regulator sets the emission standard at an industry average, the over-performers would in fact collect a subsidy that is paid for by revenue

from the under-performers. Such a **hybrid regime** might be as appealing to firms as an emission permit system with grandfathering.


3.2.4 Tradeable emission permits

Emission permits are known as cap-and-trade systems. The government allocates a fixed number of emission permits (the cap on pollution) and firms trade their allowances among each other in order to balance their emissions optimally in response to the emerging market price for the emission permits. When an emission permit is traded, the buyer is paying for the right to emit more pollution, while the seller is giving up that right because the seller does not need it. The buyer profits from the purchase because buying the permit is cheaper than reducing pollution through costly pollution abatement. The seller profits from the sale because the permit price is higher than the seller's marginal abatement cost. Thus the buyer can increase pollution until that firm's MAC equals the permit price, whereas the seller is reducing pollution until its MAC matches the permit price. Both firms gain from the trade.

Trading benefits

To illustrate the key idea of emissions trading, consider two firms A and B with different abatement costs. Let C_i denote firm i 's total abatement cost, and let Z_i denote firm i 's emission abatement. In panel A of figure 3.4 the marginal abatement cost (MAC) curves are upward sloping. Specifically, let $C_A = 100 + 1.5Z_A^2$ and $C_B = 100 + 2.5Z_B^2$, so that $MAC_A = 3Z_A$ and $MAC_B = 5Z_B$. The MAC curve for firm A is flatter than for firm B, suggesting that firm A is the 'clean' firm that finds it easy to abate, while firm B is the 'dirty' firm that finds it difficult to abate. The industry-wide $MAC = (Z_A + Z_B)/(1/3 + 1/5) = 1.875Z$, where $Z = Z_A + Z_B$. In the diagram, the industry-wide MAC can be thought of as stacking together horizontally the MAC_A and MAC_B curves. If the regulator requires abatement of $Z = 40$ units, this implies a permit price of $P = MAC(40) = \$75$. Both firms face the same emission price, and at \$75 firm A will carry out more abatement (25 units) than firm B (15 units).

Let us now consider a scenario where we start with original unabated emissions equal to 100 units, 50 from each firm. The regulator wants to reduce emissions by 40 units to a final emission level (the 'cap') of 60 units. The regulator allocates the 60 permits corresponding to this cap equally to both firms, so both firms receive 30 permits. This scenario is illustrated in panel B of figure 3.4. With 30 permits in hand, each firm still has to abate 20 emission units. When abating 20 units, firm A faces an MAC of \$60, while firm B faces an MAC of \$100. They both can do better.

Fig. 3.4 Firm Heterogeneity and Marginal Abatement Cost

If firm A increases abatement by five units, it could sell five of its emission permits that it then no longer needs. Increasing abatement from 20 to 25 will cost firm A \$337.50, or \$67.50 per unit. If it could sell those five permits for \$75, it would get paid \$375 and thus make a profit of \$37.50. If firm B were to purchase these five permits at \$75, it would pay \$375 for them. Being able to increase emissions by 5 units and reduce abatement by 5 units would allow firm B to reduce abate-

ment cost by \$437.50, which generates profits of \$62.50 for firm B. Note that the profits from the trade do not have to be equal for firms A and B. As long as their profits are positive, they will trade. Initial allocation determines which firms gain more from trade, and which will gain less from trade. How the firms allocate abatement effort among themselves is summarized in the table below:

	Firm A	Firm B
Initial unabated emission	50	50
Initial permit allocation	30	30
Initial necessary abatement Z^o	20	20
Initial $MAC(Z^o)$	60	100
Optimal $MAC(Z^*)$	75	75
Optimal abatement Z^*	25	15
Permit demand $Z^o - Z^*$	-5	+5
Final permit allocation	25	35

Initial allocation

Before emission trading can commence, permits must be allocated in some form to firms. The way in which this **initial allocation** takes place generally does not affect the efficiency of the trading outcome; this is an application of the Coase theorem. Nevertheless, the initial allocation has distributional implications and creates winners and losers.

One way to allocate emission permits is through lotteries, a type of random access. This is a cheap method with the advantage of placing all permits into the hands of firms without creating government revenue or favouring one firm over another. It is a method immune to lobbying. Nevertheless, this method is favoured neither by governments nor by businesses.

The methods most widely used involve **grandfathering** and **auctions**. Grandfathering involves the allocation based on eligibility criteria or past performance such as output and employment, or in some cases previous-year emissions. Eligibility criteria can become problematic if they can be manipulated easily or when they create perverse incentives. For example, if the emission levels of the year before the start of the permit system are used as the basis for the allocation, firms would have an incentive to increase emissions purposefully in the benchmark year in order to obtain more emission permits the next year. This is a problem of **moral hazard**. Firms generally like grandfathering because it places all permits in their hands without paying the government. Payments for traded permits end up in the hands of other firms, not the government.

Governments generally prefer auctions of emission permits because auctions generate revenue for them, just like emission taxes. In fact,

if firms get their initial allocation just right during the auction, there would be no need to engage in further emission trading.

What are the pros and cons of both methods of initial permit allocation? Grandfathering creates immediate winners and losers as the government decides the initial allocation. Governments will try to use allocation criteria that are perceived to be fair, but the allocation formula is very likely subject to intense lobbying from interested parties. At the very least the allocation formula will need to reflect the relative size of the firms, which can be captured by measures of employment, output, sales, or value added. Grandfathering will result in substantial trading volume because the government's initial allocation will be suboptimal. On the plus side, grandfathering will be politically easier to sell because it does not generate revenue for the government and thus is not perceived the same way as a tax. An initial auction of permits has the opposite advantages and disadvantages. It is more likely to provide an initial allocation that is closer to the optimal allocation, and thus there will be less trading volume during the year. As trading is costly for firms (especially smaller firms), auctions tend to have lower transaction costs than grandfathered allocations.

Smart governments may choose a hybrid allocation method that involves both auctions and grandfathering. The grandfathering component buys political acceptance as it keeps a lid on the regime's overall cost for firms. The auctioning component reduces overall transaction costs and allows governments to take the revenue and correct other distortions in the tax system that ultimately benefit firms.

Intertemporal issues

Emission permits provide companies with the right to emit a given quantity of emissions over a fixed period, typically one year. At the end of each year, each company must surrender enough allowances to cover all its emissions. Companies that exceed their allowance (including those they bought in the permit market) are fined. It can happen that a cautious firm keeps too many allowances because of variability in output and emissions. Because of the fine, a company does not want to end up with too few permits, and thus it may want to hold an inventory of permits that is slightly too large. So what should the company do with the leftover permits at year's end? And what should a company do if it ends up, unexpectedly, with too few permits at year's end and wants to avoid a penalty? Such intertemporal issues can make a permit system less attractive than an environmental tax unless there is a way to smooth out the variability in permit demand over more than the reference period for the permit. Emission volumes can be quite volatile over time due to changes in output or because of technical problems (spills and leaks). Flexibility in permit trading can be obtained through **permit banking** and **permit borrowing**.

Permit banking helps firms with an unused inventory of allowances at year's end. The firm can use up the unused allowances in the next reference period. Thus, a firm can keep a contingency reserve of allowances in case there is an unexpected need for them without having to worry that the permits expire at year's end. Permit banking allows firms to hold a permit beyond its designated date for later use.

Permit borrowing works in the other direction. If a company runs unexpectedly short of permits at year's end, it can borrow permits from next year's allotment and apply them against the excess emissions of the current year. Permit borrowing thus allows firms to use permits before their designated use period.

The use of permit banking and permit borrowing enhances the dynamic efficiency of emission permit trading. This flexibility lowers the cost of participation in the emission trading system significantly. There is a downside to this flexibility, however. If firms can borrow and bank permits, they may be able to have very large emissions in one year and much lower emissions in other years. This **temporal clustering** of emissions can lead to high peak concentrations of emissions at certain points in time. With greenhouse gas emissions such as carbon dioxide this hardly matters because temporal clustering is not an issue: the time horizon for climate change is measured in decades rather than years. However, for many local pollutants with a strong non-linear dose-response function, high peak concentrations matter. Then, banking and borrowing may exacerbate temporal clustering.


From a firm's perspective, the ability to engage in permit banking has another advantage. If the firm expects that the regulator will increasingly tighten the emission cap over years to come, the permit price will necessarily increase over time. A firm may thus find it useful to buy permits when they are still cheap and use up (or sell) them when they are more expensive. If the expected price increases are substantial, banking permits could even become a speculative investment. To keep a lid on such speculative activities, regulators typically will limit the scope of banking (and borrowing) to a few years at most.

3.2.5 *Hybrid regimes*

Hybrid regimes that involve both emission permits and emission taxes can provide beneficial effects that limit the fluctuations of prices and/or emissions. Lesser variability may have economic and ecological advantages.

Figure 3.5 illustrates two important hybrid systems. In an influential contribution, Roberts and Spence (1976) made the case for a hybrid emission permit trading system that includes a price floor and price ceiling, which confines the permit price to the range between floor and

Fig. 3.5 Hybrid Regimes



ceiling and thus dampens price volatility. The left panel in figure 3.5 shows how the price range of the permits is limited. As demand for permits shifts, the permit price fluctuates within the price range. The hybrid regime thus involves a limited amount of price uncertainty *and* a limited amount of quantity uncertainty.

A price floor can guarantee a minimum abatement effort if the permit price falls below the price floor. A permit price may drop unexpectedly due to an economic slowdown, recession, economic shock, or significant technological change. The price floor will guarantee that firms will retain a sufficient incentive to innovate (see section 3.2.8 below). In practice, a price floor can be maintained by governments' buying back permits in the open permit market, by setting a reserve price at the initial permit auction, or by charging an extra environmental tax on each used permit.

A price ceiling is implemented as a safety valve to protect firms from runaway costs if there is an unexpected shortage of permits, perhaps due to an economic boom or another (positive) economic shock. A price ceiling can be managed in different ways. One way is for the government to sell extra permits at the ceiling price. Another way is to release an extra allotment of allowances. This is shown in the right panel of figure 3.5. When the permit price reaches the price ceiling, the regulator will release an **allowance reserve** that will keep prices at the price ceiling. If permit demand exhausts even the allowance reserve, permit prices will continue to increase.


Regardless of which hybrid method is used, the rationale for using such a system is compelling because it limits uncertainty and makes the cap-and-trade system more predictable and less worrisome for participating firms.

3.2.6 Subsidies

For obvious political reasons, governments prefer popular over unpopular measures. Because subsidies inject money into the economy, they rarely receive much criticism despite the fact that governments often exhibit favouritism in their selection of subsidy recipients. Taxpayers still end up footing the bill, but the source of funding is more opaque than with a new tax or fee. But are subsidies as efficient as taxes?

Subsidies can provide the same incentive for emission reduction as emission taxes or cap-and-trade systems. However, there is one major difference in the way subsidies work. Let us assume that a government pays firms a carbon reduction fee of \$20 per tonne rather than taxing firms at the same rate. So, for every tonne of carbon emissions that a firm reduces relative to a baseline, it would receive \$20 from the government. This subsidy will have one unintended effect: it will lower the firm's overall production cost and, in a competitive environment, the firm will pass on these cost savings to its customers. The lower prices will induce consumers to buy more output, and this in turn will increase emissions along with the increased output. As output increases, the original emission reductions will be partially offset. The **scale effect** and the **technique effect** work in opposite directions.

Fig. 3.6 Feebate



It is also possible to combine subsidies and taxes into a hybrid regime known as a **feebate**—a combination of fee and rebate. Such a system has been used to promote fuel efficiency in vehicle purchases. Figure 3.6 illustrates the concept. Consider the fuel consumption x_i for vehicle i , and assume that the average fuel consumption for the entire vehicle fleet is $\bar{x}$. Then a feebate would be calculated as $\tau(x_i - \bar{x})$, where τ is the feebate rate. If a vehicle's fuel consumption is above the average, the feebate turns into a tax (it is positive), and when the fuel consump-

tion is below the average, it turns into a subsidy (it is negative). Thus a feebate penalizes gas-guzzlers and rewards fuel-thrifty hybrid cars. Different types of feebate systems have been used by the federal government of Canada and the province of Ontario. What makes feebates particularly attractive to policy makers is that they can be revenue neutral; the revenues from the fee part of the program exactly offset the cost of the rebate part of the program if the pivot point $\bar{x}$ is chosen appropriately. In figure 3.6, the area of the triangle to the right matches the area of the triangle to the left of the pivot point. The **revenue neutrality** of new environmental policies can be important to win citizens' approval. As mentioned earlier, when the British Columbia carbon tax was introduced in 2008, it was designed to be revenue neutral by providing offsetting tax cuts and a one-time 'climate dividend' cheque to each household.

3.2.7 *Second-best approaches*

Interventions such as tradeable permits and emission taxes are considered **first-best approaches** because they price the negative environmental externality directly. To the extent that regulators can assess marginal damage correctly, they will set the price on the tax or quota of permits correctly; this leads to an efficient economic outcome. However, in many instances a first-best approach is not feasible for political or other institutional reasons. In such cases, regulators often need to rely on **second-best approaches**, which involve targeting activities that are highly correlated with the negative environmental externality.

A good example of a second-best approach is the **fuel tax** on gasoline and diesel fuel. A fuel tax can be used to target carbon emissions. On one hand, the carbon content of fuel is fixed at 2.328 kg of carbon dioxide (CO₂) per liter of gasoline and 2.614 kg/l for diesel. Thus, a fuel tax is equivalent to a carbon tax for these two fuel types. However, a fuel tax only covers the transportation sector and not CO₂ emissions from industry and households. A fuel tax is therefore an imperfect substitute for a carbon tax. Nevertheless, a fuel tax also addresses other negative externalities. By discouraging driving, it can relieve traffic congestion and motor vehicle accidents in cities. As vehicles emit a variety of air contaminants such as carbon monoxide, nitrous oxides, and hydrocarbons, a fuel tax also helps to reduce air pollution. Because emissions of such air contaminants are predominantly local, the high traffic density in urban areas leads to higher ambient pollution concentrations in those areas. It is therefore sensible to have higher fuel taxes in urban areas than in rural areas. Some jurisdictions actively pursue such a strategy. For example, as of May 2012, the provincial tax on gasoline in British Columbia varies across regions. It is highest (25.5 cents) in the

densely populated south coast district that includes Metro Vancouver; it is lower (18.0 cents) in the capital district that includes Victoria, the provincial capital; and it is lowest in the rest of the province (14.5 cents). The resulting price differences are remarkably immune to arbitrage; few drivers drive dozens of kilometres to refuel, or cross the border to the United States to refuel even more cheaply.

The fuel tax is also an example of a multi-purpose environmental intervention. It has multiple targets: global carbon emissions, local air pollutants, and traffic congestion. Fuel taxes are also quite efficient to collect, and they require relatively little administrative overhead. Such administrative efficiency may be an important consideration when the first-best policy instrument is difficult or costly to implement.

Another example of a second-best approach is the fuel efficiency feebate discussed in the previous section, or the CAFE standard used in the United States. Fuel efficiency standards target the negative externality (CO₂ emissions) only indirectly because the actual emissions depend on how much each vehicle is used. A 15 l/100 km gas-guzzler that is driven only 5,000 km per year emits fewer emissions than a 5 l/100 km hybrid car that is driven 20,000 km per year. Feebates are more efficient than the CAFE standard because the fee or rebate for individual cars is directly linked to each vehicle's fuel efficiency, rather than to the fuel efficiency of a manufacturer's entire vehicle fleet.

Economists dislike second-best policies because the rationale for not pursuing first-best policies is often weak. Embracing second-best policies may lead to the *slippery slope* phenomenon: if a second-best policy is good enough, a third-best policy may be good enough too. If there are obstacles to pursuing a first-best policy, it may be better (or cheaper) to remove these obstacles than to pursue a less-efficient second-best policy.

3.2.8 *Static and dynamic efficiency*

Time plays a significant role when comparing different environmental policies. Different policies can have different dynamic effects. In general, market-based instruments have better **dynamic efficiency** than command-and-control policies because they generate persistent incentives for innovation in pollution abatement technologies. With an emission standard, once it is achieved, there is no further incentive for reducing emissions. However, with an incentive-based policy, every further emission reduction will also reduce costs further. Companies will invest in R&D for abatement technology as long as the gains from the resulting innovations are offset by the savings from lower emission charges. In other words, with an emission tax or emission permit, every unit of

emission reduction is rewarded with a tax saving or fewer permits that need to be purchased.


The four panels of figure 3.7 illustrate the dynamic effect of an emission tax and an emission permit system. Each panel shows the emission price on the vertical axis and the abatement activity on the horizontal axis. In panel A, the dashed line on the right shows the level of pre-abatement emissions. At a given emission price τ , the firm will conduct pollution abatement until its MAC equals τ . The firm's total abatement cost is therefore the (blue) triangle plus the (green) rectangle; the former captures the total cost of the actual abatement activity, and the latter captures the cost of the emission permits or emission tax.

Now consider a technological innovation in abatement technology, shown in panel B. This rotates the firm's MAC curve clockwise from MAC_1 to MAC_2 , and this generates cost savings. This will increase the amount of abatement carried out by the firm, leaving fewer emissions to be covered through permits or to be paid for through an environmental tax. In panel B, the triangular area for the total abatement cost has now increased because the firm is abating more, but the rectangular area for the emission tax or permit purchase has decreased by a larger amount. The net savings are visualized in panel C as the two highlighted triangular areas combined. The left triangle captures the savings from cheaper abatement, and the right triangle captures the net savings from reduced tax payments or reduced permit purchases. The two triangles combined capture the firm's **innovation incentive**. A firm invests into R&D for new abatement technology if the total cost of developing and installing the new technology is less than the present discounted value of the savings in abatement costs *and* taxes or permit purchases.



How do emission permits and an emission tax differ with respect to the size of the innovation incentive? With a permit system, the market price for emissions fluctuates, while it is fixed (by the regulator) with an emission tax. For an individual firm, the emission tax leads to a predictable innovation incentive. By comparison, the size of the innovation incentive varies with the emission market price in the cap-and-trade system. This effect is illustrated in panel D. As τ shifts up and down, the size of the triangular innovation incentive becomes larger or smaller (dark or light shading). The price uncertainty introduced by the cap-and-trade system reduces the innovation incentive if firms are modestly risk averse. With an emission tax, the emission price is more predictable, and thus the innovation incentive is also more predictable.

Why does it matter if the innovation incentive fluctuates over time? The firm invests in R&D up-front, while the future cost savings from the new abatement technology arrive in the future. If the firm miscalculates, the R&D cost will exceed the discounted future cost savings. A risk-averse firm that wants to avoid losses will work with a low (lower than average) market price as the planning scenario. The innovation

Fig. 3.7 Abatement Costs and Innovation



incentive penalty of emission permit systems will grow along with the price volatility of the permit market. The more volatile the market price, the less the innovation incentive.

Fig. 3.7 Abatement Cost and Innovation (continued)*Panel C**Panel D*

Emission taxes and emission permits are generally equally efficient at achieving a targeted emission reduction in the short term. However, emission taxes are dynamically more efficient than emission permits because emission taxes provide steadier innovation incentives for pollution abatement technology.

3.2.9 *Dealing with uncertainty*

For the regulator there are two types of uncertainty for incentive-based instruments: uncertainty about firms' (or plants') abatement costs, and uncertainty about marginal damage.

In the presence of uncertainty about plants' abatement cost, the regulator sets an emission tax or permit allocation where the assumed marginal abatement cost (MAC) equals marginal damage (MD). This means that the resulting emission tax or permit allocation may be sub-optimal as a result of the information asymmetry between firms and the regulator. If the regulator assumes that firms' abatement costs are higher than they actually are, a tax will be too high or a permit allocation will be too generous. A tax that is too high generates excess abatement costs, whereas the overallocation of permits generates excess damage. On the other hand, if the regulator assumes that abatement costs are lower than they actually are, the emission tax will be too low, while a permit allocation will be too restrictive. Now the tax that is too low generates excess damage, whereas the underallocation of permits generates excess abatement costs.

With uncertainty about the MAC curve, do taxes or permits perform better? This depends on the shape of the MAC and MD curves and thus on the relative size of excess abatement costs and excess damage. When the MD curve is steep relative to the MAC curve, permits are preferred over taxes, and when the MD curve is flat relative to the MAC curve, taxes are preferred over permits. A flat MD curve implies that large changes in emissions have relatively small changes in terms of damage, whereas a steep MAC curve suggests that abatement opportunities are limited and costly. This scenario is somewhat more likely than the scenario where pollution abatement is easy and environmental damage sits at a trigger point. Consequently, when there is uncertainty about firms' abatement ability, taxes are better able to get the equilibrium right at once rather than iterating towards the equilibrium in multiple rounds of policy adjustment.

When there is uncertainty about the marginal damage rather than the marginal abatement cost, taxes and permits perform equally. Both instruments generate the same equilibrium outcome.

3.3 Environmental markets in action

The International Energy Agency (IEA), an energy forum for 28 advanced economies established in 1974, maintains a very useful database of energy efficiency policies and measures for each of its member countries, including Canada and the United States. The list of activities is remarkably long and consists primarily of subsidies, accelerated depre-

ciation allowances (known as a capital cost allowance in Canada's tax system), R&D incentives, education and outreach activities, and various voluntary technical standards. In the United States, some of the most recognized programs are the Energy Star program for products and home improvement (launched in 1992) and the CAFE standard for improving fuel efficiency of vehicles (launched in 1974).

In recent years, a number of jurisdictions have introduced a **feed-in tariff** for electricity to subsidize renewable energy such as solar and wind power. One of the most lucrative such programs exists in Germany. Operators of various types of renewable energy installations (hydro, biomass, geothermal, wind, and photovoltaics) receive a guaranteed feed-in tariff when they supply electricity to the grid. In 2009 a photovoltaic installation was paid €0.3194/kWh if installed on open land and €0.3300–€0.4301/kWh if installed on buildings (depending on installation size). The feed-in tariff was guaranteed contractually for twenty years. These feed-in tariffs are very generous when compared to the typical production cost of electricity from conventional sources. As a result, solar power and wind power blossomed in Germany, from rooftop solar panels on individual homes to large wind farms. Farmers have redefined 'farming' by turning their potato fields into solar farms.

Other jurisdictions have also pursued feed-in tariffs. Spain has guaranteed €0.2694/kWh for 25 years for solar thermoelectric installations. Between 2009 and 2011, Spain completed the installation of two 150 MW, four 100 MW, and eight 50 MW solar-thermal power plants. The largest solar-thermal installation is currently being built in the Mojave desert of the United States with a capacity of 354 MW. In Canada, the Ontario Power Authority offers a feed-in-tariff of \$0.42/kWh for photovoltaics and \$0.11/kWh for wind and biogas. These incentives serve a dual purpose. On one hand they are meant to build renewable energy capacity, while on the other hand they are meant to stimulate investment and innovation in the green-tech sector of the economy. In 2010, a 97 MW photovoltaic solar farm was built by Enbridge and First Solar near Sarnia, Ontario.

Permit markets

There are several emission permit markets in operation in the United States and Europe, albeit none in Canada at this point. One of the first emission permit markets was established through the Acid Rain Program in the United States. This cap-and-trade system was established under the 1990 *Clean Air Act* and sought to reduce sulfur dioxide emissions by 50% between 1980 and 2010. Trading commenced in 1995 and initially affected 263 units at 110 mostly coal-burning electric utility plants located in 21 eastern and midwestern states. The second phase of the program commenced in 2000 and eventually encompassed over 2,000 units, virtually all power plants with more than 25 MW capacity.

Allowances are grandfathered on an emission baseline (the average fossil fuel consumed from 1985 through 1987 multiplied by 1.2 pounds of sulfur dioxide per million British thermal units).

The state of Illinois has adopted a cap-and-trade system for volatile organic compounds (VOCs) for eight counties, named the *Emissions Reduction Market System*. It was launched in 2000 with over 100 participants. Ten northeastern states agreed on a cap-and-trade system called the *Regional Greenhouse Gas Initiative* (RGGI) for carbon dioxide emissions from power plants. This system was launched in 2009 with quarterly auctions of permits and is aimed at reducing emissions by 10% below 1990 levels by 2020. Allowance prices typically range between \$2 and \$3 per ton of carbon dioxide.

By far the largest permit market in the world is the *EU Emission Trading System* (ETS) launched in 2005. The ETS now operates in 30 countries (27 EU member states plus Iceland, Liechtenstein, and Norway) and covers carbon emissions from power stations, combustion plants, oil refineries and iron and steel works, as well as factories making cement, glass, lime, bricks, ceramics, pulp, paper, and board. The total amount of emissions from these industrial sectors were originally capped by each government separately. This led to an overallocation of permits during the first (2005–2007) phase of the trading system as governments were not careful coordinating targets. As a result of the overallocation, the permit price for the 2007 allocation crashed. During the third phase of the ETS, which started in 2013, allocations are determined directly at the EU level rather than the national level. Permit trades are carried out on private futures exchanges; the largest and most liquid market is the ICE Futures Europe exchange, previously known as the European Climate Exchange (ECX). Permit prices remain quite volatile. In 2012, prices for the second ETS phase target (2008–2012) dropped to €6–8 per tonne. In 2013, prices for the third ETS phase (2013–2020) dropped to under €3 at times but have since stabilized at about €5 per tonne.

Offset Markets

Markets for **carbon offsets** have sprung up either as voluntary schemes or as compliance markets following a government mandate. Offsets are meant to take carbon dioxide or other greenhouse gases out of the atmosphere in exchange for other entities putting them into the atmosphere. This only works if the carbon dioxide is truly sequestered, such as accumulation in trees or through underground injection. To make carbon offsets effective, several criteria need to be met effectively:

- **additionality**: the emission reductions would not occur without the offset (e.g., because they are financially worthwhile anyhow);
- **permanence**: the emissions are sequestered permanently as well as irreversibly;

- **verifiability:** the reductions can be verified and audited by an independent third party (which may be difficult outside one's own jurisdiction); and
- **no leakage:** implementing the emission reduction project does not lead to higher emissions elsewhere such as through substitution effects or international trade.

Many offsets today are based on 'emissions avoided' concepts; few are based on true sequestration. Nevertheless, some progress has been made certifying offsets through the ISO 14064 standard. Box 3.2 discusses the use of carbon offsets in British Columbia.

Several airlines offer their customers to buy carbon offsets for their flights. For a Toronto-Vancouver round-trip with an estimated 703 kg of carbon dioxide emissions, Air Canada passengers can purchase carbon offsets for between \$14-\$20/tonne from a forest restoration project, a landfill gas recover project, and a tire recycling program. These three projects are among the better ones in the offsetting industry as they involve, at least in part, true sequestration. As of December 2011, Air Canada's partner Zerofootprint had sold \$285,160 in (voluntary) offsets, equivalent to 17,749 tonnes of carbon dioxide. In 2010, Air Canada consumed a total of 3.378 billion litres of turbo fuel at 2.5 kg carbon dioxide per litre. This amounts to about 8.4 million tonnes of carbon dioxide. So less than 0.02% of its emissions were offset voluntarily.

A voluntary market for carbon offsets was launched as the Chicago Climate Exchange (CCX). However, this market crashed in 2009 as firms retreated from voluntary offsets in the wake of a major economic recession.

3.4 Valuing the environment

When businesses develop new mines or build new plants, pipelines, power transmission lines, or other infrastructure, they often impact the local environment. Regulators and businesses alike need to know the cost of such activities: regulators because they may want to assess the overall cost and benefits of such a project, and businesses because they may need to compensate anyone affected by the planned new activity. Environmental economists have developed approaches to valuing the environment, but this task is not easy because the environment does not come with a simple price tag.

There are two major approaches to valuing the environment and the services provided by the environment. One approach is based on measuring or estimating the market benefits of the environment. Because this approach is based on observable market data, it is called the **revealed-preference** approach. When such market data are not observable because there are no markets for the services of the environment,

Box 3.2: Carbon Offsets in British Columbia

In British Columbia, the provincial government mandated the public sector to become 'carbon neutral.' Because the public sector cannot suddenly stop heating schools and universities, these public sector entities have to buy offsets.

The Pacific Carbon Trust (PCT) that was established as British Columbia's clearing house for carbon offsets does not necessarily achieve actual sequestration. Instead, most offsets (or 'carbon credits') are based on hypothetical reductions against a business-as-usual scenario.

For example, consider supplementing the natural gas heating of a warehouse with solar power. The owner of the warehouse offers to instal solar panels in exchange for a payment equivalent to the reduced carbon dioxide emissions times a negotiable carbon price. As public sector entities have to pay \$25 per tonne for their carbon offsets, the PCT can afford to pay up to \$25 for carbon credits. So what is the catch? Perhaps the warehouse owner would have installed solar panels anyhow. This is a problem of **free riding**. If someone offers you money to buy a hybrid-electric Toyota Prius in order to reduce emissions, you might take the money even if you would have bought the Prius anyhow because it reduces your fuel consumption and fuel costs. If the proportion of free riding is large, offsets will become ineffective.

As the \$25/tonne carbon offset price is fixed for public sector entities, a supply-demand mismatch may occur. The PCT may run out of viable projects at a price below \$25/tonne to offset all public sector GHG emissions. Then the public sector will no longer be 'carbon neutral.' Alternatively, if the PCT can supply all necessary offsets at a price below \$25/tonne, the public sector entities are overpaying. Many current offsets are in fact available well below \$25/tonne.

The PCT refuses to make public the prices it pays for individual carbon offset projects. This secrecy and lack of accountability is disturbing and calls into question the veracity of claims made by the PCT. Nurturing an 'infant industry' is no excuse for secrecy. The PCT may be more interested in developing politically opportune showcase projects than in finding least-cost carbon offsets (which are likely to be found in afforestation and reforestation projects).

In November 2013, the British Columbia government announced that it will eliminate the PCT. However, the carbon offset program will continue within the Environment Ministry's Climate Action Secretariat. A revised program may expand an exemption to health agencies and post-secondary institutions that was adopted originally for elementary and high schools.

economists rely on estimating the non-market benefits of the environment. This is called the **stated-preference** approach.

The results from any of the valuation methods are often input into a **cost-benefit analysis** (CBA); see Boardman et al. (2010a) for a thorough treatment of this subject. CBA is a policy assessment method that quantifies in monetary terms the value of all consequences of a policy to all members of society. Fundamentally, CBA is based on the principle that a society can make economic trade-offs among its members; it allows for some people to win (gain benefits) and some to lose (incur costs) from a new project. CBA also allows for intertemporal trade-offs to be made, allowing for benefits and costs to materialize at different points in time (costs now and benefits later, or benefits now and costs later). Section 3.5 explores this topic further.

3.4.1 Revealed-preference approaches

Consumers reveal preferences through market transactions. This implies that a researcher can observe market prices (and quantities) and infer how consumers value a particular resource. Sometimes the market valuation is indirect because the use of an environmental resource may involve another activity that can be measured. One such indirect method is the **travel-cost approach**. For example, it is possible to measure the value that consumers place on the use of parks and other recreational sources by observing the time they are willing to travel to get there. As time is costly, the travel time and distance implies a travel cost.

A more direct method is the use of **hedonic regressions**, an econometric tool where an observed market price is explained statistically through a number of observable characteristics. This method tries to capture the change in prices of (complementary) goods to infer consumers' willingness to pay for a cleaner environment or some other environmental amenity. Hedonic regressions are widely used in real estate appraisals. For example, a hedonic regression was used to estimate how PCB contamination of sediments in the harbour of New Bedford, Massachusetts affected real estate prices. A study by Mendelsohn et al. (1992) showed that in 1982 real estate prices were depressed by \$9,000 in the most affected zone, controlling for numerous property characteristics, including size and quality of homes.

Real estate valuations are affected by a large number of environmental issues. For example, the construction of pipelines, electrical transmission lines, roadways, bridges, railway lines, or subway tunnels affects property values. For urban environments, using real estate valuations is often the most promising avenue for estimating environmental impacts.

Yet another revealed-preference approach uses the **value of a statistical life (VSL)**. The value of a statistical life captures the trade-off between mortality risk and other goods and services. It is defined as the ratio of the marginal utility of a small reduction in mortality risk to the marginal utility of a small change in income, and it is measured in dollars per unit of risk reduced. At first glance it may appear unethical to put a price on a life, but this is a profound misunderstanding of the VSL concept (Cameron, 2010). A common misconception about VSL is that it puts a price on human life, and even provides different figures for different groups in society (such as old and young people). VSL is not about putting a price on any specific life or person. Instead, it merely captures individuals' willingness to pay for reducing mortality risk. Keeping this important distinction in mind, VSL calculations can yield very insightful policy recommendations. Workers who pursue careers that involve a higher-than-average risk of death (such as police officers, firefighters, or coal miners) typically accept a wage premium that compensates them for this increased risk. Combined with actuarial statistics on death rates, the wage premium computes into a monetary measure, the VSL. Studies of wage premia from risky jobs put the value of a statistical life at about \$5 million (in current dollars of the reference period). Based on a review of numerous studies, the US Environmental Protection Agency currently uses a VSL of \$7.4 million in 2006 dollars.

A useful example of how the VSL can be used in practice is to look at various environmental regulations and calculate the implied VSL. This can be compared with what we know from other studies about VSL. Viscusi (1996) examined numerous US environmental regulations to determine the cost of a statistical life saved, yielding widely different estimates shown in table 3.1. Putting seat belts in cars saves about 1,850 lives per year in the United States, and at a cost of \$300,000 per life saved it is relatively cheap. By comparison, banning formaldehyde in the workplace costs about \$72 billion per life saved, which is hugely expensive compared to the \$5 million VSL used in that study. The line in the middle of the table marks that threshold, indicating that the policies above the line are sensible, while the policies below are too expensive and should be reviewed.

3.4.2 Stated-preference approaches

Often there is neither a direct nor indirect market through which consumer preferences can be gauged. Many environmental resources only exhibit non-market benefits. Such resources can exhibit **use value**, **option value**, and **existence value**. The total value of an environmental resource or feature is the sum of use, option, and existence value.

Table 3.1 The Value of a Statistical Life (VSL)

US regulation (selected examples)	Agency	Initial Risk (per year)	Expected Lives Saved	Cost/Life (\$ millions)
Unvented space heaters	CPSC	2.7 in 10^5	63	0.10
Airplane cabin fire protection	FAA	6.5 in 10^8	15	0.20
Auto passive restraints/belts	NHTSA	9.1 in 10^5	1850	0.30
Hazard communication	OSHA	4.0 in 10^5	200	1.80
Grain dust	OHSA	2.1 in 10^5	4	5.30
Arsenic/glass manufacturing	EPA	8.0 in 10^4	0.11	19.20
Arsenic/copper smelter	EPA	9.0 in 10^4	0.06	26.50
Asbestos	EPA	2.9 in 10^5	10	104.20
Formaldehyde in workplace	OHSA	6.8 in 10^7	0.01	72,000.00

Source: Viscusi (1996). Initial annual risk indicates annual deaths per exposed population. Agency denotes the US agency responsible for issuing the regulation; FAA: Federal Aviation Authority; NHTSA: National Highway Traffic Safety Administration; PHSA: *Public Health Services Act*; EPA: Environmental Protection Agency; CPSC: Consumer Product Safety Commission. Valuations are expressed in 1984 current-price US dollars.

Use value is based on consumers' use of an environmental resource. For example, a clean river is more suitable for swimming, boating, fishing, and for obtaining drinking water. Option value is intertemporal in nature. For example, preserving the Amazon forest may allow for the discovery of new medicines derived from yet-undiscovered plants or animals. Option value is rooted in preserving future use and limiting current use. Existence value looks at benefits that are intrinsic to the existence of the resource, whether it is scenery, a plant or animal species, or a historic artifact. For example, an endangered species such as polar bears may have an aesthetic value (they are cute to look at) or an ecosystem value (they contribute ecosystem functions to habitat or other species).

Stated preferences are captured through a method known as **contingent valuation**; see Alberini and Kahn (2009) for a thorough discussion of this topic. Contingent valuation is a survey-based technique for the valuation of non-market resources that provide utility to people. Whenever there is a lack of a market due to the inability to transact in an environmental resource, survey-based methods may be able to help. Surveys ask people about their **willingness to pay** (WTP) to maintain or preserve an environmental feature or their **willingness to accept** (WTA) compensation for degrading or losing that environmental feature. In theory, both questions should elicit the same response. In practice, that is not the case, and the stated willingness to pay is smaller than the willingness to accept for the same environmental feature. The $WTA > WTP$ paradox is a central feature of contingent valuation. Economic theory is trying to explain this paradox.

The willingness to pay or accept is closely linked to the economic concept of **consumer surplus**, which is defined as the difference between what one is willing to pay for a good or service and what one is actually paying for the good or service. With common access, the price for most environmental resources is zero. The price that people are willing to pay for a public good such as an environmental amenity is close to the increase in consumer surplus they enjoy. WTA should be a little higher than consumer surplus because people are better off with the compensation payment, whereas WTP should be a little higher than consumer surplus because people are made poorer by paying to preserve the environmental amenity.

The $WTA > WTP$ paradox goes beyond the small deviations from consumer surplus that can be explained by a wealth effect. For example, the estimated willingness to pay for remedying a contaminated landfill is typically much less than the willingness to accept compensation for enduring the environmental hazards from the contaminated landfill.

Economists offer two leading explanations for the $WTA > WTP$ paradox. One is based on **prospect theory**, and the other is based on poor substitutability. Prospect theory was developed by behavioural researchers, most prominently Nobel laureate Daniel Kahneman and his co-researcher Amos Tversky. They demonstrated that people suffer from **loss aversion** when making economic decisions. People tend to weight losses larger than an equal amount in gains; losing \$1,000 is psychologically more intense than gaining \$1,000. Prospect theory suggests that utility is not an absolute measure (as used in neoclassical economics) but a relative measure. Consumers adopt the status quo as their reference point when evaluating (expected) gains and losses. Because of loss aversion, the marginal benefits of environmental cleanup tend to be smaller than the marginal benefits of accepting environmental degradation.

Another explanation may be rooted in the poor substitutability of monetary compensation and non-tangible costs such as the risk of death or environmental degradation. If an individual faces income constraints, he or she might not be able to pay for a large reduction in the risk of death from, say, an environmental contamination, the WTP will be low, whereas the individual may demand a very large amount in compensation for being exposed to that risk, which implies a high WTA.

Whether WTA or WTP questions are used on a contingent valuation survey also depends on how the property rights are assigned. When people own the right to an environmental amenity, it makes sense to ask them about their WTA. If, however, they do not own the environmental amenity, it makes sense to ask them about their WTP to preserve it.

Carrying out a contingent valuation survey carries with it numerous pitfalls inherent in all surveys. Among them is the risk that people articulate **protest answers** that exaggerate their true WTP or WTA. Sim-

ilarly, respondents may engage in strategic behaviour by exaggerating their true WTP and WTA if they believe that their answer on the survey will influence the decision in their favour; this leads to **strategic bias**. Another common problem is **response bias** in surveys, where respondents tell the pollster what they think they are expected to respond rather than their true belief. If the respondent believes that the survey is conducted on behalf of a particular interest group, they may be more inclined to articulate a position favourable to that interest group. People are often conflict averse and vary their opinion depending on whom they talk to. There is also the possibility of **free riding** when asked about WTP: why offer to pay anything when others are willing to pay?

Contingent valuation surveys also suffer from something called **embedding bias**. Answers to survey questions may depend on the context in which they are provided. The WTP may depend on the order in which two or more projects are embedded in the survey. This is a question of sequencing and nesting that can sometimes be overcome by asking questions in random order, or by varying the ranking of competing options in random order from one survey participant to the next. Surveys can be compromised by asking a number of leading questions that may influence the critical WTP or WTA question. A related problem arises from asking multiple WTP or WTA questions for partial but additive policies. The sum of the partial WTPs or WTAs may exceed the WTP or WTA for the joint policy.

If contingent valuation surveys are fraught with problems, why bother? The answer is simple: in the absence of market prices, surveys provide the only way to elicit useful valuation information. The key is therefore understanding, preventing, or correcting the various survey biases discussed above. Randomization of survey questions and avoidance of leading questions can help address several of the problems. The usefulness of contingent valuation is self-evident; box 3.3 illustrates a typical application.

3.5 Cost-benefit analysis

There are two major types of **cost-benefit analysis (CBA)**. Ex-ante CBA is more common and involves assisting in a decision about implementing a proposed project. Ex-post CBA is performed during or after a project's life and informs the decision about continuing or repeating a project.

Box 3.3: The Glen Canyon Dam and Recreation



Image by Christian Mehlführer, reproduced under Creative Commons License.

The Colorado River has been dammed extensively. Lake Powell was created through building the Glen Canyon Dam in 1956–66, and Lake Meade was created through building the Hoover Dam in 1931–36. A large share of water is diverted to irrigation and to nearby cities. The Colorado River used to flow into the Gulf of California in Mexico, but the river delta has shrunk to 5% of its original size.

The Glen Canyon dam provides peak-load power for residents of Nevada and California, which adversely affects downstream ecosystems in the Grand Canyon and reduces the quality of recreational rafting along the canyon. The US Bureau of Land Reclamation and the National Park Service were interested in assessing the value of recreational rafting compared to the market value of peak-load power supply. The study was meant to inform a decision about

modifying the water flow from the dam that would permit greater recreational use of the Colorado River. The two agencies conducted a contingent valuation survey and found substantial economic values for rafting with increased water flows worth \$2 million per year. This survey helped inform the *Grand Canyon Protection Act* of 1992.

Further contingent valuation surveys explored the willingness to pay for different flow regimes. The lack of the natural high spring water flows threatened endangered fish species, native vegetation, and birds. The survey revealed strong support for a more natural flow regime. This resulted in changes to the flow management, including large spills during the spring of 1996 to emulate the natural high spring flows. These large spills were repeated several times, most recently in 2004 and 2008.

CBA decision rule: Among a set of candidate projects or options, a decision maker should choose the project or option that maximizes the **net social benefit** (benefits minus costs) expressed in present-discounted terms.


The use of net social benefit in cost-benefit analysis is crucial. Net social benefit is the difference between social benefits and social costs, expressed in dollar amounts. Policy efficiency is a different concept: the ratio of benefits to costs. One policy may be more efficient than another but may still be outranked if the other policy's net social benefit is larger.

The four panels in figure 3.8 illustrate the key concepts of CBA. In panel A, a particular policy is evaluated with a benefit of \$50 million and a cost of \$40 million. The net social benefit is therefore \$10 million. However, the benefit-cost ratio of 1.2 is very low. Many parameters that were estimated to obtain the two point estimates for benefit and cost are 'noisy' and therefore may vary over a certain range. To reflect this uncertainty, it is customary to draw an **error ellipse** around each cost-benefit point estimate, based on the standard error of each of the two point estimates. In panel A, both the 50% and 95% error ellipses cut across the efficient frontier and suggest that it is possible that the policy may be inefficient—i.e., costs may exceed benefits. This policy would be considered marginal, and regulators should probably consider alternative policies that are unambiguously efficient.

In panel B, policies A and B are both efficient and the benefit-cost ratio greatly exceeds 1. Policy A has a net benefit of $\$150 - \$20 = \$130$ million, while policy B has a net benefit of $\$100 - \$40 = \$60$ million. Clearly, policy A outranks policy B both in terms of net social benefit and efficiency. The situation in panel C demonstrates that two policies can be equally efficient but still dominate each other in terms of net social benefit. Policy A has a net benefit of $\$90 - \$30 = \$60$ million, while policy B has a net benefit of $\$150 - \$50 = \$100$ million. Both policies have an equal benefit-cost ratio of 3:1 and are thus equally efficient. However, policy B generates \$40 million more in net social benefit than policy A.

Panel D in figure 3.8 illustrates the most difficult case. The policy ranking is ambiguous. Policy A has a net benefit of $\$120 - \$52 = \$68$ million, while policy B has a net benefit of $\$100 - \$40 = \$60$ million. Policy B is more efficient than policy A because its benefit-cost ratio of 2.5 is larger than the benefit-cost ratio of 2.3 for policy A. However, policy A generates \$8 million more in net social benefit. Should the regulator therefore adopt policy A? The error ellipses of the two policies overlap a fair bit, and it may not be possible to say with confidence that policy A is better than policy B. There is a significant region where both policies may deliver the same outcome. In this case, the regulator may

Fig. 3.8 Cost-Benefit Analysis and Uncertainty



in fact prefer policy B, which is 23% cheaper than policy A. This example demonstrates how important it is to carry out a rigorous analysis of the parameter uncertainty in the cost-benefit model. Simply basing a policy decision on the 'best' point estimates could lead to expensive policies that do not deliver a superior outcome.

Often the most important challenge in cost-benefit analysis is the time horizon. Costs may need to be incurred immediately while benefits accrue over many years. Table 3.2 provides the discount factors that correspond to particular discount rates and time horizons.²

² For a given discount rate r , the discount factors are $[1 - (1 + r)^{-T}]/r$ for the finite time horizon T and $(1 + r)/r$ for the infinite time horizon.

Table 3.2 Cost-Benefit Analysis and Discount Factors

Time horizon [Years]	Discount rate										
	2.0%	2.5%	3.0%	3.5%	4.0%	4.5%	5.0%	6.0%	7.0%	8.0%	10%
10	8.98	8.75	8.53	8.32	8.11	7.91	7.72	7.36	7.02	6.71	6.14
20	16.35	15.59	14.88	14.21	13.59	13.01	12.46	11.47	10.59	9.82	8.51
50	31.42	28.36	25.73	23.46	21.48	19.76	18.26	15.76	13.80	12.23	9.91
∞	51.00	41.00	34.33	29.57	26.00	23.22	21.00	17.67	15.29	13.50	11.00

For example, a project or policy with a 20-year time horizon and a discount rate of 3.5% has a discount factor of about 14. A project or policy that costs \$100 million today (and nothing henceforth) would need to generate annual benefits of $\$100/14.21 = \$7.04$ million in each of the 20 years to be considered worthwhile.

The importance of the choice of discount rate is best seen in the case of an infinite time horizon as shown in the last row of table 3.2. Suppose a particular environmental policy generates annual benefits of \$100 million in perpetuity. At a discount rate of 10%, this amount is equivalent to 11 times \$100 million today, or \$1.1 billion. At a discount rate of 4.0%, the same policy is valued at \$2.6 billion dollars today. The regulator with the 4% discount rate would be willing to spend \$2.6 billion to implement the policy, while the regulator with a 10% discount rate would only be willing to spend less than half of that amount. While economic analysis can guide us to a reasonable choice of discount rate (Boardman et al., 2010b; Moore et al., 2013), there remains great uncertainty about the choice of the appropriate time horizon. How long will a policy remain in place unchanged to generate the continual benefits? Or how long will an infrastructure project last before it will need to be replaced? These uncertainties make it necessary to engage in a rigorous analysis of possible scenarios and develop appropriate error ellipses around the cost-benefit point estimates.

If one does not want to commit to using a particular discount rate, there is another way of analyzing a cost-benefit scenario. One can ask which discount rate would drive the net social benefit to zero. This critical rate at which present-discounted benefits equal present-discounted costs is known as the **internal rate of return**. Calculating this quantity requires numerical approximation techniques. The decision to proceed with a particular project is made when the internal rate of return exceeds the cost of capital—i.e., the interest rate at which funds can be borrowed. This tool is less useful for evaluating an environmental policy because such policies may not involve investments. The internal rate of return is also not useful for ranking competing projects; it is only a tool for analyzing a single project.

3.6 International trade repercussions

Environmental policies can be affected by repercussions from a country's international trade regime. An open economy is vulnerable to opportunistic behaviour by firms. Strict environmental policies could motivate companies to shift production outside the regulated jurisdiction. This phenomenon is called the **pollution haven effect**.

It has been argued that poorer countries tend to maintain lower environmental standards than richer countries. Therefore, poorer countries may become pollution havens in response to ever stricter environmental regulations in rich countries. Theoretical and empirical economic work has demonstrated that the effect is likely small because of an offsetting effect. Richer countries also tend to be more capital intensive, and because capital-intensive industries are also pollution intensive, richer countries therefore tend to be more pollution intensive. This **factor endowment effect** works in the opposite direction of the pollution haven effect. The two sources of **comparative advantage** work against each other. Which one dominates depends on their relative strength and thus varies from country to country. Trade liberalization shifts the composition of industrial activity in a country, and thus the composition effect of trade liberalization depends on whether a country's comparative advantage tilts towards pollution-intensive industries. In addition to the **composition** effect of trade liberalization, the overall gains from trade also boost a country's income and thus generate a **technique effect** whereby richer countries adopt stricter environmental rules. The idea behind this effect is that the environment is a normal good, and that when societies grow richer they will demand more of the 'clean environment' good. Empirically, this technique effect tends to dominate the picture for a pollutant such as sulfur dioxide. However, the above statements do not apply in the same way for global pollutants such as carbon dioxide.

When a country adopts a carbon policy, there is a considerable risk of **carbon leakage**. When one country reduces carbon emissions as a result of adopting a related carbon policy, another country may increase emissions due to goods production moving from the country with the carbon policy to the country without the carbon policy. There are two transmission channels for such carbon leakage:

- **Supply-side leakage** occurs through the physical relocation of production facilities to escape pollution control regimes. This is the pollution haven effect discussed above, and is influenced strongly by a country's comparative advantage. Supply-side leakage is bounded by the cost of relocating production, plants, or entire industries to a foreign country. Furthermore, international trade adds transportation costs that may dampen the relatively small cost advantages from the absence of a carbon policy. When pollution is local rather

than global, relocation may not necessarily be a bad thing. If the firms migrate to areas with lesser population density, negative externalities may actually be reduced.

- **Demand-side leakage** occurs when local environmental regulation depresses demand for a polluting good in one region (by increasing its price), and this demand shifts to lower-priced imports from an unregulated region.

Leakage is a phenomenon that applies not only to carbon policies but potentially to any other environmental policy. Nevertheless, the potential of leakage is always associated with global pollutants because for global pollutants only the world total emissions matter, but not where the emissions occur.

Yet another international trade phenomenon is known as **reshuffling**. This phenomenon can also occur among different states and provinces within a country. Consider regulations that are imposed at the point of purchase so that consumers in one country or region are subject to the policy and others are not. For example, assume that California accounts for 10% of the world's sea bass market, and that California adopts regulation that only sea bass caught using sustainable fishing techniques can be sold there. If more than 10% of sea bass in the world market is *already* caught using sustainable methods, these fish can be diverted for sale in California. The California regulation may not have any effect at all on how fish are caught. Reshuffling becomes more likely when the share of products that already comply with an environmental policy is larger than the share of consumers who are subject to it.

Another example is a firm's decision to contract with a wind farm operator to supply 100% of its electricity from clean renewable energy, albeit delivered through the conventional electric grid. While it may appear that the firm is boosting demand for renewable energy, the wind farm operator would have sold the electricity to the grid anyhow. Electricity merely got reshuffled from the average consumer to the particular firm. In fact, electrons do not carry identification tags, and thus once electricity enters a distribution grid it is anonymous. There is no way of earmarking electricity. The firm's commitment to renewable energy is thus meaningless.

3.7 Summary

Environmental economics provides an analytic framework for designing environmental policy interventions and making decisions about the relative merit of different environmental policy options. Generally, incentive-based policy interventions such as environmental taxes and cap-and-trade systems are more efficient economically than mandated

policy interventions such as technical standards. No environmental policy is superior along all relevant dimensions, and thus a strong case can be built for hybrid environmental policies that mix elements from different policies. Long-term considerations are particularly important when considering the effect of environmental policies on innovation in pollution abatement and pollution prevention technologies. Regimes that provide predictable and stable innovation incentives (in particular, environmental taxes) have a significant advantage.

Valuing environmental amenities and natural resources can be accomplished through revealed-preference approaches based on market values, and when this is not feasible, through stated-preference approaches that are based on contingent valuation surveys. Results from such an analysis can be incorporated into a cost-benefit analysis. The CBA framework suggests that among competing projects the one with the highest total net benefit should be chosen. When two or more projects are statistically indistinguishable as top contenders, the lowest-cost project should prevail.

3.8 Study questions

1. In the presence of a negative environmental externality, how does the market equilibrium without a policy differ from a market equilibrium in which the externality has been priced correctly?
2. What is bioaccumulation, and what is biomagnification?
3. What does the 'Coase theorem' teach us about achieving efficient environmental outcomes and about the initial allocation of property rights?
4. Why do market-based interventions (green taxes, cap-and-trade) tend to deliver more efficient outcomes than mandated interventions (standards, technology mandates)?
5. What are key criteria for choosing among different environmental policy instruments?
6. Why is no single environmental policy instrument superior along all dimensions relevant to policy choice? Compare emissions standards with a green tax.
7. Explain how a green tax or a cap-and-trade system leads to distinct *scale effects* and *technique effects* for individual firms.
8. When you compare a subsidy for pollution reduction and a tax on pollution emissions, are *scale effect* and *technique effect* working in the same direction?
9. In the presence of firm heterogeneity in abatement technology (and thus marginal abatement cost), what is the optimal way of allocating abatement effort across firms, and which policy instruments can achieve this efficient outcome?

10. Does the method of issuing emission permits in a cap-and-trade system (auctioning or grandfathering) affect the efficient outcome?
11. What are the (political) trade-offs for allocating emission permits by way of auctioning or grandfathering in a cap-and-trade system?
12. What is the benefit of using *permit banking* and *permit borrowing* in a cap-and-trade system?
13. What are the practical problems associated with the European Emission Trading System? What works, what doesn't?
14. Discuss the advantages of a hybrid cap-and-trade system that includes a price floor and price ceiling, or an allowance reserve.
15. What is 'carbon leakage' and why should policy makers be concerned about it?
16. What is 'reshuffling' and why should policy makers be concerned about it?
17. Why are green taxes superior to a cap-and-trade system when a policy maker is concerned about providing incentives for innovation into pollution abatement technology?
18. Consider Germany's generous feed-in-tariff (FiT) for electricity from photovoltaic systems on house roofs and on farmers' fields. What is economically sensible about a FiT subsidy, and what is not?
19. Describe key approaches for valuing environmental assets based on market benefits (revealed preference) and non-market benefits (stated preference).
20. When using *contingent valuation* surveys, why do questionnaire respondents often indicate a greater willingness to accept compensation for environmental degradation than a willingness to pay to prevent the environmental degradation?
21. What are *hedonic regressions* and how are they used?
22. How does one calculate the 'value of a life,' and how can one use such calculations to identify the 'low hanging fruit' of environmental policies?
23. In cost-benefit analysis, what objective is being maximized: the benefit-cost ratio or the net benefit?
24. How can one compare different options in cost-benefit analysis when point estimates of costs and benefits are subject to substantial measurement error?
25. What is captured by the measure 'value of a statistical life'? How can it be used to inform policy decisions?