Background on Basic Environmental Economics Concepts*

Jacob Bradt†

September, 2021

0. Introduction

This is a primer on the economics of the environment and climate.[1] The purpose of this primer is to provide further information on the basic economic concepts which are fundamental to understanding the economic perspective of environmental and climate policy. The concepts discussed herein are either directly covered or mentioned in Professor Robert Stavins' Harvard Kennedy School Executive Education course, "Climate Change Policy: Economics and Politics" — this document will therefore be useful to participants in the course who have had relatively little previous exposure to economics in an academic setting.

1. Economic Efficiency

Much of environmental policy—perhaps even all public policy—can be thought of as asking the general question of "How much environmental protection should society undertake?" An intuitive approach to answering this question involves comparing benefits and costs, and finding where their difference is greatest. To develop this intuition, consider the case of cleaning your home: the benefit of doing so is that you get to enjoy a clean home; however, this comes at the cost of the time and effort you spend doing so. Do you choose to maintain a spotless home? Some might, but for many there comes a point at which the additional cost of maintaining such a level of cleanliness is not worth the benefit they receive.

This intuitive approach of comparing benefits and costs is a central concept in economics, that of economic efficiency **economic efficiency**. From an economists' perspective, an efficient policy or outcome is one that maximizes—i.e., achieves the greatest possible—*net benefits*. Note that this definition is different from other common uses of the term "efficiency," which is often used to describe an outcome that minimizes effort, energy, or some other input conditional on some fixed output.[2]

To better understand the concept of economic efficiency in the context of environmental and climate policy, consider the real-world issue of carbon-dioxide (CO_2) emissions from fossil-fueled electric power plants. We can frame the policy issue either in terms of the efficient level of CO_2 emissions

---

*Prepared for Harvard Kennedy School Executive Education Program, "Climate Change Policy: Economics and Politics"

†`jbradt@g.harvard.edu`

[1] This primer is largely based on material from Professor Robert Stavins' degree program course as well as Nathaniel Keohane and Sheila Olmstead's text *Markets and the Environment* [Keohane, N.O. and S.M. Olmstead. (2007). *Markets and the Environment.* Island Press.].

[2] Consider the following example from Keohane and Olmstead (2007): imagine you are choosing between an energy-efficient, top-of-the-line air conditioner that costs $500 and a model that uses more electricity but only costs $150. While the more expensive air conditioner is more energy efficient, it may not necessarily be more efficient from an economic point of view—this will depend on how often you use the air conditioner and the price of electricity.

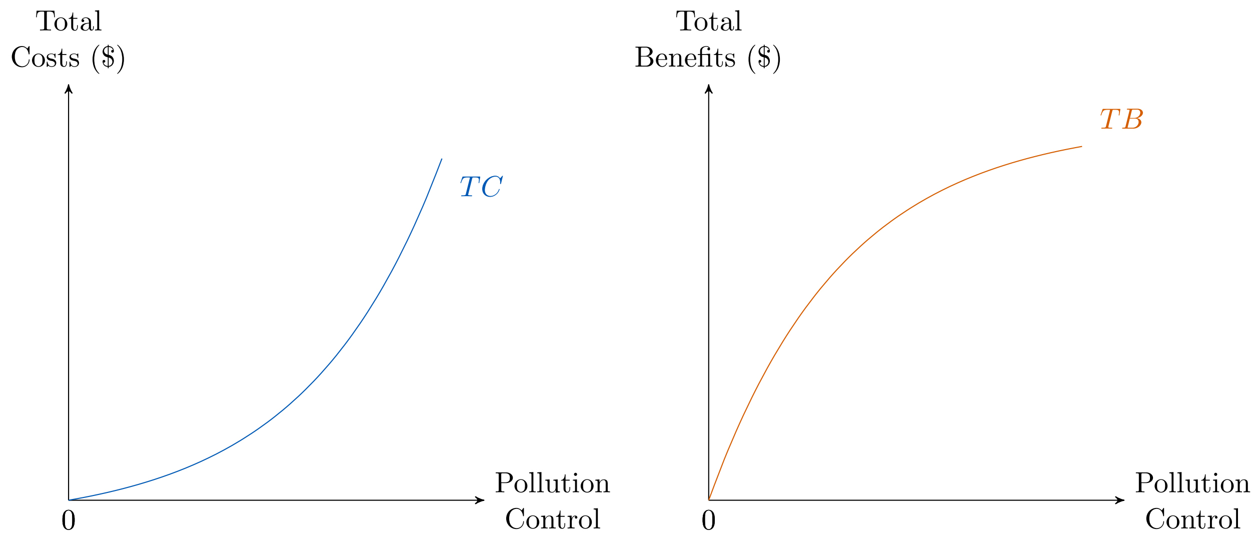

(a) Total costs of pollution abatement/control as a function of the level of pollution abatement/control.

(b) Total benefits of pollution abatement/control as a function of the level of pollution abatement/control.

Figure 1: The total costs and benefits of pollution control. The general pattern of costs increasing at a growing rate and benefits increasing at a slowing rate is nearly ubiquitous within the realm of environmental and climate policy.

or the efficient level of CO_2 emissions abatement; however it is often easier to think in terms of abatement or pollution control, which is a “good” (meaning we want more of it) than pollution, which is a “bad” (meaning we want less of it). This primer will adopt the convention of thinking about environmental policy from the perspective of the optimal level of pollution control.

To answer the question of what is the efficient level of CO_2 emissions abatement for fossil-fueled electric power plants requires thinking systematically about the costs and benefits of pollution control. Consider first the costs of CO_2 abatement. It is generally the case that the total cost of CO_2 abatement—and most forms of abatement for that matter—increases at an increasing rate for individual firms and, as a result, an industry as a whole. Why is this? Consider the case of a coal-fired power plant: it is reasonable to imagine there being some minor level of CO_2 abatement that the plant can achieve simply by improving the operation of the plant; however, for the plant to achieve substantial levels of CO_2 abatement, it might have to invest in, say, carbon capture and sequestration technology, which can come at substantial cost. Further reductions beyond this might require the plant to limit its operations or shut down entirely, an even more costly proposition. As this thought exercise demonstrates, the total costs of pollution abatement to firms and, when aggregated up, to industries generally increases at a growing rate. Figure 1a depicts this phenomenon graphically, with total costs in monetary terms on the y-axis and the level of pollution control on the x-axis.

Next consider the benefits of CO_2 abatement. The benefits of reducing greenhouse gas emissions can be thought of as corresponding to the avoided damages from global climate change. Generally, as the amount of pollution increases, environmental damages tend to increase at a more rapid rate. In the case of CO_2 emissions, minor emissions levels might be associated with small changes in precipitation patterns, whereas major emissions levels might result in more profound impacts such

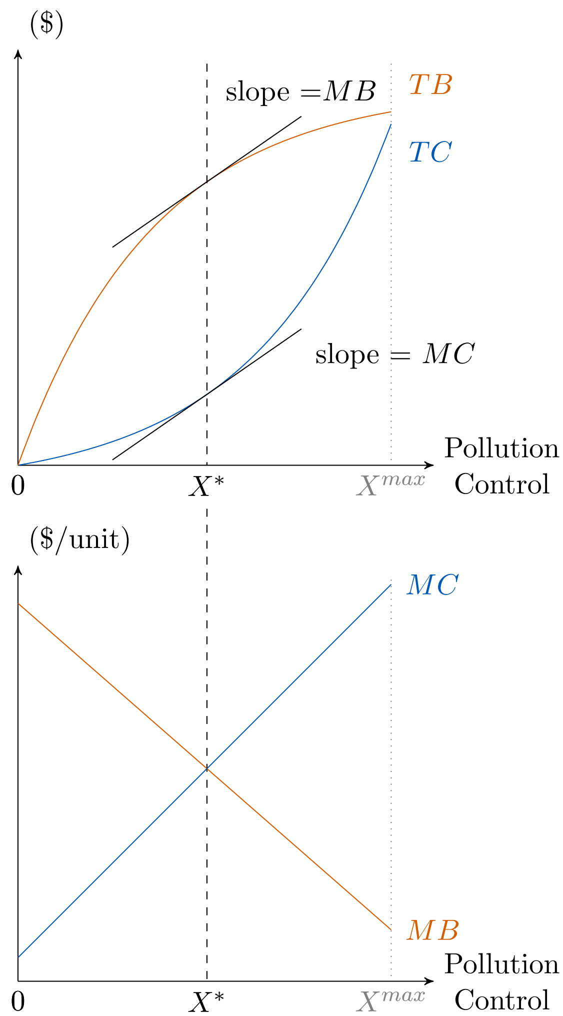

Figure 2: Putting it all together: the efficient level of pollution control or abatement (denoted as X^*) achieves the greatest possible net benefit. This occurs where the vertical distance between the total benefit and total cost curves is greatest, which is also where the marginal cost and marginal benefits are equal (which is often referred to as the *equimarginal rule*). X^{max} is the maximum amount of pollution control possible, which occurs when all pollution is eliminated.

as the melting of ice caps and resulting increases in global sea levels. This pattern of damages corresponds to total benefits from CO_2 emissions control that increase rapidl when abatement is low (and pollution is high) and increase more slowly when abatement is high (and pollution is low). Figure 1b depicts this phenomenon graphically, again with total costs in monetary terms on the y-axis and the level of pollution control on the x-axis.

With this real-world intuition for the pattern of the total benefits and costs of pollution abatement, we can now answer the question at hand: what is the efficient level of CO_2 abatement? To reiterate, this is the level of pollution control which maximizes **net benefits**, which we can note are just total benefits minus costs. It is easiest to identify the efficient level of pollution control graphically: combining the graphs of total benefits and costs into a single figure as is done in the upper panel of Figure 2, we can see that net benefits at each level of pollution control are represented by the vertical difference from the total benefit curve down to the total cost curves at that point. As abatement increases from a low level, the benefits increase more rapidly than do the costs, resulting in *increasing* net benefits. As more and more abatement is chosen, the benefits rise less rapidly, while the costs of abatement increase; these are the patterns we discussed above. Eventually the benefits increase more slowly than the costs and the net benefits *fall* as more and more abatement is done.

Given this pattern of increasing net benefits at low levels of pollution control and decreasing net benefits at high levels of pollution control, there is clearly an intermediate point at which the difference between benefits and costs must reach a maximum. This point is represented in Figure 2 as X^*, which is by definition the efficient level of pollution control. What can we learn from Figure 2? In addition to showing what maximizing net benefits might look like graphically in a general pollution control context, we can see that the efficient level is both greater than zero and less than the maximum possible level of abatement. This is a function of the general patterns we previously described about how total benefits and costs vary with the level of pollution control.

An alternative way to describe the benefits and costs of pollution control is in terms of **marginal costs and benefits**. Marginal cost can be defined as the cost of an incremental unit of abatement: if we abate 10 units of pollution, then the marginal cost is the cost of the 10th unit of pollution. Similarly, marginal benefit refers to the benefit from the final unit of abatement. We can think of marginal cost and benefit in our graphical framework as the slope of the total cost and benefit curves, respectively. Why is this? because they measure the cost and benefit of one additional unit of abatement.

As noted previously, it is common—if not ubiquitous—in the context of environmental pollution that total costs of pollution control increase at an *increasing* rate and total benefits increase at a *decreasing* rate. What does this mean? The incremental cost of pollution control increases with the total amount of control and the incremental benefit of pollution control decreases with the total amount of control. This is equivalent to saying that marginal costs increase and marginal benefits decrease with the overall level of pollution control. This is exactly what is depicted in the lower panel of Figure 2.

Examining Figure 2, we can see that it is drawn so that the marginal cost and benefit curves intersect at X^*, the efficient level of pollution control. Why is this? Well to the left of X^*—at higher levels of abatement—the benefits of pollution control rise faster than the costs. This means that the marginal benefits are greater than the marginal costs. To the right of X^*—at higher levels of abatement—the costs of pollution control rise faster than the benefits. This means that the marginal costs are greater than the marginal benefits. Putting these observations together, we see that at the efficient level of abatement, the benefit and cost curves must have the same slope.

This is known as the **equimarginal rule**: the efficient level of abatement, X^*, occurs where marginal benefit equals marginal cost: $MB(X^*) = MC(X^*)$. The equimarginal rule says that the efficient level of pollution control is where the extra benefit of the last unit of abatement undertaken equals its extra cost. Beyond that point, the additional costs of further abatement will outweigh the benefits. Thus, we can state our definition of an efficient outcome or policy in two equivalent ways:

1. An efficient outcome/policy maximizes net benefits.

2. An efficient outcome/policy equates marginal benefits and costs.

2. Dynamic Efficiency and Net Present Value (NPV)

The above definition of an efficient outcome or policy—one which maximizes net benefits/equates marginal benefits and costs—applies easily to a case where the costs and benefits are incurred at a single point in time. What happens when the costs and/or benefits of an outcome or policy occur over time, as is often the case? Then we must introduce the additional rules of **dynamic efficiency**.

Whereas in a static setting (a setting where benefits and costs occur a single time period) an efficient policy or outcome maximizes net benefits, in a dynamic setting (a setting where benefits and/or costs occur in multiple time periods) an efficient policy or outcome maximizes the **present value** of net benefits to society. This means that in a dynamic setting, we must convert all of the benefits and costs of a potential environmental policy into their dollar value *today* before summing them up. This ensures that we are using a common measure to compare costs and benefits across time.

Why might the value of benefits and costs vary over time? Consider a thought experiment: if you were asked whether you prefer $100 today or the same amount 50 years from now, you would almost certainly prefer the money today since you could invest it and earn a return over the next half century. The intuition underlying this thought experiment shows what is known as the time value of money, which is the reason that we must **discount** benefits and costs expected to occur in the future when considering the dynamically efficient policy. Discounting assigns weight to present versus future payoffs and in so doing converts monetary values to a common unit: present value. Once we discount streams of future costs and benefits into their present value, we can then identify the efficient policy as that which maximizes the difference between the present value of benefits and costs, also known as the **net present value**. We can write this mathematically as

$$NPV = \sum_{t=0}^{t=T} \frac{B_t - C_t}{(1+r)^t}$$

where t indexes time periods, B_t are the total benefits in period t, C_t are the total costs in period t, and r is a **discount rate**. What is this formula doing? The numerator calculates the net benefit in a given period and the denominator applies the appropriate discount factor that converts dollars in that period to their present value. The symbol $\sum_{t=0}^{t=T}$ indicates that we are summing the present value of net benefits from period $t = 0$ to period $t = T$.

Consider a simple example. Imagine that there is a CO_2 emissions abatement policy that will cost $100 today but yields $50 in benefits over the next 20 years. Assuming a discount rate of 3%, the NPV of such a policy is calculated as

$$NPV = \sum_{t=0}^{t=20} \frac{B_t - C_t}{(1+0.03)^t} = -100 + \frac{50}{(1+0.03)^1} + \frac{50}{(1+0.03)^2} + \cdots + \frac{50}{(1+0.03)^{20}} = \$643.87$$

What does this tell us? For $100 invested today, we receive today's equivalent of more than $600.

What is the correct discount rate, r, to use when calculating NPV? The perhaps unsatisfying answer is that *it depends*. For analysis of private sector projects, the relevant discount rate is the opportunity cost of capital (funds) to the specific firm. In the case of public sector projects or policies, we can define the **social discount rate**, which is the relative valuation placed by society on future consumption that is presently sacrificed. Unfortunately, there is no simple, unanimously accepted, and conceptually-correct approach to identifying the correct social discount rate. This remains an active area of research, though in the U.S. a common discount used for public sector analysis is 3%.

3. Additional Welfare Concepts

Economic efficiency is concerned with the overall net benefits to society from a policy. This says nothing about who gains and loses from a policy. For instance, a policy aimed at reducing SO_2 emissions from coal-fired power plants may benefit outdoor recreationists due to the resulting reduction in acid rain; however, it may very well harm communities that mine sulfur-rich coal.

In simple, common terms, efficiency is about maximizing the size of the pie, while **distributional equity** is about who gets what share of the pie. This metaphor highlights a potential conflict between efficiency and distributional equity: efficiency ignores distributional impacts entirely! To demonstrate why efficiency should remain a benchmark for setting policy, it is necessary to introduce two additional concepts from welfare economics:

- **Pareto efficiency**: the concept of "Pareto efficiency" can be seen as an attempt to evaluate policies based both on the size and share of the pie. A policy is "Pareto efficient" if—and only if—no member of society could be made better off by an alternative policy without making at least one person worse off.

- **Kaldor-Hicks criterion**: under the "Kaldor-Hicks" or "potential Pareto" criterion, a policy should be undertaken if the sizes of the resulting gains and the sizes of losses are such that those who gain could fully compensate those who lose and still be better off themselves.

As a means of evaluating policies, the concept of Pareto efficiency has some appeal; however, when it comes to applying this criterion in the real world, it becomes clear that it is much to strict: using this as a guide gof policy would almost always favor the status quo. Take the case of climate policy: some level of greenhouse gas abatement has clear and likely substantial benefits to many; however, there are clearly certain communities (for example, communities that employ individuals in the oil and gas industries) who might lose as a result of such a policy. If we are to only undertake Pareto efficient policies, there is very little that would change from the status quo in most settings, environmental and climate policy included.

This is why the Kaldor-Hicks or potential Pareto criterion is appealing in practice. As long as a policy creates sufficient gains relative to the status quo to in theory compensate those who lose as a result of the policy while still making everyone better off, then it is admissable according to this criterion. What does this criterion really say? The winners from a policy can only compensate the losers (even in theory) if the benefits from the policy are greater than the losses. Thus, a policy satisfies the potential Pareto criterion if and only if it produces greater net benefits than the alternative. Satisfying the Kaldor-Hicks criterion is equivalent to maximizing net benefits.

4. Understanding and Measuring Costs and Benefits

Finding the efficient outcome or policy requires knowing the costs and benefits of the outcome or policy. First consider the issue of defining and measure costs. In economic terms, the true costs of any activity are the **opportunity costs**: what you give up by doing one thing instead of another. For example, the true cost of attending an executive education program course is not simply the fee that you pay to enroll, but also the value of the next best use of the foregone time when attending the program. More broadly, the prices of inputs to production reflect their values in alternative uses: the production of electricity requires capital to pay for the construction of the generating unit, labor to operate the plant, and fuel for the generating unit, all of which could have been put to alternative uses. This principle applies to reducing polution: removing SO_2 from a power plant's emissions requires capital to build the scrubber and labor and materials to operate it. Devoting these resources to pollution control results in less time spent on other endeavors, such as improving operating efficiency.

Two general approaches are used to measure the direct costs of pollution control. The first uses information that firms from industries that are required to report cost data to the government. The second approach uses data on revenues and production costs to estimate the foregone profits or diminished productivity associated with environmental policy.

Outside of the direct costs of pollution control, there are several additional types of costs which may be relevant to environmental policies depending on the setting. These include:

1. Government regulatory costs: the costs to the government of administering, monitoring, and enforcing regulations; these are frequently small relative to the direct, private sector compliance costs.
2. Social welfare costs: losses in consumer and producer welfare due to increases in prices or decreases in output of regulated goods and services.
3. Transitional costs: costs of reallocating resources due to regulation, such as plant closings or disruptions in production.
4. Indirect costs: regulation can affect changes in market structure, product quality, or investment, all of which leads to non-trivial costs.

How do we define and measure benefits? From an economic point of view, a person's value for a particular good can be determined by what they would willingly give up in exchange for that good. Within economics, this measure of benefits is referred to as **willingness to pay**. For most goods, measuring benefits is straightforward: we can gather data on market prices and the quanties purchased and estimate individuals' willingness to pay. This approach does not work in the case of environmental quality, however, because environmental goods such as clean air or clean water are not typically traded in markets. Economists have developed two basic strategies to circumvent this lack of price and quantity data: the first uses observed behavior in related markets to infer willingness to pay for environmental quality. A common approach is to use observed market prices to infer the implicit prices for environmental amenities that are bundled with private goods, such as home sales. The second basic approach is to simply ask individuals how much they would be willing to pay for a given environmental good.

Knowing benefits and costs can help to identify the economically efficient policy. It can also help to compare alternative policies, neither of which necessarily maximizes net benefits relative to all other possible policies. This exercise, known as **benefit-cost analysis** is common in practice since it can be difficult to assess the universe of potential policies to find the efficient option. The basic idea is simple: policy 1 is more efficient than policy 2 if the net benefits are greater under policy 1.

5. Market Efficiency and Market Failures

How does an economist define a market? A **market** is a decentralized collection of buyers and sellers whose interactions determine the allocation of a good or set of goods through exchange. To understand market outcomes, you need to describe the behavior of sellers and buyers and economists do so using supply and demand curves.

Consider first the demand curve: a **demand curve** summarizes how much buyers in the aggregate will buy at a given market price, holding other factors constant. Demand curves capture a fundamental fact referred to as the "law of demand:" as the price of a good falls, some people who were already buying the good decide to consume more, while some people who choose not to buy the good at higher prices start to purchase it. A demand curve therefore implicitly traces out consumers' willingness to pay; as a result, we can think of the demand curve as consumers' marginal benefit curve.

Consider the supply side of the market. A **supply curve** summarizes the relationship between price and quantity on the supply side, representing how much of a good suppliers are willing to sell at a given price. Supply curves capture a fundamental fact referred to as the "law of demand:"all else equal, the quantity supplied of a good rises when the price of the good rises. The supply curve traces the marginal cost curve of the industry: the amount a firm is willing to produce and sell at a given price will be a direct function of their marginal cost of production.

What happens when supply and demand interact as they do in a market? The market will eventually settle where demand equals supply, a combination of quantity and price that is referrred to as the **market equilibrium**. Since the demand and supply curves trace out the marginal benefit and cost curves, respectively, we can apply our definition of efficiency: under certain conditions, market outcomes must be efficiency. What conditions must be met for a market equilibrium to be efficient? It turns our that three conditions must hold:

1. The market must be competitive, meaning that firms and consumers take prices as given.
2. Firms and consumers must both have good information about the quality of the good or service being traded.
3. The market must be "complete" meaning that all relevant costs and benefits must be borne by the market participants (consumers and suppliers).

If any of the above conditions are not met, an economist would say that a **market failure** exists. Such situations are common in the environmental realm and give rise to the need for policy intervention. Indeed, climate change can be viewed as a market failure.

Economic activity often leads to undesirable by-products, such as water pollution or greenhouse gas emissions, that impose indirect costs on consumers or firms not participating in a given market. Economists refer to such phenomena as **negative externalities**: an externality results when the actions of one individual (or firm) have a direct, unintentional, and uncompensated effect on the welfare of other individuals or the profits of other firms. Note that this definition does not necessarily require that the effects are adverse; it is possible to have both positive (beneficial) and negative (harmful) externalities.

Most cases of environmental or climate policy involve negative externalities. However, there are important cases of market failures within the environmental realm which do involve positive externalities, namely the provision of **public goods**. Public goods are goods that are shared by all and owned by no one. These goods have two foundational properties: first, they are nonrival, meaning the amount of any individual's consumption does not diminish the amount available for others. Second, they are nonexcludable, meaning that individuals cannot be prevented from enjoying a public good. The provision of public goods often leads to a market failure in the real world due to the free-rider problem: given these fundamental properties, the cost of providing a public good is borne by those who provide it; however, the benefit is enjoyed by all. Thus, private individuals and firms, when left to their own devices, will not have a substantial incentive to supply public goods, leading to underprovision relative to what would be socially optimal.

Glossary

abatement A reduction in pollution; pollution control.

bad Anything that someone does not want. Something of which people prefer less.

benefit-cost analysis A comparison of expenditures and advantages in dollar terms resulting from various actions. Often used to compare alternative public sector policies, none of which necessarily maximize net benefits relative to all other possible policies.

demand curve A graph of the relationship between the price of a good and the quantity demanded.

discount rate Factor used to convert monetary amounts to ones for a standardized year. Does not capture uncertainty in outcomes or inflation over time, but does capture the time value of money.

distributional equity The concept or idea of fairness in economis, particularly in regard to taxation or welfare economics.

dynamic efficiency The principle that economic efficient outcomes in dynamic settings (settings with benefits and costs in multiple periods) maximizes net present value.

economic efficiency An efficient outcome or policy maximizes net benefits. An efficient policy or outcome equates marginal benefits and costs.

equimarginal rule The definition of economic efficiency which states that the efficient policy or outcome (for example, the level of pollution abatement) equates marginal benefits and costs.

free-rider problem Problem that occurs when someone thinks he may be able to enjoy something without paying for it, and fails to contribute ever a portion of the cost.

good Anything that someone wants. Something of which people prefer more.

Kaldor-Hicks criterion The principle that a policy should be undertaken if the size of the resulting gains and the size of the resulting losses are such that those who gain could (in theory) fully compensate those who lose and still be better off themselves. The "potential Pareto criterion" is satisfied if it would be Pareto efficient provided that the winners from the policy compensated the loser, but such transfers do not actually have to take place.

law of demand The claim that, other things equal, the quantity supplied of a good rises when the price of the good rises.

law of demand The claim that, other things equal, the quantity demanded of a good falls when the price of the good rises.

marginal benefit The increase in total benefits as one more unit is produced; in pollution control contexts the increase in total benefits as one more unit of pollution is abated.

marginal cost The increase in total costs as one more unit is produced; in pollution control contexts the increase in total costs as one more unit of pollution is abated.

market A group of buyers and sellers of a particular good or service.

market failure A situation in which a market left on its own fails to allocate resources efficiently.

market equilibrium THe combination of quantity and price for which demand equals supply.

net present value The difference between the present value of the future stream of benefits and the present value of the future stream of costs. The sum of present-discounted stream of net benefits.

net benefits The difference between the benefits and the costs of a given outcome or policy.

opportunity cost Whatever must be given up to obtain some item.

Pareto efficiency The principle that policies should be undertaken if no member of society could be made better off by an alternative policy without making at least on person worse off. A policy is "pareto efficient" if—and only if—no member of society could be made betteer off by an alternative policy without making at least one person worse off.

present value Value today of a monetary amount in the future.

public good A good that cannot be withheld from people even if they don't pay for them (non-excludable). A good which, if made available to one person, automatically becomes available to all others in the same amount.

social discount rate The relative valuation placed by society on future consumption that is presently sacrificed.

supply curve A graph of the relationship between the price of a good and the quantity supplied.

willingness to pay The maximum amount that a buyer will pay for a good.