CHAPTER 1

Environmental Economics and the Theory of Externalities

Our aim in this book is to provide a comprehensive treatment of graduate level environmental economics in a single volume, using a style of presentation that integrates the many sub-areas of inquiry that have come to define the field. To this end we begin in this chapter by introducing the field of environmental economics via its roots in neoclassical welfare theory and the theory of externalities. Environmental problems and policy challenges stem, of course, from human uses of the environment and natural resources. This statement alone does not provide the basis for our study, however. Rather, it is the way that humans use the environment and the impact this use has on the well-being of others that interests us and defines the field. Our starting point therefore is the notion that one person's interactions with the environment can have direct and unsolicited effects on another person, without compensation or other recognition of the impact. To use the classic example, a factory owner whose plant sits next to a laundry impacts the launderer by dirtying the air he needs to produce clean linen. The launderer suffers as a result of the actions of the factory owner, without recourse or compensation. A contemporary example involves the leaching of nitrogen fertilizer from agricultural fields into underground aquifers, from which surrounding communities draw drinking water. Users of groundwater for drinking suffer due to the actions of the farmer, again without recourse or compensation.

These two examples serve to illustrate the types of problems considered in environmental economics and hint at both the positive (describing what is) and normative (describing what ought to be) aspects of study. From a positive point of view, we might be interested in understanding how existing institutional structures lead the self-interested factory owner and farmer to undertake actions that have negative consequences for others. From a normative perspective, we might be interested in suggesting policy interventions that help mitigate these consequences. In either case we are dealing with a potential misallocation of resources that affects the level of well-being that members of society can obtain. It is in this sense that environmental economics falls under the rubric of welfare economics and the theory of externalities, dealing specifically with the failure of market economies to properly account for the environmental ramifications of economic activity. As an aside, we are also dealing with behavioral interactions between humans and the natural environment, or humans and environmental policy. In this sense, environmental economics is also closely tied to both the broader environmental and public policy sciences. In the remainder of this chapter we focus on developing the welfare theoretic basis for environmental economics. We delay discussion of the field's links to policy and environmental science until the next chapter.

1.1 MARKET FAILURES

The starting point for developing the welfare-theoretic basis for environmental economics is the definition of a normative criterion that we will use to judge the desirability of different economic outcomes. The criterion we use is Pareto optimality, which has both virtues and weaknesses as a normative basis for policy recommendations. Its virtue is its simplicity, illustrated by the following definition.

Definition 1.1
An economic outcome is said to be Pareto optimal if a reallocation of resources cannot make at least one person better off without making another person worse off.

From this it follows that there is the potential for a Pareto improvement if one person's well-being can be improved without decreasing that of another person. A second virtue is the lack of income distribution judgments imbedded in the criterion. By saying nothing about whose well-being matters more, Pareto optimality sidesteps the tricky issue of equity and fairness among economic agents, focusing instead on efficiency – making the economic pie as large as possible, regardless of how it is sliced. This virtue, however, is also a vice, in that outcomes that favor a small number of people while leaving many in need can be judged optimal based on the relatively weak Pareto condition. The trade-off is clear: by saying nothing about income distribution, Pareto optimality as a normative criterion has the potential for admitting rather perverse outcomes into the "desirable" category.

This weakness in the normative criterion is partially offset by the statement of the first and, more directly, second fundamental theorems of welfare economics. Informally, the first welfare theorem states that if markets are complete and perfectly competitive, a decentralized price system coupled with self-interested behavior provides an allocation of resources among society's individuals that is Pareto optimal. The second welfare theorem states the converse of this. If markets are complete and perfectly competitive (and certain regularity conditions hold), any Pareto optimal allocation can be supported by a price system arising from an appropriate redistribution of income, via lump sum taxes and transfers. The appeal of the first and second welfare theorems is obvious: if the conditions are met, society can achieve the largest size pie (efficiency) simply by allowing the free market to function. If society perceives some inequity in how individuals fare in the market system (how large a slice of pie some are getting), simple transfers of income can be used to appropriately adjust the distribution of well-being to something more palatable.

Besides the obvious difficulty in arranging suitable transfers of income, the problem with the first and second welfare theorems is that their conditions are often not met. The two examples from above are cases in point. There is no market intervening between the factory owner and launderer to sort out how much dirty air will be allowed in the vicinity of the linen. Likewise, there is no market determining the "proper" quantity of nitrogen that ends up in local drinking water. These missing markets arise from the absence of well-defined property rights for clean air and potable groundwater. Without property rights, there can be no self-interested owner who makes decisions on how much air and water contamination to allow, and for what price, meaning the factory owner and farmer are not obliged to pay for their use of the environment as an input. In these two examples, we cannot conclude from the first welfare theorem that the free market leads to an efficient allocation. Thus, rather than using the fundamental welfare theorems as justification for non-intervention in markets, environmental economists use them as a point of departure for understanding when free markets are unlikely to deliver efficient outcomes. Judging the validity of the first welfare theorem's result involves establishing the extent to which its conditions hold, which takes us to the theory of externalities and public goods.

For environmental problems, complete markets are the exception rather than the norm, owing to the related problems of externalities and public goods. There are several ways to define an externality using either informal description or formal mathematics. We rely on the following definition from Baumol and Oates (1988, p. 17) to motivate our analysis.

Definition 1.2
An externality exists when agent A's utility or production function depends directly on real variables chosen by another agent B, without an offer of compensation or other attention given to the effect on A's well-being.

A key aspect of the definition rests on the notion of *real* variables. The examples of the laundry and factory owner and farmer and water consumers are externalities in that real outcomes – smoke emissions and nitrogen leaching, respectively – directly impact the production or utility functions of the victims. We can contrast this with a different type of interaction in which, for example, the factory owner's demand for labor drives up the wages the launderer must pay his workers. Here again the launderer is impacted by the factory owner, but the interaction has an indirect path that is filtered through the labor market, and manifests itself only through the price of labor. This latter type of interaction is not an externality according to our definition.

Many externalities in environmental economics have a structure that is similar to that of public goods. Public goods, by definition, are goods that are at least partially non-rival and non-exclusive. Non-rival means multiple people can simultaneously enjoy the services of the good; non-exclusive means that none of these people can be *prevented* from enjoying the services of the good. It is via the latter characteristic that public goods cause problems for the conditions of the first and second welfare theorems. If a good is non-exclusive, a meaningful competitive market price cannot arise in a decentralized system, since individuals can have the good for "free" even if they do not elect to "buy" some positive amount of it. For example, it is possible to use the services of National Public Radio in the United States, even if one does not make a contribution to the annual fund drive. This is the classic *free rider* problem associated with public goods. Thus public goods, like externalities, lead to a type of missing market situation that violates the conditions of the first and second welfare theorems. This is relevant for environmental economics, in that many environmental problems and their solutions share characteristics of public goods (or bads). For example, smog exists frequently in Mexico City due to unfortunate combinations of geography, weather and vehicle traffic – and it broadly affects people with respiratory problems. Efforts to reduce smog in the city would benefit these same people, regardless of their contribution to the control effort. In this sense, air pollution is a public bad, and a reduction in air pollution is a public good. As such, voluntary efforts to improve air quality will lead to its under-provision, since people have incentive to free ride rather than actively contribute.

The definitions and examples of externalities and public goods highlight the importance of these concepts for understanding environmental problems. The failure of the fundamental welfare theorems to hold in the presence of externalities and public goods provides the analytical starting point for the study of environmental economics. Before using these concepts in a formal analysis, however, we consider the following intuitive description of market failure in an environmental context.

1.2 DESCRIPTIVE EXPLANATION OF MARKET FAILURE

We illustrate the concepts of Pareto optimality and market failure via the familiar parable of a small, isolated, and primitive island economy. To begin, suppose there is a single person *R* living on the island who extracts resources, engages in production, and consumes goods provided directly by the natural system, as well as those he produces himself using the island's resources. For concreteness, suppose further that the island contains a small grove of slow-growing timber, which provides the only source of wood for constructing shelter and consumer goods, as well as habitat for the island's only game animals. In this initial setup, there can be no externalities and any observed outcome is efficient according to our Pareto criterion, since *R*'s actions by definition affect only his own well-being. This is worth elaborating on. Person *R* can elect to cut the stand of timber and eat the entire population of game animals immediately, and the outcome will still be considered optimal. While outside observers might object to the removal of an entire animal population and clear cutting a forest on moral grounds, our normative criterion considers only the human society's well-being based on its members' preferences. For the single person society, this criterion admits only the well-being of the individual *as he defines it* – and therefore any observed outcome will be Pareto optimal.

This is of course an uninteresting case from the perspective of economics, but it serves to illustrate how the concept of an externality is directly tied to human interactions and consumer sovereignty, and is thus an amoral concept, as opposed to a moral concept based on higher order judgments of right and wrong. To illustrate the economic concept of an externality, suppose now that the island is inhabited by a second person *F*. With this step, a number of interesting economic problems present themselves. These include issues of specialization, exchange, ownership of property, income distribution, and the potential for externalities. Suppose that *R* and *F* agree to joint ownership of the grove of trees. Suppose as well that they are different in their preferences, in that *R* does not much care for meat, but would like to have a large shelter, while *F* does not mind sleeping outside, so long as he can eat meat every day. In this setup, the two men see different, and competing, uses for the shared grove of trees: *R*'s large shelter will require a reduction in habitat for game animals, thereby decreasing the quantity of animals available to *F*. In the language from above, it is now the case that self-interested behavior on the part of both people will not lead to an efficient outcome – the externality has caused a market failure. Without some type of coordination, *R* will over-cut trees and *F* will be left to suffer without his daily meat.

In the island parable, the solution to this problem is obviously some type of coordination between the two people. This might take one of two forms. First, suppose *R* was the first to arrive on the island and as such has established undisputed control over the grove of trees. Person *F* might offer something to *R* – perhaps advice and help on constructing his shelter out of stone rather than wood – in return for his agreement to leave the grove of trees in place. In this scenario, the conditions of the first welfare theorem hold, in that the clear division of property rights leads to an *exchange mechanism* (bartering over advice and labor) for the resource that formerly had a missing market. As such, the negotiated outcome is efficient, and it will have come about through non-cooperative market-like coordination. Second, *R* and *F* might engage in cooperative or collective decision-making as a household or clan might, voluntarily weighing both sets of preferences to arrive at an agreed upon division of uses of the grove. In the special case where their preferences are identical, the collective agreement would only require mutual restraint, to avoid over-use of the common resource. In the case of heterogeneous preferences, the collective decision will reflect the men's sense of equity and their negotiating skills. Regardless, the cooperative agreement balances the competing uses of the grove, and provides an efficient allocation of the resource. Both coordination strategies – non-cooperative/market and cooperative/collective – serve to "internalize" the externality, allowing *R* and *F* to reach an efficient outcome. The two approaches can, however, result in substantially different distributions for the two men's well-being.

How is this parable relevant for environmental economics as a whole? Most types of environmental problems result from a failure of coordination due to missing markets, incomplete property rights, and/or the inability of affected individuals to make collective decisions. The solution to environmental problems often involves designing coordination mechanisms that can work at the scale of the problem at hand. We begin to formally analyze this type of coordination failure (and potential solutions) for a modern economy in the next section.

1.3 A FORMAL MODEL OF EXTERNALITIES

Consider a simple modern economy with two individuals, a dirty good, a clean good, and labor as the only factor of production. Define the utility function for each person by $U_i(x_i,z_i,E)$ for $i=1,2$ where x_i and z_i are consumption levels of the two goods and E is an exogenous (to the individual) level of pollution emissions. Production of x causes the emissions. Define the production function for x by $x=f(l_x,E)$ where both labor input l_x and emissions E have a positive marginal product. This setup treats emissions as an input, implying that a reduction in pollution reduces the output of x, by decreasing a productive factor. While it is also possible to develop the model by treating x and E as joint products, this approach reduces notational clutter. The clean good z is produced using only labor according to the production technology $z=g(l_z)$. Labor employed in the economy is constrained by the work time endowment l such that $l_x+l_z=l$.

1.3.1 Pareto Optimality

Our first task is to derive the conditions for a Pareto optimal allocation of resources in this economy. This involves finding the consumption levels, factor allocations, and pollution amount that maximizes one person's utility in a way that makes the other no worse off than a given benchmark. The allocation must also obey the inherent technology and factor endowment constraints. The problem is given analytically by

$$\max_{x_1,x_2,z_1,z_2,l_x,l_z,E} U_1(x_1,z_1,E)+\lambda_u\left[U_2(x_2,z_2,E)-\bar{u}_2\right]+\lambda_x\left[f(l_x,E)-x_1-x_2\right] \\ +\lambda_z\left[g(l_z)-z_1-z_2\right]+\lambda_l\left[l-l_x-l_z\right], \tag{1.1}$$

where we are maximizing person 1's utility subject to the constraint that person 2 obtains at least utility level $\bar{u}_2$. For clarity of exposition, we state and interpret the first-order conditions for maximization in stages. Beginning with the consumption levels we have

$$\frac{\partial U_1(\cdot)}{\partial x_1}=\lambda_x,\quad \frac{\partial U_1(\cdot)}{\partial z_1}=\lambda_z \tag{1.2}$$

and

$$\lambda_u\frac{\partial U_2(\cdot)}{\partial x_2}=\lambda_x,\quad \lambda_u\frac{\partial U_2(\cdot)}{\partial z_2}=\lambda_z, \tag{1.3}$$

from which we can derive the expression for *efficiency in consumption* as

$$\frac{\partial U_1(\cdot)/\partial x_1}{\partial U_1(\cdot)/\partial z_1}=\frac{\lambda_x}{\lambda_z}=\frac{\partial U_2(\cdot)/\partial x_2}{\partial U_2(\cdot)/\partial z_2}. \tag{1.4}$$

From this we see that efficiency in consumption requires that the solution to Eq. (1.1), denoted $\{x_1^*,x_2^*,z_1^*,z_2^*,l_x^*,l_z^*,E^*\}$, must be such that the marginal rate of substitution between goods is equal for both individuals. Taking derivatives with respect to labor we obtain

$$\lambda_x\frac{\partial f(\cdot)}{\partial l_x}=\lambda_l,\quad \lambda_z\frac{\partial g(\cdot)}{\partial l_z}=\lambda_l, \tag{1.5}$$

which implies that the shadow value of the marginal product of labor for each of the goods should be equal to the shadow price of labor. This is the *efficiency in production* condition. From these expressions, and the result in Eq. (1.4), we can also state the condition for *efficiency in exchange*

$$\frac{\partial U_i(\cdot)/\partial x_i}{\partial U_i(\cdot)/\partial z_i}=\frac{\lambda_x}{\lambda_z}=\frac{\partial g(\cdot)/\partial l_z}{\partial f(\cdot)/\partial l_x}. \tag{1.6}$$

This shows that for an allocation to be Pareto optimal, it must be the case that the slope of the production possibility curve is equal to the slope of each person's indifference curve.

Finally, we complete the Pareto optimal characterization by analyzing the first-order condition with respect to emissions

$$\frac{\partial U_1(\cdot)}{\partial E} + \lambda_u \frac{\partial U_2(\cdot)}{\partial E} + \lambda_x \frac{\partial f(\cdot)}{\partial E} = 0. \tag{1.7}$$

The intuition of this condition is best shown by a simple transformation. Dividing both sides of the expression by λ_x, and substituting the first and third terms in Eq. (1.2), leads to

$$-\frac{\partial U_1(\cdot)/\partial E}{\partial U_1(\cdot)/\partial x_1} - \frac{\partial U_2(\cdot)/\partial E}{\partial U_2(\cdot)/\partial x_2} = \frac{\partial f(\cdot)}{\partial E}, \tag{1.8}$$

which has a clear interpretation. While a reduction in E increases utility for both people directly, it decreases utility indirectly, since a smaller quantity of x is available for consumption. The amount of utility decrease is dependent on the intensity of individual preferences for x. As such, the optimum emission level should balance the utility cost of emissions against the utility benefit of x. This trade-off is shown in Eq. (1.8). The terms on the left represent each person's marginal willingness to give up x for reduced E in consumption, while the term on the right is the physical cost in x of reducing E. Since a reduction in E benefits both people, the marginal cost in x of a reduction in E needs to be compared to the *sum* of both people's marginal willingness to give up x to have the reduction. When the marginal cost is equal to the marginal willingness to pay for reduced E (both as measured in x), the outcome is efficient.

Equation (1.8) also serves to illustrate the symmetry in particular cases between externalities and public goods. Although we have described E as an externality, a reduction in E also has the characteristics of a public good. The benefit that person 1 receives from a reduction in E does not impact or interfere with person 2's benefits, and neither person can be prevented from enjoying the reduction. Viewed in this light, Eq. (1.8) is also the Lindahl-Samuelson condition for the efficient allocation of a public good.[1]

Finally, it is worth emphasizing that $\{x_1^*, x_2^*, z_1^*, z_2^*, l_x^*, l_z^*, E^*\}$ is just one of many possible Pareto efficient outcomes, since it is conditional on a particular reference level of utility for person 2. A different maintained division of utility among the two agents – i.e. a different distribution of well-being – would lead to a different solution generally, and a different optimal level of pollution in particular.

1.3.2 The Competitive Market

The Pareto optimum conditions provide a baseline against which other potential allocations can be assessed. It is reasonable to ask if the free market outcome meets the Pareto criteria, when the output and factor markets are competitive. To study this, suppose p_x and p_z denote the prices of x and z, respectively. In addition, let w denote the price of

[1] The condition is named after the work by Lindahl (1919, 1958) and Samuelson (1954a, 1954b).

labor, and y_i income for person i. Acting as price takers, the individuals maximize utility, and the firms maximize profits. The utility maximization problem for each person i is

$$\max_{x_i, z_i} U_i(x_i, z_i, E) + \lambda_i \left[y_i - p_x x_i - p_z z_i \right], \tag{1.9}$$

with first-order conditions

$$\begin{aligned} \frac{\partial U_i(x_i, z_i, E)}{\partial x_i} &= \lambda_i p_x \\ \frac{\partial U_i(x_i, z_i, E)}{\partial z_i} &= \lambda_i p_z, \quad i = 1, 2. \end{aligned} \tag{1.10}$$

From this we can derive the condition for efficiency in consumption as

$$\frac{\partial U_1(\cdot)/\partial x_1}{\partial U_1(\cdot)/\partial z_1} = \frac{p_x}{p_z} = \frac{\partial U_2(\cdot)/\partial x_2}{\partial U_2(\cdot)/\partial z_2}. \tag{1.11}$$

The profit maximization problem for firm z is

$$\max_{L_z} \left\{ p_z g(l_z) - w l_z \right\}, \tag{1.12}$$

which leads to the first-order condition $p_z \partial g(\cdot)/\partial l_z = w$. Of more interest is the behavior of the firm that produces x. In the market situation, the firm selects both its labor and emission inputs to maximize profit, treating emissions as a free factor of production. The profit maximization problem is

$$\max_{L_x, E} \left\{ p_x f(l_x, E) - w l_x \right\}, \tag{1.13}$$

and the first-order conditions for a solution are

$$\begin{aligned} p_x \frac{\partial f(\cdot)}{\partial l_x} &= w \\ p_x \frac{\partial f(\cdot)}{\partial E} &= 0. \end{aligned} \tag{1.14}$$

Two observations arise from equations (1.12) through (1.14). First, for the market solution, efficiency in labor use and efficiency in exchange are met, in that the value of the marginal product of labor in producing both goods is equal to the wage rate, and the slope of the production function is equal to the slope of each person's indifference curve (the price ratio). Second, and more importantly for our purposes, the condition for an efficient allocation of emissions is not met. The condition for E in Eq. (1.14) is different from Eq. (1.8). The unregulated free market provides no mechanism that encourages the firm that produces x to treat E as scarce, or to account for the impact of its choice on consumers. Instead, the firm emits pollution until the value of the marginal product of pollution is zero – the pollution price the firm faces. Given that the externality is in effect a missing market, and complete markets is a condition for the first fundamental welfare theorem to hold, it is unsurprising that the competitive market equilibrium is not Pareto efficient.

1.3.3 Market Intervention

Much of environmental economics is concerned with designing policy to correct this type of market failure. The intellectual starting point comes from Pigou (1920, 1932), who suggested that if the government can correctly compute and impose a fee (subsequently to be known as the Pigouvian tax) which the polluting firm must pay for each emission, decentralized market behavior will lead to the Pareto optimal result. To see this, define

$$\tau^* = -p_x\left[\frac{\partial U_1(\cdot)/\partial E}{\partial U_1(\cdot)/\partial x_1} + \frac{\partial U_2(\cdot)/\partial E}{\partial U_2(\cdot)/\partial x_2}\right], \tag{1.15}$$

and suppose firm x is required to pay τ^* for each unit of pollution emitted. The firm's modified objective function is now

$$\max_{L_x,E}\left\{p_x f(l_x,E) - wl_x - \tau^* E\right\}, \tag{1.16}$$

and the first-order conditions are

$$\begin{aligned} p_x\frac{\partial f(\cdot)}{\partial l_x} &= w \\ p_x\frac{\partial f(\cdot)}{\partial E} &= \tau^* = -p_x\left[\frac{\partial U_1(\cdot)/\partial E}{\partial U_1(\cdot)/\partial x_1} + \frac{\partial U_2(\cdot)/\partial E}{\partial U_2(\cdot)/\partial x_2}\right]. \end{aligned} \tag{1.17}$$

Comparing condition (1.8) to (1.17), we see that the market outcome conditions do indeed match the Pareto optimum conditions, when pollution is priced at τ^*. This suggests that the government can, in principle, intervene in the market to provide the proper incentive for the polluting firm to internalize the externality. The proper intervention depends, however, on a particular distribution of income. This is apparent from Eq. (1.17), which makes clear that the correct emissions tax depends on solutions to the individuals' utility maximization problems, which are themselves functions of the income (and more generally, endowment) distribution. Changes in income or factor endowments will therefore lead to different tax and emissions levels, which will still be Pareto efficient. This application of the second fundamental welfare theorem is easily overlooked when we refer to "the" optimal pollution level.

The idea of a Pigouvian tax is further illustrated in panel A of Figure 1.1, which relates levels of E to the marginal benefit and marginal cost of additional E. The marginal benefit is reflected in the value of the marginal product of E in producing x (denoted VMP_E), which is the firm's factor demand for E. Without the tax, the firm chooses emissions level $\hat{E}$, where the value of marginal product is zero. The cost of pollution is the marginal utility loss of additional E, denoted by $MUC1_E$ and $MUC2_E$ for the two people, measured as the value of additional x that is needed to compensate for the loss due to E. Because E is non-rival, the total marginal utility cost of additional pollution is found by summing the costs to the individuals, which results in the curve we have labeled MUC_E. The vertical summation of these curves further connects the notion of an externality to the public good (bad) case. Intuitively, MUC_E represents the two individuals' collective marginal willingness to accept compensation in exchange for additional E. The efficient level of pollution is found by balancing the firm's value of E in production, against the non-market cost of pollution to the individuals. The point E^*, where the marginal willingness to accept and the value of the marginal product are equal, is the optimum amount. The polluting firm can be induced to emit at this level if they face an emission fee in the amount τ^*.

Panel B of Figure 1.1 illustrates how the externality and Pigouvian tax can also be viewed from a public good perspective. As described above, pollution emissions E are a bad. We can, however, define its mirror image – the elimination of pollution – as a good. Let $A = \hat{E} - E$ be the units of pollution removed, where $\hat{E}$ is the baseline emissions level shown in panel A. Using this notation, we have $E = \hat{E} - A$, which can be substituted into the households' and firm's utility and profit functions in equations (1.9) and (1.16), respectively. The level of abatement (pollution removal) A is found by differentiating with respect to A, and the optimality conditions are reflected in the figure. The only difference from panel A is that the curves now slope in opposite directions. For the firm, increased abatement implies a marginal cost, because it involves reduction of a productive input. For the households, increased abatement corresponds to fewer emissions, and so it is a benefit. The firm's curve, labeled MC_A, is the marginal cost of providing abatement, and the curve labeled MUB_A is the households' collective marginal willingness to pay (demand) for abatement. Because abatement is a public good, demand is found by summing the individual households' marginal willingness to pay curves. The optimal abatement level is labeled as A^*, which is linked to panel A by the relationship $A^* = \hat{E} - E^*$. The Lindahl-Samuelson optimal price per unit for abatement is τ^*, of which person 1 should contribute τ_1^* and person 2 τ_2^*, so that $\tau^* = \tau_1^* + \tau_2^*$.

Figure 1.1 and Eq. (1.17) serve to foreshadow many of the themes we will discuss in the balance of the book. From the perspective of designing environmental policy, Eq. (1.17) establishes a shadow price of emissions that a polluting firm must face if its emissions are to be efficiently reduced. *How* the environmental authority imposes this shadow price on firms via environmental policy – e.g. emission taxes, pollution permits, standards – is our theme as we begin the book. From the perspective of policy implementation, Eq. (1.17) also illustrates the measurement challenges faced by environmental authorities. The shadow price of pollution is a function of individuals' preferences for the environment, which are not directly observable. The conceptual and empirical techniques necessary for measuring preferences are considered later in the book.

1.4 POINTS GOING FORWARD

Our approach in this book will adhere fairly strictly to the neoclassical economic paradigm as we explore the design and implementation of environmental policy. Even in a