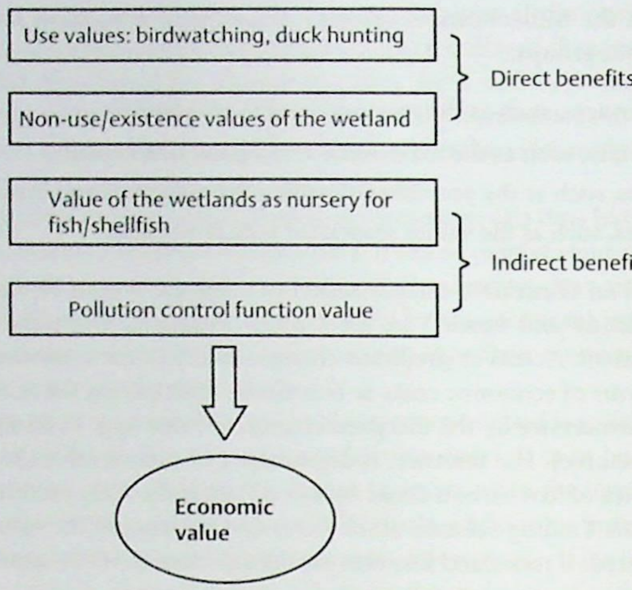

Figure 3.4 The economic value of a wetland.

in the catchment, then this could be valued by looking at the additional costs to water companies of cleaning up the water prior to supplying it to customers. If moorlands are replaced by forests, then lost sheep production could be valued using the market price of lambs.

This idea of 'valuing ecosystem services' can be related to a further categorization system of environmental benefits that has become popular in the literature (Pearce and Turner, 1989). Take, as an example, the preservation of a wetland that is important to birds, but that also functions as a nursery for fish/shellfish and as a natural pollution control plant. How might the *total economic value* of this wetland be described (Figure 3.4)?

Consider first what we have called direct benefits; that is, direct sources of utility. Some of those who benefit from the wetland in this way may participate in activities that make the wetland valuable to them, such as birdwatching or duck hunting. Such benefits are often known as *use values*, since they require actual participation to enjoy them. Use values may be consumptive (hunting) or non-consumptive (birdwatching). However, people other than those who actually visit the wetland may derive benefits, in terms of the utility they get from just knowing that the wetland is preserved. These types of benefit have become known as *non-use* or *existence values*. They may be motivated by selfish reasons, or by altruism, either for other members of the current generation, or for future generations. Existence values may be particularly high for unique, irreplaceable natural assets, such as the Grand Canyon in the United States or Kakadu National Park in Australia.

The sum of use and existence values gives the total direct benefits of preserving the wetland. The wetland's role as a nursery for fish and shellfish could be evaluated by estimating biological models of the contribution that the wetland makes to fish/shellfish populations, and then by looking at the economic (commercial) value of these species. Changes in these

Table 3.2 Values provided by tropical coastal and marine ecosystems

Direct values	Indirect values	Existence and bequest values
Fishing:	Nutrient retention and cycling:	Cultural heritage:
• Aquaculture	• Flood control	• Resources for future generations
• Transport	• Storm protection	• Existence of charismatic species
• Wild resources	• Habitat for species	• Existence of wild places
• Water supply	• Shoreline stabilization	
• Recreation		
• Genetic material		

Source: Adapted from Heal et al. (2005).

economic values, in terms of gains/losses in consumers' surplus[4] and producers' profits from some change in the wetland, could be calculated. Finally, the wetland's pollution control function could be valued either by using the value of avoided pollution damages (say, from sedimentation of coral reef fisheries, or from nutrient enrichment), or the pollution control costs that would have to be incurred to replace the role currently being played by the wetland. The sum of avoided pollution and/or pollution control costs, plus the value of commercial fisheries, would give the indirect benefits of preserving the wetland. Adding the wetland's direct and indirect benefits gives its total economic value. Table 3.2 shows an illustration of a similar breakdown of benefits for tropical coastal ecosystems, such as a mangrove wetland (note that this employs a somewhat different categorization of benefits than we have used).

3.3 Why Place Economic Values on the Environment?

Economic estimates of the value of changes in environmental and resource quality may be useful in a number of contexts. The main use that we review here is cost–benefit analysis (CBA; also called benefit–cost analysis in the USA), although only a brief explanation of this method is provided (for more details, see Hanley and Barbier, 2009). Other uses of environmental values are noted at the end of this section.

3.3.1 Cost–benefit analysis

Economists have long been intrigued by the problem of how to decide whether one outcome is better than another from society's point of view (see Box 3.1). Ideally, we would like to find decision-making rules that give consistent outcomes; outcomes that are the same when applied in the same circumstances. We would also like to find a method that is democratic, in some sense, and practical; and that can be shown to be consistent with economic theory. Welfare economics developed out of the search for such a method. Cost–benefit analysis developed from welfare economics as a practical application of a decision-making rule that

[4] As Chapter 2 explains, the consumers' surplus is the difference between the most you are willing to pay for something and the price you actually pay. In a market-traded good, it is the area underneath the demand curve but above the equilibrium price.

BOX 3.1 What Role Should CBA Play in the Policy Process?

One of the earliest uses of cost-benefit analysis in public policy appraisal was in the assessment of water resource projects such as new dams or flood control investments in the USA. The 1936 Flood Control Act stated that the federal government should undertake public investments in flood alleviation if 'the benefits to whomsoever they may accrue are in excess of the estimated costs'. As Banzhaf (2009) explains, this eventually resulted in fierce debate amongst economists in the USA as to the proper role of CBA in the policy appraisal process. Essentially, the argument was between those who viewed the philosophical and theoretical structures of CBA to be robust enough for the outcome of a CBA to be viewed as actually determining whether a particular project should go ahead (this group was associated particularly with a think tank called Resources for the Future, which is still in existence today) and those who were sufficiently uneasy about the principle of making interpersonal utility comparisons in dollar terms to conclude that CBA should only inform the decision-making process. This latter group, most identified with Harvard University's Water Program, felt that political or 'expert' judgement should always be decisive, and that the role of the CBA analyst was no more than to make clear the trade-offs involved in deciding whether or not to proceed with a project. Moreover, the Harvard team viewed CBA as focusing too closely on a single objective–economic efficiency–in contrast to the multiple objectives of public policy.

These arguments over the correct role of CBA within the public policy process, and of the advantages and disadvantages of CBA relative to other project and policy appraisal methods, have continued ever since. However, CBA emerges as a remarkably robust institution, perhaps because of its apparent simplicity of approach: add up the benefits to society of a particular project and compare these with the costs. This still makes sense to a lot of economists and policy analysts as a way of thinking about decisions.

could be used to decide between different policy options or projects in terms of their net contribution to social well-being. CBA consists of identifying the impacts of a project or policy, valuing these impacts in terms of their effects on social well-being, and then comparing the good effects (benefits) with the bad effects (costs). Costs and benefits are expressed in monetary terms to allow comparison. The links with welfare economics come in terms of how benefits and costs are measured (e.g. using the principles of WTP and opportunity costs), and with the basis on which the difference between benefits and costs can be viewed as a proxy for the underlying change in net social welfare. This basis is often referred to as the *Kaldor–Hicks compensation test*. This asks: Could the gainers (those who benefit from a project) compensate the losers, and still be better off? Acceptance of this principle as the basis for evaluating contributions to social well-being in turn requires us to accept that: (i) all relevant benefits and costs can be expressed in the same units; and (ii) benefits and costs can be compared with each other, so that any cost (loss) can always be compensated by some offsetting benefit (gain). Clearly, not all would agree with these statements (see, e.g., Aldred, 2006).

3.3.1.1 The stages of a CBA

(1) Project/policy definition

This involves setting out exactly what is being analysed; whose welfare is being considered; and over what time period. In terms of 'whose welfare', the usual answer is that it is national well-being that is considered, although defining the 'relevant population' is often a difficult issue. For instance, if a new dam in Indonesia would threaten an internationally rare habitat, should the costs to foreign conservationists be counted?

(2) Identify physical impacts of the policy/project

Any project/policy has implications for resource allocation: for example, labour used to build access roads to a new hydroelectric dam; additional electricity production; land used up in the creation of the reservoir; or less pollution being generated from a coal-fired power station that can now be closed early. The next stage of a CBA is to identify these outcomes in physical magnitudes. Frequently, these changes in resource allocation will not be known with certainty—For example, how many tonnes of pollution will be displaced? Once physical impacts have been identified and quantified, it is then necessary to ask which of them are relevant to the CBA. Essentially, anything that impacts on the quantity or quality of resources, or on their price, may be said to be relevant, if these impacts can be traced back to a link to the well-being of the relevant population. Since we specify relevant impacts in terms of utility impacts, it is not necessary to restrict attention to market-valued impacts, since non-market value changes (such as an improvement in air quality) are relevant if they affect people's utility.

(3) Valuing impacts

One important feature of CBA is that all relevant effects are expressed in monetary values, so that they can then be aggregated. The general principle of monetary valuation in CBA is to value impacts in terms of their marginal social cost or marginal social benefit. Here, 'social' means 'evaluated with regard to the economy as a whole'. But where are these marginal social benefits and costs derived from? Under certain conditions, this information is contained in market prices, as Chapter 2 explains. Market prices contain information on both the value to consumers of a particular product (say, electricity) being supplied, and the costs to producers of supplying it. The market wage rate, similarly, shows both the value of labour to employers and the value of leisure to workers. Assuming that the impacts of the project are not large enough to actually change these prices, then market prices are a good first approximation to the values of benefits and costs. But markets often 'fail', as Chapter 2 shows. Moreover, for some environmental goods such as biodiversity and river water quality, no market at all exists from which a price can be observed. In this case, methods described in Chapter 4, and based on the principles of valuation outlined earlier in this chapter, will need to be employed.

(4) Discounting of cost and benefit flows

Once all relevant cost and benefit flows that can be expressed in monetary amounts have been so expressed, it is necessary to convert them all into *present value* (PV) terms. This necessity arises out of the time value of money, or time preference. To take a simple example, suppose that an individual is asked to choose between receiving £100 today and receiving that same £100 in one year's time. The more immediate sum might be preferred due to impatience (I want to spend the money right now). Alternatively, I may not want to spend the money for a year, but if I have it now, I can invest it in a bank at an interest rate of, say, 10 per cent and have £100 × $(1 + i)$ = £110 in one year's time, where i is the rate of interest. Benefits are more highly valued the sooner they are received. Similarly, a sum of money to be paid out, or any kind of cost, seems less onerous the further away in time we have to bear it. A bill of £1 million to repackage hazardous wastes seems preferable if paid in 100 years time rather than in 10 years time. This is nothing to do with inflation, but more to do with the expectation that we might expect to be better off in the future. Box 3.2 gives an example

BOX 3.2 An Illustrative CBA

Consider a project to support the construction of a new renewable energy source. A plan is proposed to build a small-scale wind farm in a scenic area. The initial construction costs are estimated to be £750,000. Following start-up, annual maintenance costs of £5,000 are expected throughout the 15-year lifespan of the plant. At the end of this 15-year period, the wind farm will need to be dismantled and the site restored, at an expected cost of £35,000. Every year after the initial construction year, the site will produce electricity with a market value of £150,000, which for now we take to be a constant flow in real terms (we will ignore the effects of inflation here). Objectors have protested about the visual impact of the windmills, and so the government has commissioned a contingent valuation study of local residents. The results suggest that the mean annual compensation demanded by locals is £25 per household: 2,000 households are thought to be affected (this mean is calculated across both those against the project and those in favour).

It is easy to set up a basic CBA of this project. The initial ('year zero') construction costs are not discounted, as they occur at the start of the project. Maintenance costs are then discounted each year using the relevant discount factor, for a discount rate of 6 per cent. Annual environmental costs of (£25 × 2,000) are also discounted over the 15 years of the project, and are assumed to stop when the site is restored at a cost of £35,000 in year 15. This year-15 cost also needs to be discounted. Annual benefits of £150,000 get discounted each year, you can see how the present value of this fixed amount falls each year as we move forward in time. The following table shows all workings:

Year	Discount factor at 6% discount rate, $(1.06)^{-t}$	Benefits (£)	Present value of benefits (£)	Costs (£)	Present value of costs (£)
0	1		0	750,000	750,000
1	0.9433	150,000	141,495	55,000	51,881
2	0.8899	150,000	133,485	55,000	48,944
3	0.8396	150,000	125,940	55,000	46,178
4	0.7921	150,000	118,815	55,000	43,565
5	0.7472	150,000	112,080	55,000	41,096
6	0.7049	150,000	105,735	55,000	38,769
7	0.6650	150,000	99,750	55,000	36,575
8	0.6274	150,000	94,110	55,000	34,507
9	0.5918	150,000	88,770	55,000	32,549
10	0.5583	150,000	83,745	55,000	30,706
11	0.5267	150,000	79,005	55,000	28,968
12	0.4969	150,000	74,535	55,000	27,329
13	0.4688	150,000	70,320	55,000	25,784
14	0.4423	150,000	66,345	55,000	24,326
15	0.4172			35,000	14,602
Total discounted benefits/costs			1,394,130		1,275,779

As we can see, the total present value of costs is £1,275 million, whilst the total present value of benefits is £1.394 million. Thus the net present value of the project is positive, at £118,351, which passes the CBA test at the 6 per cent discount rate. Note how little the site renovation cost of £35,000 amounts to in present-value terms: less than half its future value.

BOX 3.3 Discounting and the Discount Rate

Discounting means placing a lower value on benefits and costs the further away in time they occur. Why might this make sense? Two main reasons have been given for discounting. These revolve around:

- the productivity of capital, and
- preferences.

These motivate two possible choices for the discount rate to be used in public-sector policy and project appraisal:

- The social opportunity cost of capital, *r*, and
- The rate of social time preference, *s*.

Economies grow over time for many reasons, but an important one is that by building up the stock of capital, an economy increases its potential output. The act of investing in a new factory is expected to generate a flow of returns over time to the owner of that capital, in terms of annual sales of goods produced. Across the entire economy, invested capital generates a positive rate of return, meaning that the value of consumption goods in year $t + 1$ that could be produced should all of the resources of an economy be invested in year t will be greater than the maximum value of consumption goods that could be produced in year t. However, capital is scarce: investing £1 million in a new factory means that we cannot invest the same £1 million in another scheme. Choosing to invest in a particular scheme thus involves an opportunity cost, which is the return on capital forgone from some other use (in particular, from its most profitable alternative). Across the economy as a whole, we could rank investment projects in terms of their rates of return. These rates of return show the net benefits from investing resources rather than consuming. At the margin, this is known as the opportunity cost of capital, which can be used to measure the *social opportunity cost of capital, r*.

The other motivation for discounting is that 'pure time preference'—the desire for benefits to come sooner rather than later—is a fundamental feature of human desires. Various motivations have been suggested for time preference: impatience; the fact that we might not be around in the future to collect on benefits; that future benefits are less certain than present-day benefits; and that we might expect to be richer in the future, and thus will regard each extra pound of income as less valuable than we do today. An important distinction is between a discount rate that applies to individual well-being and that which might be applied to collective well-being. We could refer to the former as being a reflection of individual time preference and the latter as a reflection of the *rate of social time preference, s*.

For more detailed discussion of the discount rate and of alternative approaches to discounting, see Hanley and Barbier (2009: ch. 7).

of discounting of benefits and costs, whilst Box 3.3 discusses the concept of discount rates in more detail.

How is this time effect taken into account, and how are cost and benefit flows made comparable regardless of when they occur? The answer is that all cost and benefit flows are *discounted*, using a discount rate that is here assumed to be a market rate of interest, *i*. The present value (PV) of a cost or benefit (X) received in time t is then calculated as follows:

$$PV(X_t) = X_t\left[(1+i)^{-t}\right]. \tag{3.1}$$

The expression in square brackets in equation (3.1) is known as a discount factor. Discount factors have the property that they always lie between 0 and +1. The further away in time a cost or benefit occurs (the higher the value of t), the lower is the discount factor. The higher the discount rate i for a given t, the lower is the discount factor, since a higher discount rate means a greater preference for things now rather than later.

(5) Applying the net present value test

The main purpose of CBA is to help select projects and policies that are efficient in terms of their use of resources. The criterion applied is the *net present value* (*NPV*) test, which is how the Kaldor–Hicks compensation principle is implemented. This test simply asks whether the sum of discounted gains exceeds the sum of discounted losses. If so, the project can be said to represent an efficient shift in resource allocation, given the data used in the CBA. The *NPV* of a project is thus as follows:

$$NPV = \Sigma B_t(1+i)^{-t} - \Sigma C_t(1+i)^{-t}, \tag{3.2}$$

where the summations (indicated by the 'Σ' symbols) run from $t = 0$ (the first year of the project) to $t = T$ (the last year of the project). Note that no costs or benefits before year 0 are counted. The criterion for project acceptance is as follows: accept if $NPV > 0$ (i.e. is positive). Any project passing the *NPV* test is deemed to be an improvement in social welfare.

(6) Sensitivity analysis

The *NPV* test described above tells us about the relative efficiency of a given project, given the data input to the calculations. If this data changes, then clearly the results of the *NPV* test will change too. But why should data change? The main reason concerns uncertainty. In many cases in which CBA is used, the analyst must make predictions concerning future physical flows (e.g. the quantity of electricity produced per year) and future relative values (e.g. the wholesale price of electricity). None of these predictions can be made with perfect foresight. When environmental impacts are involved, this uncertainty may be even more widespread; for example, if a policy to reduce global greenhouse gas emissions is planned, then the impacts of this in terms of avoided damage may be subject to a wide range of predictions. Therefore, an essential final stage of any CBA is to conduct a sensitivity analysis. This means recalculating the *NPV* when the values of certain key parameters are changed.

3.3.1.2 Why is CBA useful?

In one very important sense, the practice of CBA addresses what might be called the fundamental economic problem: how to allocate scarce resources in the face of unlimited wants. Resources are scarce because the sum total of demands on them exceeds their availability, and using up scarce resources in one way imposes an opportunity cost on society in that we cannot use those same resources for some other purpose. For example, a proposal in 2007 to expand irrigated agriculture on the Canterbury Plains in New Zealand suggested diverting water from two rivers to a newly constructed reservoir that would then be used to supply irrigation schemes for dairy farmers. However, if land is used up to create a reservoir, that same land cannot also be used for sheep farming. If water is taken from a river to supply a reservoir and then to irrigate dairy farms, that same water is not available in the river to maintain in-stream ecological quality, or to support water-based recreation such as kayaking.

Society might find it useful, in determining whether to allow such schemes, to know whether the economic benefits of irrigated dairy farming were bigger or smaller than the costs of reservoir construction, lost sheep farming output, losses in river ecological quality and losses in kayaking opportunities.

CBA is a decision-aiding tool that conveys this manner of useful information to decision-makers. Not only does CBA allow a comparison of the benefits and costs of particular actions, reflecting therein the scarcity of resources, but it also allows for ordinary people's *preferences* to be included in government decision-making. Economic values in a CBA depend partly on what people like (their preferences), what they are prepared to give up to have more of what they like (their WTP), and what they can afford to pay (their budget constraints). In a sense, CBA is an exercise in economic democracy, since every citizen gets an economic vote in terms of his or her WTP. CBA is also a formal way of setting out the impacts of a project or policy over time, of organizing debate over an issue, and of identifying who enjoys the gains and who suffers the losses from such undertakings. It is also, as Arrow et al. (1998) have noted, a good way of ensuring consistency and perhaps transparency in public-sector decision-making.

3.3.1.3 Uses of environmental CBA

One way in which CBA can be useful is as part of the *policy appraisal* process. Worldwide, much of the funding for environmental valuation studies has come from government departments and agencies with responsibilities for environmental policy design and implementation (e.g. in the United Kingdom, with the Forestry Commission and the Environment Agency); or with responsibility for policies that impact on the environment (e.g. roads policy). Within the European Union, CBA is an important aspect of implementing the Water Framework Directive and the REACH directive on chemicals registration. Within both the UK and the USA, CBA is also a part of the process by which the costs to the economy of new government regulations—for example, the costs of setting stricter targets for recycling of waste (garbage)—are regularly assessed.

Much early work on CBA was carried out in a *project appraisal* context. A good early example is its use in assessing the development of hydroelectric power in the USA (see Box 3.4 and Krutilla and Fisher, 1985). CBA is also used by public forest authorities in the UK in assessing the net benefits of alternative forest management regimes, and in the appraisal of major transport projects, such as new rail lines. The World Bank also has a long history of using CBA in project appraisal. Many governments worldwide have official guidelines on how CBA should be applied to the appraisal of public-sector projects (see, e.g., http://www.hm-treasury.gov.uk/data_greenbook_index.htm, which describes the procedures followed in the UK for both policy and project appraisal).

3.3.2 Other uses of environmental valuation

3.3.2.1 Setting environmental taxes

Environmental valuation has been used in the UK for setting eco-taxes; for example, with regard to the landfill tax and the tax on quarrying. Valuation has been used in both justifying

BOX 3.4 An Example of Cost-benefit Analysis of Hydropower Regulation

Kotchen et al. (2006) carry out a CBA of the re-licensing of two hydroelectric dams in Michigan. The policy context involves a move to reduce the environmental impacts of hydropower operations, by changing how rivers are managed. The changes investigated by Kotchen et al. involve managing releases from dams and reservoirs in a way that more closely parallels natural fluctuations in water levels, rather than timing releases to coincide with maximum electricity demands. This change imposes costs in terms of lost electricity output on hydro operators at peak periods, which must be compensated for with more expensive output from other sources–here, from fossil fuel-powered generation. The gain is an environmental one–in this case, an increase of about 270,000 salmon per year emigrating from the Manistee River to Lake Michigan. Due to the mix of fossil fuel power supplied to the grid, there is also an environmental gain from reduced net air pollution, since the peak-period demands are met from less polluting natural gas-powered generation rather than the more polluting coal sources.

The costs to producers of the change in operations is given by the differences in marginal costs per kilowatt hour (kWh) between hydro-derived and fossil fuel-derived electricity. This implies that the annual costs for the two dams rise by about $310,000. For air pollution, the authors consider five pollutants, including NO_x, CO_2, and SO_2. Changes in air pollution between the two water management regimes are then converted into dollars using estimates from the literature of marginal damage costs, reporting a range of possible values. Finally, changes in migrating salmon numbers are converted into changes in predicted catches for recreational anglers, and then valued using travel cost-derived estimates of the value of recreational fishing (see Chapter 4).

The conclusion of the study is that the benefits of changing the way in which the river system is managed for hydropower produces benefits that are larger than costs. Annual losses in electricity output imply costs in the range of $219,132–$402,094, with a best guess of $310,612. Annual benefits from emission reductions are in the range $67,756–$246,680, whilst gains in recreational fishing are worth $301,900–$1,068,600, with a most likely estimate of $738,400. The authors conclude that 'the benefits exceed the costs of the switch ... even ignoring the air quality benefits entirely, the best estimate of recreational fishing benefits exceeds the upper bound of producer costs'. In this case then, changing how water resources are managed to reduce adverse environmental impacts seems to pass the cost-benefit test.

a tax and determining its level. However, the application of an estimate of the average external cost at the current level of activity does not constitute the Pigovian tax that it is made out to be, since this would typically measure the marginal external cost at the optimal level of externality. Here, the crucial issue would appear to be to know how marginal damages vary with the level of the externality-causing activity.

3.3.2.2 Environmental damage claims

Under the CERLA and Oil Pollution Acts in the USA, firms that cause accidental environmental damages can be sued in the courts by states and by the Federal government to recover the monetary value of such damage. The most famous such incident to date has been the *Exxon Valdez* oil spill in 1989; another example was the accident involving the oil tanker *American Trader* in February 1990, which spilled up to 400,000 gallons of crude oil into the Pacific Ocean off the coast of California (Dunford, 1999). The State of California sued the 'responsible parties' for damages. Interestingly, in this case, the two parties produced alternative sets of estimates of these damages, which resulted in the main from the temporary closure of a number of beaches. The State of California's experts estimated the value of a lost day's recreation on the beach to be around $15 per visit; unsurprisingly, the defendants' economists came up with a lower value of $4–8 per visit. Arguments also raged about how many visits were lost.

3.3.2.3 National accounting

In Chapter 6, we will discuss the issue of making 'green' adjustments to national accounting figures such as gross national product, as a way of producing a better measure of welfare and an indicator of sustainability. Governments can use environmental valuation in calculating such green adjustments to the national accounts; for example, to take account of changes in the level of pollution between this year and the previous year, or to value changes in water quality in a nation's rivers.

Summary

Changes in the quantity or quality of environmental resources have economic value if they have an impact on utility. We can base our measures of these values upon either the most that people are willing to give up to acquire some (desirable) change or the least they are willing to accept to forgo it (for an undesirable change, then we can use either the most people are willing to pay to prevent it, or the lowest compensation they would accept to put up with it). Governments can and have made extensive use of environmental valuation in a number of contexts, particularly cost–benefit analysis in the context of policy and project appraisal. Other uses also exist, including the settling of environmental damage claims and the calculation of green tax rates.

Tutorial Questions

3.1 Why do economists use WTP as a measure of the value someone places on an environmental good or ecosystem service?

3.2 What determines how WTP for a particular change in environmental quality (e.g. a fall in air pollution) may vary across people?

3.3 Several studies have shown that WTP differs substantially from WTA for a given change in environmental quality. Why might this be? Does it matter, from a policy analysis viewpoint?

3.4 What is meant by the economic value of an ecosystem service? How could changes in pressures on an ecosystem (such as an estuary or a forest) produce a change in ecosystem service values? What would we need to know to be able to estimate the monetary value of such a change?

3.5 Explain how you would undertake a cost–benefit analysis of a planned new nuclear power station. Why might the choice of discount rate be particularly crucial to the net present value calculation in this instance? How could uncertainty over (i) the future price of electricity and (ii) future waste storage costs be included in this analysis?

3.6 What are the advantages and problems of applying cost–benefit analysis to environmental policy decisions?

References and Additional Readings

Aldred, J. (2006). 'Incommensurability and monetary valuation', *Land Economics* 82(2): 141–61.

Arrow, K., Cropper, M., Eads, G., Hahn, R., Lave, L., Noll, R., Portney, P., Russell, M., Schmalensee, R., Smith, V.K., and Stavins, R. (1998). 'Is there a role for benefit–cost analysis in environmental, health and safety regulation?' *Environment and Development Economics* 2: 196–201.

Banzhaf, H.S. (2009). 'Objective or multi-objective? Two historically competing visions for benefit cost analysis', *Land Economics* 85(1): 3–23.

Barbier, E.B. (2009). 'Ecosystems as natural assets', *Foundations and Trends in Microeconomics* 4(8): 611–81.

Dunford, R. (1999). 'The American Oil Trader spill', *AERE Newsletter* (May): 12–20.

Hanemann, M. (1991). 'Willingness to pay and willingness to accept: how much can they differ?' *American Economic Review* 81: 635–47.

Hanley, N., and Barbier, E.B. (2009). *Pricing Nature: Cost–Benefit Analysis and Environmental Policy Appraisal* (Cheltenham: Edward Elgar).

Heal, G.M., Barbier, E.B., Boyle, K.J., Covich, A.P., Gloss, S.P., Hershner, C.H., Hoehn, J.P., Pringle, C.M., Polasky, S., Segerson, K., and Shrader-Frechette, K. (2005). *Valuing Ecosystem Services: Toward Better Environmental Decision Making* (Washington, DC: The National Academies Press).

Kotchen, M., Moore, M., Lupi, F., and Rutherford, E. (2006). 'Environmental constriants on hydropower: an *ex post* benefit cost analysis of dam re-licensing', *Land Economics* 82(3): 384–403.

Krutilla, J.V., and Fisher, A.C. (1985). *The Economics of Natural Environments* (Washington, DC: Resources for the Future).

Loomes, G., Starmer, C., and Sudgen, R. (2003). 'Do anomalies disappear in repeated markets?' *Economic Journal* 113 (March): C153–66.

Pearce, D.W., and Turner, R.K. (1989). *Economics of Natural Resources and the Environment* (London: Harvester Wheatsheaf).

Sterner, T. (2011). *Fuel Taxes and the Poor* (Washington, DC: Resources for the Future).

Van der Wal, R., Bonn, A., Monteith, D., Reed, M., Blackstock, K.L., Hanley, N., Thompson, D., Evans, M., Alonso, I., Allott, T., Armitage, H., Beharry, N., Glass, J., Johnson, S., McMorrow, J., Ross, L., Pakeman, R.J., Perry, S., and Tinch, D. (2011). 'Mountains, moorlands and heaths', in *The UK National Ecosystem Assessment Technical Report* (Cambridge: UNEP-WCMC): 106–59.

Willig, R. (1976). 'Consumers surplus without apology', *American Economic Review* 66: 589–97.