CHAPTER 4

Sustainable development

Definitions

Many definitions of sustainable development (SD) (often incompatible with each other) have been suggested and debated in the literature. What this suggests is that the debate has exposed a range of approaches which differ because they are linked to alternative environmental ideologies (see Box 2.1). From the ecocentric perspective, the extreme deep ecologists seem to come close to rejecting even a policy of 'modified' development based on the sustainable use of nature's assets. For them only a minimalist development strategy is morally supportable. From the opposite technocentric perspective, other analysts argue that the concept of sustainability contributes little new to conventional economic theory and policy. Given this worldview, the maintenance of a sustainable economic growth strategy over the long run merely depends on the adequacy of investment expenditure. Investment in natural capital is not irrelevant but it is not of overriding importance either. A key assumption of this position is that there will continue to be a *very high degree of substitutability between all forms of capital* (physical, human and natural capital). The classification scheme set out in Box 2.1 labelled these two positions as *very weak sustainability* and *very strong sustainability* respectively.

The most publicized definition of sustainability is that of the World Commission on Environment and Development (WCED) (the 'Brundtland Commission', 1987). The Commission defined SD as: 'development that meets the needs of the present without compromising the ability of future generations to meet their own needs' (WCED, 1987, p. 43).

On the basis of this SD definition both *intergenerational equity* and *intragenerational equity* concerns must be met before any society can attain the goal of sustainability. Social and economic development must be undertaken in such a way as to minimize the effects of economic activity (on resource sources and waste assimilation sinks – see Chapter 1) whenever the costs are borne by future generations. When currently vital activities impose costs on

the future (e.g. mining of non-renewable minerals – see Chapter 16) *full compensation* must be paid (e.g. performance or assurance bonds yielding financial aid, or new technologies allowing resource switching say from fossil fuels to solar power, etc. – see Chapter 11).

The Commission also highlighted 'the essential needs of the world's poor, to which overriding priority should be given'. In other words, SD must allow for an increase in people's standard of living (broadly defined) with particular emphasis on the wellbeing of poor people, while at the same time avoiding uncompensated and significant costs on future people.

The Commission also took a fairly optimistic view of the possibilities for decoupling economic activity and environmental impact (see Chapter 3) and in terms of our classification system has put itself into the *weak sustainability* camp. Recall that the *strong sustainability* supporters, while not dismissing decoupling, argue that modifications to the scale of the economy (the throughput of matter and energy) will also be required. The amount of scale reduction is debated within the strong sustainability camp (which is a fairly 'broad church').

SD, it is generally agreed, is therefore economic development that endures over the long run. Economic development can be measured in terms of Gross National Product (i.e. the annual output of goods and services adjusted for exports and imports) per capita, or real consumption of goods and services per capita. In a later section we will argue that, in fact, the traditional GNP measure needs to be modified and extended if it is to measure SD. But for the moment SD is defined as at least non-declining consumption, GNP, or some other agreed welfare indicator.

The conditions for sustainable development

A more difficult task is to determine the necessary and sufficient conditions for achieving SD. Fundamentally, how do we compensate the future for damage that our activities today might cause? The answer is through the transfer of **capital bequests**. What this means is that this generation makes sure that it leaves the next generation a stock of capital no less than this generation has now. Capital provides the capability to generate wellbeing ('justice as opportunity' and the 'Lockean Standard' notions are relevant in this context – see Chapter 2) through the creation of goods and services upon which human wellbeing depends.

Weak sustainability (WS)

Under this interpretation of SD, it is not thought necessary to single out the environment (natural capital) for special treatment, it is simply another form of capital. Therefore, what is required for SD is the transfer of an *aggregate capital stock* no less than the one that exists now (this then is the **weak sustainability constant capital rule**). We can pass on less environment so long as we offset this loss by increasing the stock of roads and machinery, or other man-made (physical) capital. Alternatively, we can have fewer roads and factories so long as we compensate by having more wetlands or mixed woodlands or more education. WS is, as we pointed out in Chapter 2, based on a very strong assumption, **perfect substitutability** between the different forms of capital.

Strong sustainability (SS)

Under this interpretation of SD, perfect substitution between different forms of capital is not a valid assumption to make. Some elements of the natural capital stock cannot be substituted for (except on a very limited basis) by man-made capital. Some of the functions and services of ecosystems are essential to human survival, they are life support services (biogeochemical cycling) and cannot be replaced. Other ecological assets are at least essential to human wellbeing, if not exactly essential for human survival – landscape, space, and relative peace and quiet. These assets are **critical natural capital** and since they are not easily substitutable, if at all, the SS rule requires that we protect them.

Measuring sustainable development

Another way of looking at the idea that SD means generating human wellbeing now without impairing the wellbeing of future generations is to think about a *sustainable flow of income*. This is a level of income that the nation can afford to receive without depreciating the overall capital stock of the nation. The danger is that a failure to adequately account for natural capital and the contribution it makes to economic welfare and income will lead to misperceptions about how well an economy is really performing. This danger is real because the current system of national accounts used in many countries fails, in almost all cases, to treat natural capital as assets which play a vital part in providing a flow of output/income over time. Extended national accounts (i.e. not restricted to market-based outputs, incomes and expenditure, as measured in the Gross National Product concept) are required in order to improve policy signals relating to SD.

Two adjustments are required, one for the depreciation of natural capital (changes in quantity) and the other for degradation of the natural capital stock (changes in quality). A framework to reflect the use of natural resources at the national level is in the process of being agreed by the United Nations

Statistical Office. However, the theory and practice of making these adjustments is complex and they are not discussed further here (we provide some suggested reading at the end of the chapter). Instead we present a simple test for SD which yields data which is at least indicative of national sustainability. The test is, however, far from a definitive sustainability indicator, but it is based on modified national accounting information.

Simple indicator of sustainable development

One SD rule states that an economy must save at least as much as the sum of the depreciation on the value of man-made and natural capital (Pearce and Atkinson, 1992). An analogy with a business is useful in this context. If our business consistently failed to save enough money to plough back into the business, to replace machinery and buildings as they wear out (depreciate), we might stay afloat for a while but not long term – our business would be unsustainable. The same is true for any economy, its national savings ratio (savings over some measure of income like Gross Domestic Product (GDP)) must be greater than or equal to depreciation in the natural capital and man-made capital stock, if it is to pass our simple sustainability test. Box 4.1 illustrates some sustainability indicators for a selection of countries. Nothing definitive is being claimed since the data available is not always comprehensive and the test itself is 'static' and ignores factors such as technological change, population growth and international trade.

Precautionary principle and safe minimum standards

For some analysts supportive of the strong sustainability position, sustainability constraints (such as the critical natural capital protection rule) should be seen as expressions of the so-called **precautionary principle** and one similar to the notion of **safe minimum standard** (SMS). The SMS concept is one way of giving shape to the intergenerational social contract idea we discussed in Chapter 2. Somehow we have to trade off using resources to produce economic benefits and conservation of resource stocks and flows to guarantee sustainable benefit flows. The trade-off decisions have to be taken within a context of *uncertainty* and possible **irreversibilities** (i.e. decisions once taken result in changes that are physically impossible to reverse or prohibitively expensive to reverse, e.g. loss of tropical forests and complex wetlands). To satisfy the *intergenerational social contract* (via the constant capital rule and capital bequests), the current generation could rule out in advance, depending on the costs (strictly known as the social opportunity costs, i.e. what society has to give up or forgo), development activities that could result in natural capital depreciation beyond a certain threshold of damage cost and

Box 4.1 Test for weak sustainable development

An economy is sustainable if it saves more than the depreciation on its human-made and natural capital.

Country	Gross savings ratio (S/Y)	Depreciation of human-made capital (d_M/Y)	Depreciation of natural capital (d_N/Y)	Sustainability indicator (Z)
Finland	28	15	2	+11
Germany	26	12	4	+10
Japan	33	14	2	+17
UK	18	12	6?	0?
USA	18	12	4	+2

Notes and sources:
Y denotes that the values are expressed as a percentage of GDP.
S/Y is taken from World Bank, *World Development Reports* –
d_M/Y is taken from the UN *System of National Accounts* (UNSO, 1990).
The test takes the form,

$$Z >= S/Y - d_M/Y - d_N/Y$$

Z must be greater than or equal to zero for sustainability.

irreversibility (i.e. loss of critical natural capital, life support services, keystone species and processes) – see Box 4.2. The compatibility between SMS and strong sustainability is not, however, quite complete. SMS says conserve unless the benefits forgone are very large. SS says that, whatever the benefits forgone, loss of critical natural capital is unacceptable.

Sustainable livelihoods

Any sustainable strategy for the future will have to confront the question of how a much larger total global population can gain at least a basic livelihood in a manner which can be sustained. For the people of the South, many of their livelihoods will have to endure in environments which are fragile, marginal and vulnerable. Sustainable livelihoods can only be promoted via policies which reduce vulnerability – e.g. flood protection to guard against sea-level rise induced by climate change due to global warming (see Chapter 19); measures to improve food security and to offset market and intervention failures such as inappropriate resource pricing and uncoordinated development policies (see Chapters 5, 6 and 23).

Box 4.2 Safe minimum standards approach to sustainability

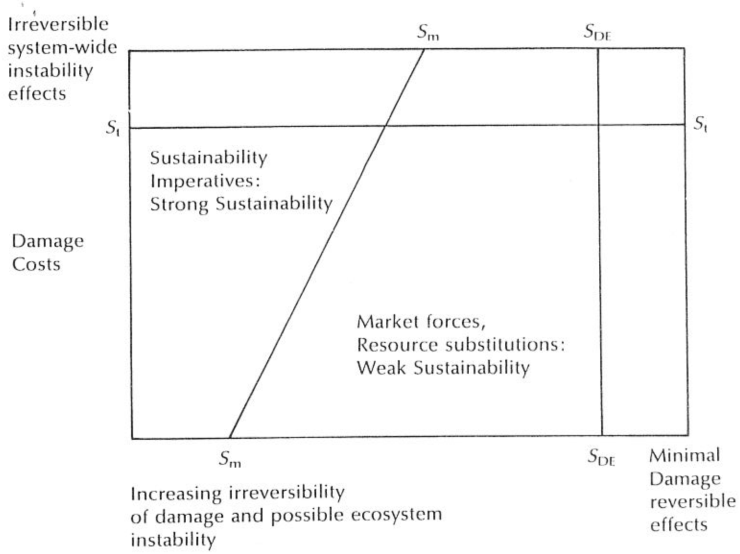

The line S_m–S_m represents some hypothetical safe minimum standard trade-off decision. Supporters of the weak sustainability (technocentric) position(s) might favour a line such as S_t–S_t; while very strong sustainability advocates, such as the deep ecologists, might favour a line such as S_{DE}–S_{DE}.

Source: Adapted from B. Norton, Georgia Institute of Technology, quoted in Toman (1992).

Sustainable development: operational principles

A number of rules (which fall some way short of a blueprint) for the sustainable utilization of the natural capital stock can now be outlined (roughly ordered to fit the VWS to VSS progression):

1. Market and intervention failures related to resource pricing and property rights should be corrected.
2. Maintenance of the regenerative capacity of renewable natural capital (RNC) – i.e. harvesting rates should not exceed regeneration rates – and avoidance of excessive pollution which could threaten waste assimilation capacities and life support systems.
3. Technological changes should be steered via an indicative planning system such that, switches from non-renewable (NRNC) to RNC are fostered; and

Box 4.3 Sustainability practice

Sustainability mode (overlapping categories)	Management strategy (as applied to projects, policy or course of action)	Policy instruments (most favoured)		
		Pollution Control and Waste Management	Raw Materials Policy	Conservation and Amenity Management
VWS	Conventional Cost–Benefit Approach: Correction of market and intervention failures via efficiency pricing; potential Pareto criterion (hypothetical compensation); consumer sovereignty; infinite substitution	e.g. pollution taxes, elimination of subsidies, imposition of property rights		
WS	Modified Cost–Benefit Approach: Extended application of monetary valuation methods; actual compensation, shadow projects, etc.; systems approach, 'weak' version of safe minimum standard	e.g. pollution taxes, permits, deposit-refunds; ambient targets		
SS	Fixed Standards Approach: Precautionary principle, primary and secondary value of natural capital; constant natural capital rule; dual self-conception, social preference value; 'strong' version of safe minimum standard	e.g. ambient standards; conservation zoning; process technology-based effluent standards; permits; severance taxes; assurance bonds		
VSS	Abandonment of Cost–Benefit Analysis: or severely constrained cost-effectiveness analysis; bioethics	standards and regulation; birth licences		

Source: R. K. Turner (1993)

efficiency-increasing technical progress should dominate throughput-increasing technology.

4. RNC should be exploited, but at a rate equal to the creation of RNC substitutes (including recycling).
5. The overall scale of economic activity must be limited so that it remains within the carrying capacity of the remaining natural capital. Given the uncertainties present, a precautionary approach should be adopted with a built-in safety margin.

Box 4.3 summarizes some of the measures and enabling policy instruments that would be involved in any application of an SD strategy. Succeeding chapters in this book cover these various elements in greater detail.

Conclusions

Although it has been defined in many different, and sometimes contradictory, ways the concept of sustainable development does have both relevance and meaning. Weak and strong versions of the concept can be distinguished, and a rudimentary measure of sustainability can be calibrated. How precisely sustainability principles can be translated into operational practice remains more uncertain. But the framework for general sustainability rules has been set out and will require adaptation to specific economic and environmental circumstances.

Further reading

The basic idea of sustainable development and the constant capital rule are covered in:

D. W. Pearce, A. Markandya and E. B. Barbier, *Blueprint for a Green Economy*, Earthscan, London, 1989

and in the context of developing countries by

D. W. Pearce, E. B. Barbier and A. Markandya, *Sustainable Development: Economics and Environment in the Third World*, Earthscan, London, 1990.

The strong sustainability position is set out in:

R. Costanza and H. Daly, 'Natural capital and sustainable development', *Conservation Biology* **6**: 37–46, 1992.

Modified national income accounting is discussed in:

P. Bartelmus *et al.*, 'Integrated environmental and economic accounting: framework for a SNA satellite system', *Review of Income and Wealth* **37**: 111–48, 1991.

C. Bryant and P. Cook, 'Environmental issues and the national accounts', *Income Trends*, No. 469: 99–122, HMSO, London, 1992.

H. Daly and J. Cobb, *For the Common Good*, Greenprint, London, 1990.

On the safe minimum standard see:

R. C. Bishop, 'Economics of endangered species', *American Journal of Agricultural Economics* **60**: 10–18, 1978.

References

D. W. Pearce and G. Atkinson, 'Are National Economies Sustainable? Measuring Sustainable Development', CSERGE GEC Working Paper 92–11, University College London and University of East Anglia, 1992.

M. A. Toman, 'The difficulty of defining sustainability', *Resources* **106**: 3–6, 1992.

R. K. Turner (ed.) *Sustainable Environmental Economics and Management: Principles and Practice*, Belhaven, London, 1993, Chapter 1.

World Commission on Environment and Development, *Our Common Future*, Oxford University Press, Oxford, 1987.

PART II

The causes of environmental degradation

CHAPTER 7

Cost–benefit thinking

The idea of cost–benefit analysis

Meeting individuals' preferences

Everyone is used to making decisions on the basis of a balance of gains and losses, advantages and disadvantages. The idea behind such a weighing up is that we only do those things that yield us *net* gains, and that, where we have to choose between alternatives, we choose that one which offers the *greatest net gain*. Instead of speaking of gains we could speak of benefits, and instead of losses we can talk of costs. This is the simple foundation of **cost–benefit analysis** (or, in the United States, benefit–cost analysis). However, cost–benefit analysis (CBA) defines costs and benefits in a particular way, and it goes beyond the idea of an *individual*'s balancing of costs and benefits to *society*'s balancing of costs and benefits.

Costs and benefits are defined according to the satisfaction of wants, or preferences. If something meets a want, then it is a benefit. If it detracts from wants, it is a cost. Put more formally, anything is a benefit that increases human wellbeing, and anything is a cost that reduces human wellbeing. For the economist, whether wellbeing has increased or not is to be discovered by looking at people's preferences. If an individual states a preference for situation A to the present condition, then the net benefits of moving to A must be positive for that individual. *Why* A is preferred is not the immediate concern, although no-one would argue that the individual should be allowed to get to situation A if it involves some immoral or illegal act. Subject to the wider considerations about the 'morality' of giving people what they want, CBA functions on the basis that a 'better' allocation of resources is one that meets people's preferences (wants). For the individual, i, we could write that i should accept a proposal to move to situation A if

$$[B_A - C_A] > 0 \tag{7.1}$$

where B is benefit and C is cost, and benefits and costs are measured in terms of people's wellbeing.

Getting a social decision rule

To make this a *social* decision rule for accepting the move we need to know what everyone else prefers. If everyone prefers the move to A, we have no problem. If many people prefer the move and the rest simply don't care (they are indifferent) we also have no problem. In this latter case, those who prefer the move are 'better off' (by definition) and those who are indifferent are neither better off nor worse off. But what happens if some people are made better off (they prefer to move) and some worse off (they prefer being where they are)? In order to decide whether society as a whole is better or worse off we need to compare individuals' gains and losses. Many economists believe that this is not possible. They say it is not possible to compare individuals' wellbeing (in the jargon, they say 'interpersonal comparisons of utility' are impossible – utility is just another word for wellbeing here). But we do in fact make such comparisons on a regular basis. We do judge how other people feel – by how they look, how they behave and what they say. We might add that all policy decisions actually *do* involve such comparisons because it is virtually impossible to find a policy that makes *everyone* better off – someone always loses. If such comparisons can be made, then we could develop some rules for comparing the extent to which the wellbeing of individuals will change because of a given policy.

Willingness to pay (WTP)

In order to pursue the idea of aggregating individual preferences, we should first ask how we might *measure* the gains and losses in wellbeing. One way to do this would be to observe how people vote in a referendum, but this will not tell us about the strength of the preference for or against something. Another way is to look at what people are *willing to pay* for something. A measure of an individual's preference for a good in the market-place is revealed by their willingness to pay (WTP) for that good. A somewhat less familiar idea is that people are WTP to *avoid* something they do not like, or that they might be willing to accept (WTA) compensation for tolerating something they do not like. Chapter 8 looks more closely at these concepts: for now it is necessary only to get an intuitive idea of the link between WTP and human wellbeing.

By looking at what people are WTP *for* a benefit, or WTA to tolerate a cost, we have found a way of measuring the strength of individuals' preferences. Box 7.1 investigates the idea of WTP in a little more detail. The WTP concept

Box 7.1 Willingness to pay

Panel (a) Demand curve.

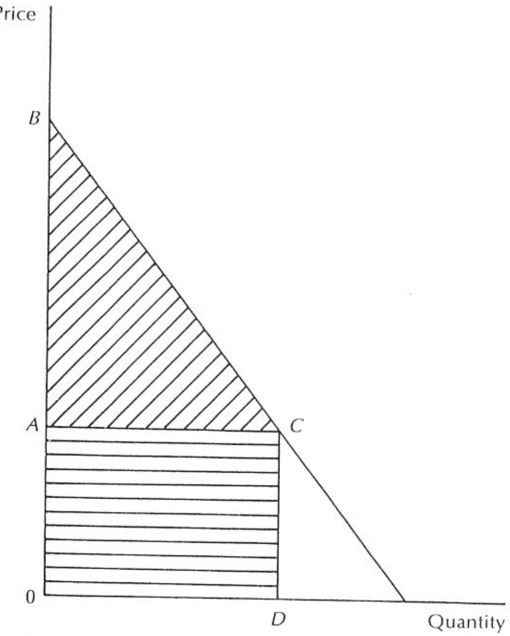

Panel (a) shows a demand curve for a product. This could be a product sold in the market-place (a marketed good) or it could be one which does not have a market (a non-marketed good). In the latter case, it is necessary to find out the demand curve by special means (see Chapter 8).

Suppose the price that we observe is $0A$. The quantity demanded is then $0D$. We can think of the demand curve as a 'willingness to pay curve': it shows the WTP for extra units of the good and is therefore a *marginal willingness to pay* curve. The amount that individuals actually pay in the market-place (or the amount they would pay if there was a market-place) is given by the amount of expenditure $0ACD$. But there are WTP higher prices for the first units, e.g. WTP is $0B$ for the first unit, down to DC for the last unit bought. So WTP exceeds *actual* payment. If we add the excess of WTP over $0A$ (the price paid) for each unit, we get the shaded triangle ABC. This is known as the *consumer's surplus*: it is the benefit received over and above the amount actually paid. The *gross* WTP is $0ACD + ABC = 0BCD$, and this is made up of the actual payment made and the consumer surplus. Alternatively, we could call $0BCD$ the gross WTP and ABC the net WTP. The important concept here is the *net* WTP, since this is a measure of the net gain that consumers secure.

is what we need to solve the problem of aggregating across individuals when some prefer a situation like A above, and some do not. For example, the picture may be as follows:

Individual 1: WTP to move to A = £10
Individual 2: WTP to move to A = £8
Individual 3: WTA to tolerate a move to A = £6
Individual 4: WTA to tolerate a move to A = £5

Individuals 1 and 2 are gainers, and 3 and 4 are losers. Is society as a whole better or worse off from the move to A? The rule we could use is:

$$(WTP_1 + WTP_2) - (WTA_3 + WTA_4) > 0 \tag{7.2}$$

which would give us the result:

$$(10 + 8) - (6 + 5)$$

which comes to +7 which is >0, so the move is socially worthwhile.

Hypothetical compensation

Using this rule to declare the move to A to be socially worthwhile may seem odd. After all, two people are better off and two are worse off. This hardly seems like a good criterion for saying things have improved overall. But to see that there is a net social gain, suppose individuals 1 and 2 are required to compensate individuals 3 and 4 for their losses. Then, individual 1 could give 8 to individual 3, leaving him with 3. Individual 1 would be better off than he was without the move, and individual 3 would be no worse off, having been compensated for the loss. Individual 2 could compensate individual 4 and the same situation would arise. Now it turns out that individuals 3 and 4 are no worse off, and 1 and 2 are better off. This is known technically as a **Pareto improvement**[1]. Or at least, it is an *actual* Pareto improvement if the compensation is actually paid, and a *potential* Pareto improvement if the compensation is 'hypothetically' paid[2].

The social cost–benefit rule embodied in Equation (7.2) provides the foundation of CBA. Instead of WTP we can write *B* for benefit, and instead of WTA we can write *C* for cost. Then Equation (7.2) looks just like Equation (7.1), but it is for the *aggregate* of individuals that make up society, rather than for one individual. So, (7.2) can be rewritten as:

$$\Sigma_i[B_i - C_i] > 0 \tag{7.3}$$

where Σ means 'sum of', and the subscript refers to each individual. We now need to modify this rule for the introduction of time.

Introducing time

Discounting

Equation (7.3) is timeless. It does not indicate the time period over which costs and benefits are being added up. But, as we all know, changes in a situation could involve costs and benefits occurring over long periods of time, or occurring immediately, after which they disappear, or occurring later on. We need to add up benefits in each time period. The simplest way to do this would be to say that a benefit of £1 to individual 1 in year 1 should be added to a £1 benefit in year 2, and so on. This would be correct *if the individual concerned did not care when benefits and costs occurred*. But typically people do care. They prefer to have benefits now rather than later, and costs later rather than now. They have what we call **time preference**. Since the very rationale of CBA is that preferences count, it is essential to take account of the preferences for time in the same way as we do with the measures of WTA and WTP. This brings us to the phenomenon of **discounting**.

The process of discounting can best be understood by looking at the mechanism of compound interest. It we invest £1 at 5 per cent annually compounded, it will be worth £1.63 in ten years' time. Conversely, 61 pence invested now, at the same rate of interest, will be worth £1 in ten years' time. We would then refer to 0.61 as the **present value factor** for a ten-year period when the **discount rate** is 5 per cent. Given this direct relationship between discounting and compound interest, it is evident that the higher is the discount rate, the lower will be the discount factor, and the faster will the discount factor fall, as the time horizon is extended. From this we can see that discounting is really only compound interest back-to-front. If £1 now is worth:

$$£1\ (1 + r)^5$$

in five years' time at a rate of interest of *r* per cent, then £1 in five years' time must be worth *now*:

$$\frac{£1}{(1 + r)^5}$$

This is the *present value* of £1 in five years' time. If we invest this sum at *r* per cent compound interest for five years it will come to:

$$\frac{£1}{\quad} \cdot \frac{(1 + r)^5}{(1 + r)^5} = £1$$

Box 7.2 provides some worked examples of calculations of present value. More generally, the present value of £1 in year *t* is:

$$\frac{£1}{(1 + r)^t}$$

Box 7.2 Computing present values

The basic formula for computing a present value is:

$$\frac{B_t}{(1+r)^t}$$

where B_t is the benefit in year t and r is the discount rate. The same formula would hold for costs – simply substituting C for B. The general formula for computing a present value of a set of benefits *and* costs that occur through time, known as a **net present value** (NPV) would be:

$$\Sigma_t \frac{B_t - C_t}{(1+r)^t}$$

The CBA rule then, is that for any policy or project, the NPV should be positive.

To illustrate the above rule, consider a project that has the following sequence of costs and benefits:

	Year 1	Year 2	Year 3	Year 4	Year 5
Cost	30	10	0	0	0
Benefit	0	5	15	15	15
Net benefit	−30	−5	15	15	15

Note that costs appear as minuses and benefits as pluses. Suppose the discount rate, r, is 10 per cent (which is written as 0.1). Then the computation is:

$$-30/1.1 - 5/(1.1)^2 + 15/(1.1)^3 + 15/(1.1)^4 + 15(1.1)^5$$

Typically this would be done with a discounted cash flow computer program, but in this case the calculations are simple:

$$-27.3 - 4.1 + 11.3 + 10.3 + 9.3 = -0.5$$

The NPV is negative and therefore the project is not worthwhile. Notice that without the discounting procedure, benefits of 45 exceed costs of 35. Discounting can therefore make a big difference to the ultimate decision to accept or reject a project.

The choice faced by a decision-maker may be a simple 'accept/reject' decision like the one above, but it may also be one of choosing between competing alternatives, e.g. a hydroelectric power plant or a coal-fired plant or a nuclear power station. If each option has a positive NPV, the choice should be made on the basis of the highest NPV. Yet another decision context arises where a number of projects can be chosen but the budget available is limited. The rule then is to *rank* the projects according to the *ratio* of the PV of benefits to the PV of costs (the 'benefit–cost ratio') and work down the ranked list until the budget is exhausted.

The practice of *discounting* arises because individuals attach less weight to a benefit or cost in the future than they do to a benefit or cost now. Impatience, or 'time preference', is one reason why the present is preferred to the future. The second reason is that, since capital is productive, £1's worth of resources now will generate more than a £1's worth of goods and services in the future. Hence an entrepreneur would be willing to pay more than £1 in the future to acquire £1's worth of these resources now. This argument for discounting is referred to as the 'marginal productivity of capital' argument, the use of the word 'marginal' indicating that it is the productivity of additional units of capital that is relevant.

Time and the cost–benefit rule

We can now modify the CBA rule in Equation (7.3). With time incorporated into the approach, we have as our decision rule that any move is potentially worthwhile if:

$$\Sigma_t(B_t - C_t)\,(1+r)^{-t} > 0 \qquad (7.4)$$

where the subscript t now refers to time and we take it for granted that B and C are now aggregated across individuals. In order to highlight the environment, we need to remind ourselves that both B and C in Equation (7.4) *include* environmental costs and benefits. That is, if there is a change in environmental quality that people like (their preferences for the change are positive), then we need to find their WTP for it. This WTP will appear as part of the benefit of any policy that includes the change in environmental quality. Equally, if the change in the environment is for the worse, then we seek the WTP to avoid that change, or the WTA to put up with it. This value will appear as part of C in Equation (7.4).

Environment and the cost–benefit rule

To highlight environmental costs and benefits, we will write them separately as E. Equation (7.4) then becomes:

$$\Sigma_t\,(B_t - C_t \pm E_t)\,(1+r)^{-t} > 0 \qquad (7.5)$$

Equation (7.5) is the fundamental equation of cost–benefit analysis: it tells us that for any project or policy to be regarded as potentially worthwhile, its non-environmental benefits (B) less its non-environmental costs (C) plus or minus the value of the environmental change (E), all discounted to a present value, must be positive.

Box 7.3 illustrates an actual cost–benefit study.

Box 7.3 Cost–benefit in practice: case study of lead in gasoline

Under Executive Order 12291 of 1981 in the United States, government agencies were required to use 'Regulatory Impact Analysis' (RIA) and to adopt regulatory processes that would maximize 'the net benefits to society'. The Order was the first to establish the net benefit objective as the criterion for adopting regulatory processes, although its adoption has been circumscribed by existing laws relating to other objectives.

Benefit–cost analysis played an important role in the adoption of regulations concerning lead in gasoline. Ambient lead concentrations were thought to be linked to serious health effects, including retardation, kidney disease and even death. The Environmental Protection Agency (EPA) conducted a benefit–cost study with the results shown in Table 1.

The regulation involved reducing lead in gasoline from 1.1 grams per gallon (g.p.g.) to 0.1 g.p.g. The costs of the rule are shown as 'total refining costs'. Refinery costs increase because lead has traditionally been used to boost octane levels in fuel, and other means would have to be found to achieve this. The benefits included:

(a) improved children's health;
(b) improved blood pressure in adults;
(c) reduced damages from misfuelled vehicles, arising from hydrocarbon, NO_x and CO emissions; and
(d) impacts on maintenance and fuel economy.

Children's health

The EPA study found that blood lead levels closely tracked trends in gasoline lead. Medical costs for the care of children would be reduced by reducing lead concentrations, and there would be less need for compensatory education for IQ-impaired children. These savings are shown as 'children's health effects' in Table 1.

Table 1 Year-by-year costs and monetized benefits of final rule, assuming partial misfuelling (milliions of 1983 dollars)

	1985	1986	1987	1988	1989	1990	1991	1992
Monetized benefits								
Children's health effects	223	600	547	502	453	414	369	358
Adult blood pressure	1724	5897	5675	5447	5187	4966	4682	4691
Conventional pollutants	0	222	222	224	226	230	239	248
Maintenance	102	914	859	818	788	767	754	749
Fuel economy	35	187	170	113	134	139	172	164
Total monetized benefits	2084	7821	7474	7105	6788	6517	6216	6211
Total refining costs	96	608	558	532	504	471	444	441
Net benefits	1988	7213	6916	6573	6284	6045	5772	5770
Net benefits excluding blood pressure	264	1316	1241	1125	1096	1079	1090	1079

Adult blood pressure

Blood lead levels were thought to be associated with blood pressure and hypertension. Medical costs would be saved if these illnesses could be reduced. Moreover, some heart attacks and strokes would be avoided. A value of a 'statistical life' of $1 million was used for the latter. The resulting values show up in the 'adult blood pressure' row of the table. They are seen to be high because of the involvement of mortality-avoidance in this benefit.

Other pollutants

Reducing lead in gasoline also reduces other pollutants. This is because making unleaded fuel the 'norm' reduces the risk of 'misfuelling', i.e. using leaded fuels in vehicles designed for unleaded fuels. The mechanism whereby misfuelling is reduced is through the higher cost of leaded fuels at the new low-lead concentration. This deters drivers from purchasing the leaded fuel. As misfuelling is reduced, so emissions of HC, NO_x and CO are reduced. Damage done by these pollutants was estimated by studies of ozone pollution damage (ozone arises from HC and CO emissions), but estimates were also made of the value of the equipment destroyed by misfuelling. The figures appearing in the row 'conventional pollutants' in the table, are, in fact, an average of the two methods.

Maintenance and fuel economy

Maintenance costs for vehicles were expected to fall due to reduced corrosive effects of lead and its scavengers on engines and exhaust systems. Fewer engine tune-ups and oil changes would be needed, exhaust systems would last longer. Fuel economy was expected to rise as the new technologies to raise octane levels to what they were previously also increase the energy content of fuels. There would also be reduced fouling of oxygen sensors. Maintenance benefits outweighed fuel economy benefits by around 6 to 1.

The net benefits from reducing lead in gasoline are seen to be substantial, even if the blood pressure benefits (which dominate the aggregate benefits) are excluded. Indeed, inspection of the table shows that the regulation would be worthwhile *even if all health benefits were excluded.* In the event, the blood pressure benefits were excluded from the final decision because the research establishing this link was judged too recent to permit adequate review. The lead regulation was also of interest because of the introduction of a 'lead permits system' to reduce the financial burden on the refining industry. Essentially, this allowed 'lead quotas' to be traded between refiners. Refiners who found it easy to get below the limit were allowed to sell their 'surplus' lead rights to refiners who found it expensive to get back to desirable octane levels without lead. The particular feature of the lead-in-gasoline benefit–cost study that made it a powerful aid to decision-making was the clear-cut nature of the net benefits even when uncertainties about benefits were allowed for.

Discounting and the environment

Since discounting attaches a *lower* weight to benefits and costs in the future, it has some unfortunate effects as far as the environment is concerned. Box 7.4 shows what is sometimes known as the 'tyranny' of discounting. Examples of how discounting militates against the interests of future generations are found in the following cases:

(a) Where the *environmental damage* done by, say, a project occurs far into the future, discounting will make the present value of such damage considerably smaller than the actual damage done. Examples of such projects might be the possible damage from stored nuclear waste and long-lived persistent micropollutants such as heavy metals.

(b) Where the *benefits* of a project accrue to people 50 or 100 years hence, discounting will lower the value of such benefits and make it difficult to

Box 7.4 The 'tyranny' of discounting

A simple example shows the way in which discounting can shift heavy costs to future generations. Imagine a nuclear waste repository where safety conditions are of the highest order. Some nuclear waste remains very radioactive for hundreds of years. As with many waste dumps in the past, societies often 'forget' where waste has been dumped, or relax their controls. If this happens with a nuclear waste repository, then it could produce a grave risk for future generations due to leakage. Hypothetically, imagine that the cost of such a leakage is £1 billion (£1000 million) in 100 years' time. If the discount rate is 8 per cent, say, then the present value of the damage from the leaked material 100 years from now would be:

$$\text{£1000 million} \times \frac{1}{(1.08)^{100}} = \text{£1000 million/2200} = 0.45 \text{ million} = \text{£450 000}$$

The £1 billion damage is recorded in the CBA as only £450 000!

justify the project or policy. An example might be afforestation, particularly in slow-growing hardwood trees in temperate climates.

(c) Where the decision to extract a resource is affected by the discount rate. Chapter 16 will show that exhaustible resources are more likely to be used up quickly the higher is the discount rate, leaving less for future generations.

Because of this apparent discrimination against the future, environmentalists frequently object to discounting *per se*.

There is, in fact, no unique relationship between high discount rates and environmental deterioration. High rates may well shift the cost burden to future generations for the reasons given above, but, as the discount rate rises so the overall level of investment falls, thus slowing the pace of economic development in general. Since natural resources are required for investment, the demand for such resources is lower at higher discount rates. High discount rates may also discourage development projects that compete with existing environmentally benign uses, e.g. watershed development as opposed to existing wilderness use. Exactly how the choice of discount rate impacts on the overall profile of natural resource and environment use is thus ambiguous. This point is important because it indicates the invalidity of the more simplistic generalizations that discount rates should be lowered to accommodate environmental considerations.

One approach to determining a discount rate is through the formula:

$$s = p + u.c$$

where s is the social time preference rate (society's rate of discount); p is the 'pure time preference rate', i.e. the discount rate that arises simply because people prefer the present over the future ('impatience'); c is the growth rate of real consumption per head and u is a measure of the rate at which the extra wellbeing (or 'utility') arising from consumption declines as consumption rises. (In technical language, u is known as the elasticity of the marginal utility of consumption.) Suppose, for simplicity, that $u = 1$. Then:

$$s = p + c$$

Many people feel it is simply wrong to discount the future because today's generation is 'impatient'. If we reject pure time preference on these grounds we are left with:

$$s = c$$

i.e. the social rate of discount simply becomes the growth rate of the economy (measured as growth in per capita consumption).

Environmentalists point to the *presumed* positive value of c in the social time preference rate formula. First, they argue that there are underlying 'limits' to the growth process (see Chapter 3). We cannot expect positive growth rates of, say, 2–3 per cent for long periods into the future because of natural resource constraints or limits on the capacity of natural environments to act as 'sinks' for waste products. A second concern highlights the problems of particular regions. Chapter 3 showed that in some low income countries real per capita consumption *fell* in the 1980s. That is, c was negative. Does this mean the social discount rate should be negative? Arguably, it should, although past negative growth may not be relevant to a discount rate based on expected future growth. More significantly, the pure time preference component of a social discount rate could be argued to be very high. Real borrowing rates in poor economies are often of the order of 10–15 per cent and offer a first guess at personal time preference rates. So, the approach described above for deriving a social discount rate is inconclusive.

An alternative approach is to look at the **opportunity cost of capital**. This is obtained by looking at the rate of return on the best investment of similar risk that is displaced as a result of the particular project being undertaken. It is only reasonable to require the investment undertaken to yield a return at least as high as that on the alternative use of funds. For example, if private business can earn 10 per cent rate of return, governments should earn as least as much, otherwise the money is better allocated to the private sector. This is the basic justification for an opportunity cost discount rate.

The environmental literature has made some attempts to discredit discounting on opportunity cost grounds. As these criticisms are fairly complex, they are not discussed here (see Markandya and Pearce, 1991). But the general feeling is that opportunity cost arguments lead to rates that are 'too high' and these discriminate against the environment.

It is widely accepted that a benefit or cost should be valued less the more

uncertain is its occurrence. After all, if you are offered £1 now for certain, or £1 next year with some uncertainty as to whether you would actually get the £1, you are virtually certain to prefer the £1 now. The types of uncertainty that are generally regarded as being relevant to discounting are:

(a) uncertainty about whether an individual will be alive at some future date (the 'risk of death' argument);
(b) uncertainty about the preferences of the individual in the future; and
(c) uncertainty about the size of the benefit or cost.

The risk of death argument is often used as a rationale for the impatience principle itself, the argument being that a preference for consumption now rather than in the future is partly based on the fact that one may not be alive in the future to enjoy the benefits of one's restraint. The argument against this is that although an individual may be mortal, 'society' is not and so its decisions should not be guided by the same consideration.

Second, uncertainty about preferences is relevant to certain goods and perhaps even certain aspects of environmental conservation. Economists generally accept that the way to allow for uncertainty about preferences is to include **option value** in an estimate of the benefit or cost; see Chapter 8.

The third kind of uncertainty is relevant but the difficulty is in allowing for it by adjusting the discount rate. Such adjustments assume that the scale of risks is increasing exponentially over time. Since there is no reason to believe that the risk factor takes this particular form, it is inappropriate to correct for such risks by raising the discount rate. This argument is in fact widely accepted by economists but the practice of using risk-adjusted discount rates is still quite common among policy-makers.

The extent to which the interests of future generations are safeguarded when using positive discount rates is a matter of debate. For example, generations 'overlap' – parents, children and grandchildren exist at the same time. Each generation has 'altruism'. Altruism is said to exist when the wellbeing of the current generation is influenced not only by its own consumption, but also by the utility of future generations. This is modelled by assuming that the current generation's wellbeing is also influenced by the wellbeing of the second generation and the third generation. Each generation's discount rate is therefore likely to take account of the next generation's interests. This approach goes some way towards addressing the question of future generations but it does so in a rather specific way. Notice that what is being evaluated here is the current generation's judgement about what the future generations will think is important. It does not therefore yield a discount rate reflecting some broader principle of the rights of future generations. The essential distinction is between the first generation judging what next generations want (**selfish altruism**) and the first generation using resources so as to leave the next ones with the maximum scope for choosing what they want (**disinterested altruism**).

Conclusions on discounting

We have spent some time on the issue of discounting, although we have hardly done justice to the many and frequently complex arguments that are debated on the issue. There are no firm conclusions. We have to remember that the *purpose* of discount rates is to discriminate against the future in order to reflect the underlying value judgement of CBA, namely, that individuals' preferences count. Most of the criticism of discounting rejects the underlying rationale for discounting. It introduces a separate and contradictory objective, namely being fair to the future. It should not, therefore, be too surprising to find that we cannot find a convenient rule for adjusting the discount rate to reflect these contradictory objectives.

A sustainability approach

Chapter 4 introduced the idea that all forms of capital should be maintained intact in some sense, or even enhanced. And it introduced the much stricter idea that the stock of *natural* capital should be kept intact. Is it possible to build these ideas into CBA? Requiring that no project or policy should contribute to environmental deterioration would be absurd. But requiring that the overall *portfolio* of projects should not contribute to environmental deterioration is not absurd. One way to meet the sustainability condition is to require that any environmental damage be *compensated* by projects specifically designed to improve the environment. So, imagine an agency with ten road building projects each of which does some harm to the environment. The idea would be that, subject to assurances that environmental damage was being minimized in each case, the *cumulation* of environmental damage would be offset by a project specifically aimed at enhancing the environment. This may well mean that nine of the projects go ahead, but the tenth is abandoned in favour of the environmental project.

Conclusions

The sustainability approach has some interesting implications for project appraisal. One of these is that the choice of discount rates problem largely disappears. The goal of adjusting discount rates to capture environmental effects is better served by the sustainability condition. If this is correct, the sustainability condition deserves more investigation. Although it may have quite radical implications, it offers the prospect of avoiding belabouring the 'tyranny of discounting' and of asking that all ethical and environmental concerns be accounted for by discount rate adjustment.

Notes

1. Vilfredo Pareto (1848–1923) was an Italian economist who formulated various rules about when society could be said to be better or worse off in terms of welfare (wellbeing). His work laid the foundation for modern **welfare economics**.
2. The variant of the Pareto rule in which people are hypothetically compensated is known as the Kaldor–Hicks rule after Lord Nicholas Kaldor and Sir John Hicks.

Further reading

A. Markandya and D. W. Pearce, 'Development, the environment, and the social rate of discount', *World Bank Research Observer*, **6** (July): 137–52, 1991.

The elements of cost–benefit analysis are explained in

D. W. Pearce, *Cost–Benefit Analysis*, Macmillan, Basingstoke, 1986.

The literature on discounting is quite complex and there are no easy expositions. See:

D. W. Pearce, A. Markandya and E. B. Barbier, *Blueprint for a Green Economy*, Earthscan, London, 1989, Chapter 6.

CHAPTER 8

Valuing concern for nature

Imputing values for non-market goods and services

In the market-place, the individual has fairly clear information (depending on the degree to which advertising is informative rather than persuasive) on which to base their valuation and choices. The product tends to be visible, its characteristics are generally well known, and it has a market price. The individual, on the basis of the available information, weighs up quantity, quality and price on offer. But as we have seen in Chapters 1 and 5, environmental goods and services often have no market price tag and a considerable amount of uncertainty can surround their true value and significance. Many of the environmental assets are also public goods and this is another characteristic that makes it difficult for markets to evolve in such assets.

Given that in the real world, individual consumers and policy-makers have to make trade-offs, it is fundamentally important to know what is being traded-off against what. It is not possible to know this unless we have some idea of the *economic value of environmental assets*. To make comparisons involving an *unpriced good or service*, it is necessary to **impute a value**. The discipline of economics has developed techniques whereby such values can be imputed.

As we saw in Chapters 1 and 7, in the market-place individuals exercise choice by comparing their *willingness to pay* with the price of the product. Imputing values involves finding a willingness to pay measure in circumstances where markets fail to reveal this information. We also noted in Chapter 2 that the values that count belong to those actually exercising the choice: the current generation. But it is a particular feature of environmental costs and benefits that they often accrue to people in generations yet to come (intergenerational incidence of costs and benefits). Counting only the current generation's preferences biases choice against future generations unless there is some built-in mechanism to ensure that current generations choose on

behalf of future generations and take their interests into account. We have reviewed the various ethical arguments about intergenerational fairness in Chapter 2, and the further bias against future generations that can be present in economic decision-making because of the practice of *discounting* in Chapter 7. In this chapter we will concentrate on the problems that economists face when they attempt to impute values for non-market environmental assets. Our central message in this chapter is:

> economic (monetary) valuation of non-market environmental assets may be more or less imperfect given the particular asset together with its environmental and valuation contexts; but, invariably, some valuation explicitly laid out for scrutiny by policy-makers and the public, is better than none, because none can mean some implicit valuation shrouded from public scrutiny.

One final caveat is necessary before we analyze the concept of **total economic value** (TEV) and the range of imputation methods that have been developed to estimate the TEV of environmental assets.

The caveat concerns the idea of **intrinsic value in nature** which we looked at in Chapter 2. The TEV which we will be concerned about is related to valuation of preferences held by people (anthropocentric and instrumental value) and does not encompass any value which may intrinsically reside 'in' environmental assets.

In Box 8.1 we take a simplified and hypothetical example of a woodland resource which, among other things, is providing a place for people to enjoy some recreation. The recreation service is initially an unpriced environmental service, but we could observe some consumer behaviour which will give us a clue to demand for recreation, i.e. the number of recreational visits which an average individual makes to this woodland park every year. We can establish a demand curve by varying the 'price' of a visit and seeing how many visits our average consumer makes. How do we impute or simulate the price variable? In practice, a variety of methods can be utilized, some of which we will review later in this chapter, but we start with a simple approach and assume the imposition of an entrance fee.

Suppose that we find that at an entrance fee in excess of £15.00 the individual is unwilling to make any visits. However, at a fee of £15.00 the individual would make just one visit per year, i.e. the maximum WTP for the first visit is £15.00. Now suppose that the maximum entrance fee that the individual is WTP for a second visit is, say £8.50. Maximum WTP for further trips will similarly decline, say £4.00 for a third visit, £2.00 for a fourth and just £0.50 for a fifth visit. Indeed let us suppose that the individual would only make a sixth visit if there was no entrance fee, i.e. his/her maximum WTP for the sixth visit is £0.00. So if, in reality, there is in fact no entrance fee for walking in the woodland then our average individual will indeed make six

Box 8.1 The demand for recreational visits to a wood

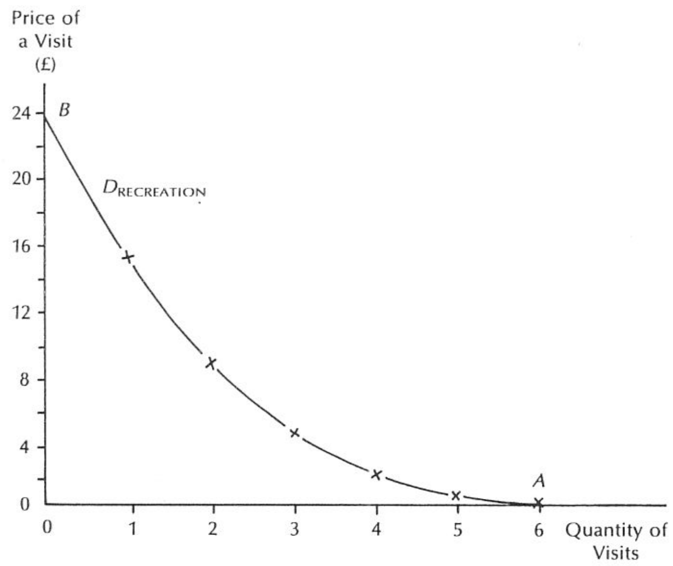

Table 1 WTP, price paid and consumer surplus for a visit

(1) Visit No.	(2) WTP (£)	(3) Price Paid (£)	(4) Consumer Surplus (£)
1	15.00	0.00	15.00
2	8.50	0.00	8.50
3	4.00	0.00	4.00
4	2.00	0.00	2.00
5	0.50	0.00	0.50
6	0.00	0.00	0.00
Totals (£)	30.00	0.00	0.00
	Total Value (total WTP)	= Total Price Paid	+ Total Consumer Surplus

The actual entrance fee paid for a visit is zero. Here the demand curve ($D_{RECREATION}$) shows the maximum amount which a visitor would be WTP for each visit to a wood (column 2 of Table 1). Summing this for each visit gives the visitor's total value for all visits made (area *0AB*). Here there is, in fact, no (or rather a zero) market price for a visit, i.e. entrance is free (another way of looking at this is to imagine that the supply curve runs along the horizontal axis, i.e. at zero price, any amount of visits are supplied). Therefore the total price paid is zero (i.e. the price line is also a flat line running along the horizontal axis) and the total value area is equal to the total consumer surplus.

visits per annum (i.e. all visits where the WTP is equal to or exceeds the price (here zero) of a visit).

Plotting out these price/quantity of visit combinations, we obtain our demand curve $D_{RECREATION}$ in Panel (b) of Box 8.1. We can now calculate the total value of all visits as the total of the amounts which the individual would have been willing to pay for each of the visits in turn, i.e. £15.00 + £8.50 + £4.00 + £2.00 + £0.50 + £0.00 = £30.00. Now because there is in reality no entrance fee to the wood (i.e. no supply constraint) then the actual price paid by the individual for a visit is zero. We can now calculate total consumer surplus as the difference between the total value and the actual price paid, i.e. £30.00 – £0.00 = £30.00. Here we can see that, for unpriced goods, total consumer surplus equals the total value of the good. This is often the case for many environmental goods. Therefore in Panel (b) the total value is (as always) given by the area under the demand curve (*0BA*) which, for an unpriced good only, is also equal to the total consumer surplus provided by the good.

Total economic value

Environmental economists have gone a considerable way towards a classification of economic values as they relate to natural environments. The terminology is still not fully agreed, but the approach is based on the traditional explanation of how value occurs (i.e. it is based on the interaction between a human subject, the valuer, and objects – things to be valued). Individuals have a number of held values which in turn result in objects being given various assigned values. In order, in principle, to arrive at an aggregate measure of value (total economic value) economists begin by distinguishing **user values** from **non-user values**.

Box 8.2 illustrates the use and non-use values which a multiattribute

Box 8.2 The total economic value of woodland

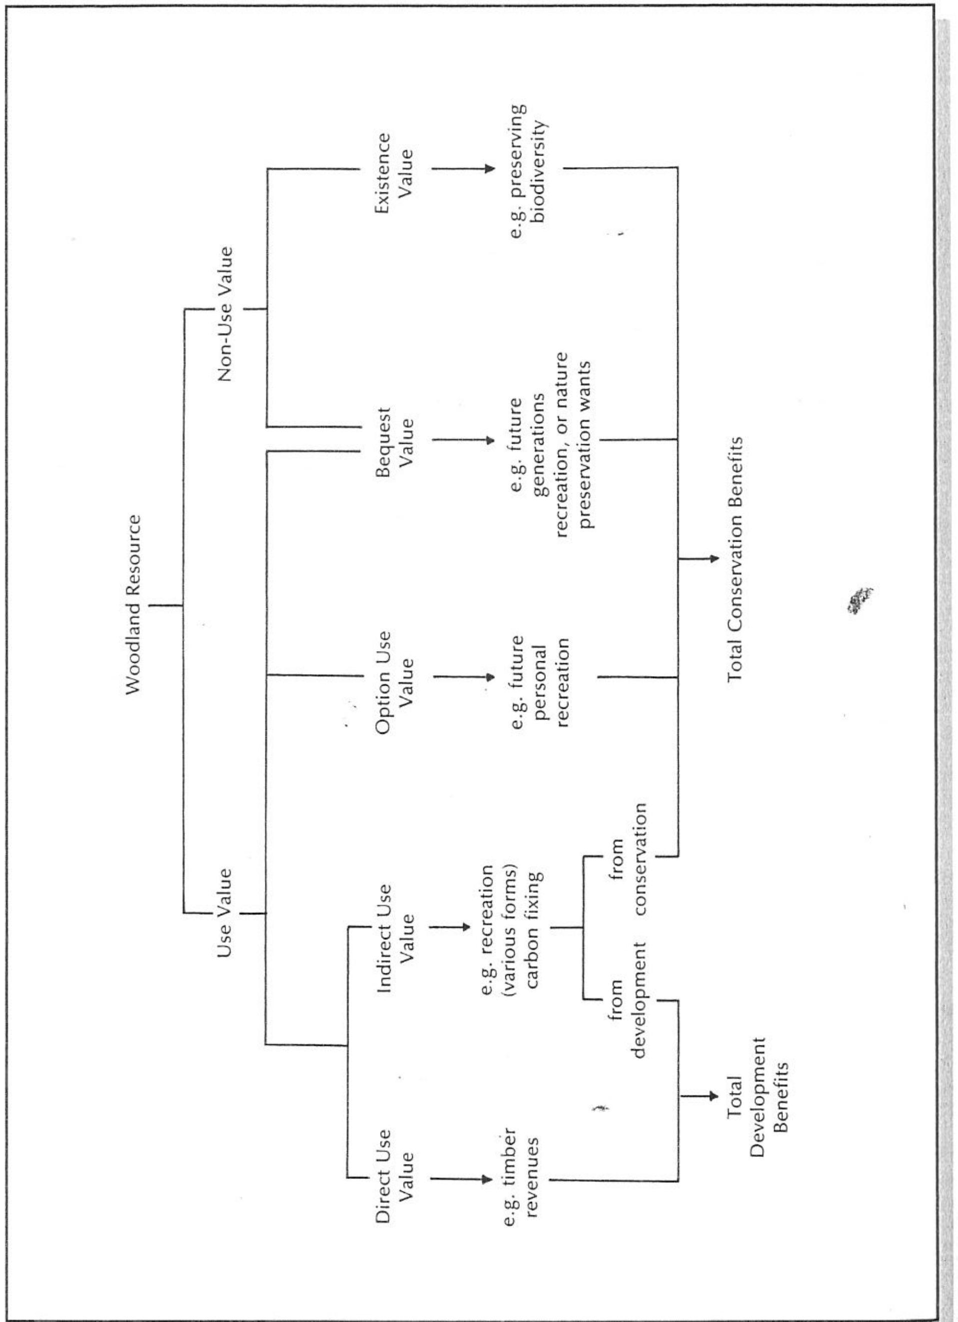

environmental asset such as, for example, a woodland provides. By definition, use values derive from the actual use of the environment. Slightly more complex are values expressed through options to use the environment (*option values*) in the future. They are essentially expressions of preference (willingness to pay) for the conservation of environmental systems or components of systems against some probability that the individual will make use of them at a later date. Provided the uncertainty concerning future use is an uncertainty relating to the 'supply' of the environment, economic theory indicates that this option value is likely to be positive. A related form of value is **bequest value**, a willingness to pay to preserve the environment for the benefit of one's descendants. It is not a use value for the current individual valuer, but a potential future use value or non-use for their descendants.

Non-use values are more problematic. They suggest non-instrumental values which are in the real nature of the thing but unassociated with actual use, or even the option to use the thing. Instead such values are taken to be entities that reflect people's preferences, but include concern for, sympathy with, and respect for the rights or welfare of non-human beings. These values are still anthropocentric but may include a recognition of the value of the very existence of certain species or whole ecosystems. *Total economic value is then made up of actual use value plus option value plus existence value.*

During the 1980s more extensive use of monetary valuation methods was combined with technical improvements in techniques. The result is a large literature consisting of a wide diversity of valuation case studies.

Some scientists have argued that the full contribution of component species and processes to the aggregate life support service provided by ecosystems has not been captured in economic valuation. There does seem to be a sense in which this scientific critique of the *partial* nature of economic valuation has some validity; not in relation to individual species and processes but in terms of the prior value of the aggregate ecosystem structure and its life support capacity. Thus the aggregate ecosystem could be said to possess **primary value**. The prior existence of a 'healthy' ecosystem is necessary before the range of use and non-use values, linked to the ecosystem's structure and functions (see Box 8.2, the woodland ecosystem), can be utilized by humans. We would therefore label all use and non-use values as **secondary values**. It is the various components of total secondary value that are included in total economic value (TEV). But the primary value of the total system is *not* encompassed by TEV.

It is also the case that TEV may fail to fully capture total secondary value (underestimation of 'true' TEV). This is because some ecosystem functions and processes are difficult to analyze scientifically, as well as to value in monetary terms. The **indirect use** values of ecosystems (see Box 8.2) are often surrounded by uncertainty, and the distinction between these values and non-use values is far from clear cut once we realize how complicated and interrelated natural environments actually are in practice. This has recently led some economists to coin a new term, replacing non-use values with **passive use** values. This latter term does seem to better capture the fuzzyness and uncertainty that surrounds the distinction between use and non-use values.

Approaches to valuation

Box 8.3 illustrates *one* (a number of different typologies exist in the literature) way in which the various approaches and methods of monetary valuation can be classified, in the context of environmental resources. Two basic approaches are distinguished, those which value a commodity via a demand curve (Marshallian or Hicksian) and those which do not and therefore fail to provide 'true' valuation information and welfare measures. These latter methods are, however, still useful heuristic tools in any cost–benefit appraisal of projects, policies or courses of action.

The **dose–response approach** requires the existence of data linking human, plant or animal physiological response to pollution stress. If, for example, a given level of pollution is associated with a change in output then it is usually the case that the output can be valued at market or shadow (adjusted/proxy market) prices (loss of crop output from air pollution). But for situations involving human health, complex questions relating to the value of a human life have to be addressed (strictly, analysts seek to value the increased risk of illness or death).

The **replacement cost** technique looks at the cost of replacing or restoring a damaged asset and uses this cost as a measure of the benefit of restoration, (e.g. costs of cleaning building soiled by air pollution). But application of this technique does require careful thought. It is a valid approach in situations where it is possible to argue that the remedial work must take place because of some other constraint. For example, where there is a water quality standard that is mandatory, then the costs of achieving that standard are a proxy for the benefits of reaching the standard.

Another situation where the replacement cost approach is valid would be where there is an overall constraint (a 'sustainability constraint') not to let environmental quality decline. Wetland ecosystems, for example, have been heavily depleted across the globe and are now 'protected' by an international convention (agreement) known as the Ramsar Convention. In these circumstances, wetland replacement costs (these could be wetland restoration elsewhere in a region, wetland relocation, or new wetland creation) might be allowable as a first approximation of the benefits of future wetland conservation, or wetland loss. The so-called **shadow project** approach relies on such constraints. It argues that the cost of any project designed to restore an environment because of a sustainability constraint (see Chapter 7) is then a minimum valuation of the damage done.

Box 8.3 Methods for the monetary evaluation of the environment

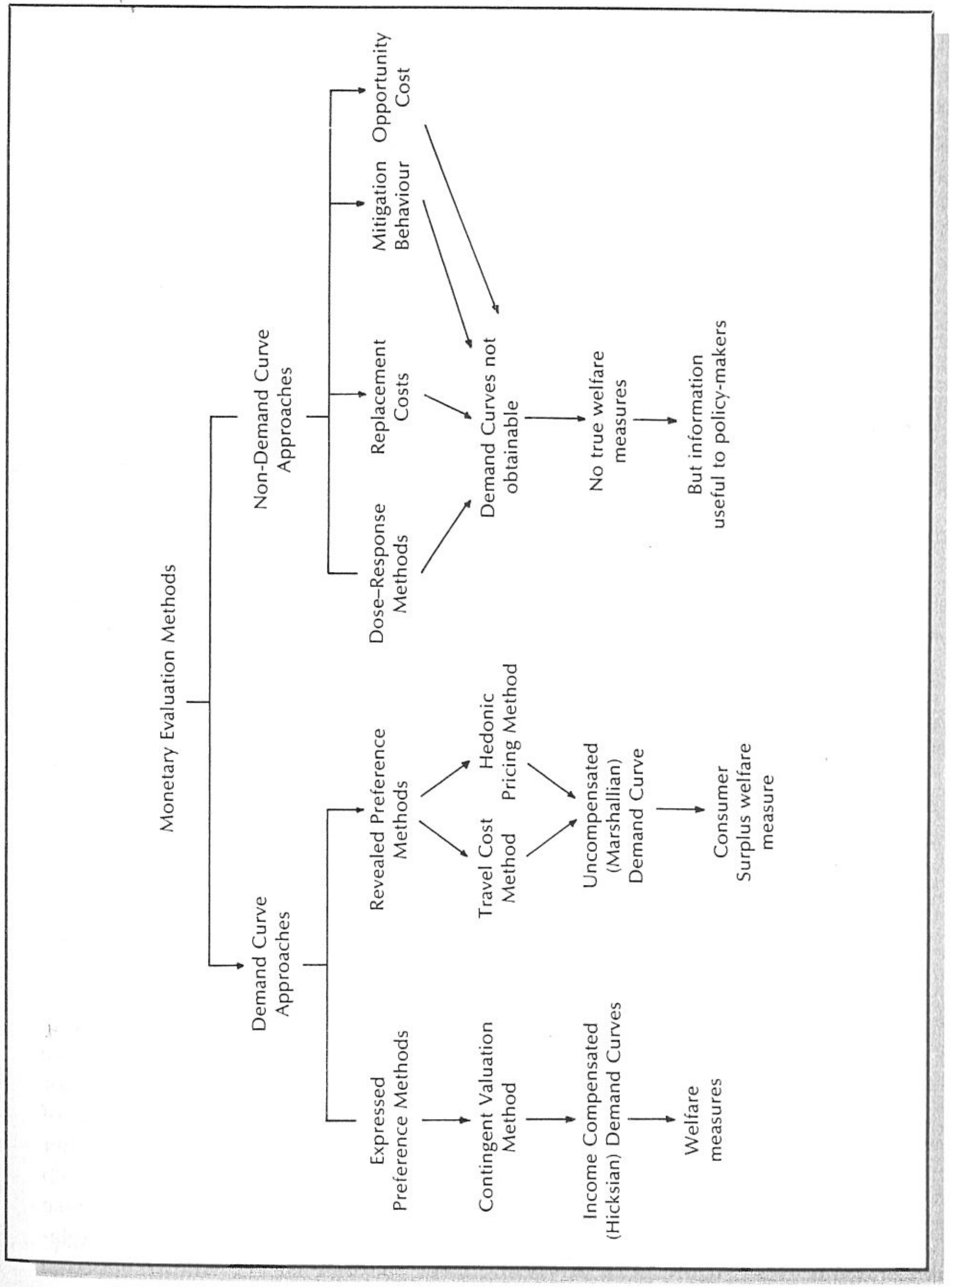

Mitigation behaviour (in terms of **avertive expenditures**) can sometimes be observed in the pollution context. Householders may purchase insulation to 'defend' their homes from noise pollution, as a substitute for a reduction in noise at source.

In the **opportunity cost approach** no direct attempt is made to value environmental benefits. Instead, the benefits of the activity causing environmental degradation – say, drainage of a wetland to allow intensive agriculture – are estimated in order to set a benchmark for what the environmental benefits would have to be for the development (agriculture) *not* to be worthwhile. While this is not a valuation technique, it has proved to be a very useful aid to decision-makers, for example, much of the recent loss of wetlands in Europe due to the operation of the Common Agricultural Policy represents a socially inefficient result because of the heavily subsidized nature of the drainage investments and arable crops that replaced the wetland. Such conversions have now all but ceased as subsidies have been withdrawn or lowered.

Box 8.3 indicates that there are two basic types of demand curve evaluation; firstly, demand can be measured by examining individuals' *stated (expressed) preferences* for environmental goods (elicited via questionnaires); secondly, demand can be *revealed* by examining *individuals' purchases of market priced goods* which are necessary in order to enjoy associated environmental goods.

The **travel cost method** (TCM), which is a revealed preference method, can be used to estimate demand curves for recreation sites and thereby value those sites. The underlying assumption of the TCM is a simple one, that the incurred costs of visiting a site (e.g. petrol costs) in some way reflect the recreational value of that site. Questionnaires are used to ask visitors to the recreational sites where they have travelled from. From visitors' responses, we can estimate their travel costs and relate this to the number of visits per year. Not surprisingly, this relationship generally shows a typical downward sloping demand curve relationship between the cost of a visit and the number of visits taken, i.e. people living a considerable distance from a recreational site (facing high travel costs) make few visits per year, while those living near the site (with low travel costs) tend to make more frequent visits.

Of course, other factors than just travel cost can affect how often people visit a site. For example, if we compare two individuals, one rich and one poor, living the same distance from a site (i.e. facing identical travel costs) we would not be surprised if the rich person made more visits than the poor person. Because of this, analysts usually take into account the income of visitors as one factor explaining the number of visits per year. Other explanatory factors include the number of alternative sites available to each visitor, their personal interest in the type of site, etc. Nevertheless, once these adjustments are made, the analyst can then see the demand curve relationship between the price of visiting a recreational site (i.e. the travel cost)

and the number of visits made. Box 8.4 shows a graph illustrating typical results from a TCM survey with each circle recording the travel costs per visit and number of visits made per year for one visitor (a real life study will typically interview several hundred visitors). From this information, statistical techniques can estimate the 'demand curve' for the site, i.e. the representative relationship between the price of a visit to the site (travel costs) and the number of visits made. This demand curve illustrates, for a typical visitor, how many visits would be made at any particular visit price. The demand curve is then used to obtain the total recreational value for the site. We can multiply this figure by the total number of visits made to the site per year to get an estimate of the total annual recreation value of the site.

The TCM seems at first a relatively straightforward technique based upon the defensible assumption that recreational value must be related to travel cost. However, in practice there are numerous problems with this technique a few of which we raise here.

1. *Time costs.* The underlying assumption of the TCM is that travel costs reflect the recreational value of visiting a site. A simple TCM might assume that the only travel cost is related to petrol expenses, however time is also valuable to people in that time spent during a long car journey cannot be spent doing anything else. There is, therefore, a value of time (a 'time cost') which should be added to the travel cost as a reflection of the true recreational value which the visitor gets from visiting a site. So ignoring time costs is generally believed to lead to a significant underestimate of the recreational value which people obtain from visiting a site. However, what is the value of time? Can we put a price on an hour spent in a car? There have been many attempts to estimate a value of time, for example, by comparing the travel time of differing methods of commuting to work with the costs of those differing measures. However, no real consensus has yet been achieved. A further complication is that many people enjoy travelling, for them the journey to a recreational site is not a cost and may even be a benefit. In such cases we should subtract the time benefit of the journey from its travel costs, i.e. simple TCMs (based on travel costs) may now be overestimating the recreational value of sites in such cases.

2. *Multiple visit journeys.* If an individual visits several sites during a single day's journey but is asked to answer a TCM questionnaire at one of them, then how should analysts apportion the visitor's travel costs? During the day the visitor may have incurred high travel costs, however, only a portion of these reflect the recreational site in question. Conventionally analysts have tried to use a percentage of the day's total travel costs, sometimes asking the visitor to set that percentage. However, the margin for error in this context is uncertain.

3. *Substitute sites.* One visitor may travel 20 miles to visit a site which they particularly enjoy whereas another who has comparatively little enthusiasm for the site may travel the same distance from another direction simply

Box 8.4 Evaluating woodland recreation using the TCM

An on-site questionnaire is used at the forest gate or car park to record how often visitors visit the wood, what their travel costs of a visit are, what their income is, etc. Adjusting for factors such as income we can examine the relationship between travel costs and visits per year as illustrated in Panel (a).

Panel (a) The relationship between the number of visits to a site and the price of the visit.

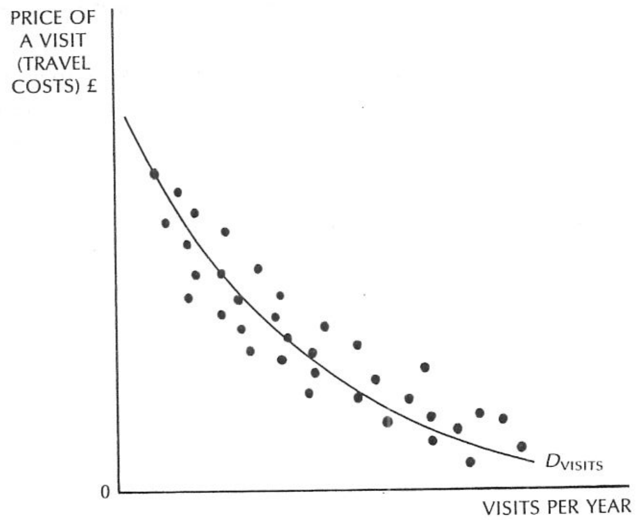

The demand curve D_{VISITS} shows the overall trend relationship between travel costs and visit rates for all the visitors interviewed. Using this information we can estimate the average visitor's total recreational value for the site. Multiplying this by the total number of visitors per annum allows us to estimate the total annual recreational value of the wood itself.

A large scale TCM study of forest recreation in the United Kingdom was undertaken in the late 1980s. Table 1 summarizes results from this study, showing average recreational value per visit at various forests, the annual number of visits to these forests and (multiplying these two together) the estimated total annual recreational value of the forests in question.

Table 1 Estimates of forest recreational values derived from the travel cost method

Location	Average recreation value per visit (£)	Annual number of visits	Total annual recreational value of the forest (£)
Cheshire (Delamere)	1.91	225 000	429 750
Ruthin (Clwyd)	2.52	48 000	120 960
Brecon (Coed Taf)	2.60	41 000	106 600
New Forest (Bolderwood)	1.43	68 000	97 240
Loch Awe (Inverliever)	3.31	3 000	9 930
Lorne (Barcaldine)	1.44	10 000	14 400
Newton Stewart (Glen Trool)	1.61	70 000	112 700
Buchan (Bennachie)	2.26	84 000	189 840
Aberfoyle* (Queen Elizabeth)	2.57	145 000	372 650

Note: All 1988 prices except * 1987.
Source: Adapted from Benson and Willis (1990).

because there is no other available site near their home. Using the simple TCM approach would yield the result that both visitors held the same recreational value for the site, which is clearly incorrect. Some analysts have tried to allow for this by asking visitors to name substitute sites; however, this is both statistically complex and open to error.

4. *House purchase decision*. It may well be that those who most value the recreational attributes of various sites will choose to buy houses near those sites. In such cases, they will incur relatively low travel costs visiting the sites they value so highly, i.e. travel cost will be a gross underestimate of recreational value. Interestingly although this problem has been recognized for many years, it is only very recently that analysts have attempted to include this factor in their questionnaires.

5. *Non-paying visitors*. TCM studies often omit any visitors who have not incurred travel costs to reach the site, e.g. those who have walked from nearby homes. However, this group may well put a very high value on the site.

In summary then, the TCM is grounded upon a simple and fairly well-founded assumption that travel costs reflect recreational value. It allows us to estimate a demand curve and thereby recreational value. However, in practice there are a number of application problems which need to be addressed before we can accept monetary evaluations produced by this technique.

The hedonic pricing method (HPM)

The hedonic price method (HPM) attempts to evaluate environmental services, the presence of which directly affects certain market prices. In practice, by far the most common application of HPM is to the property market. House prices are affected by many factors: number of rooms, size of garden, access to workplace, etc. One such important factor will be local environmental quality. If we can control for the non-environmental factors, e.g. by looking at houses with the same number of rooms, similar garden size, similar accessibility, etc., then any remaining difference in house price can be shown to be the result of environmental differences. For example, in a recent study in Gloucestershire the presence of open water near to a house was shown, on average, to be responsible for a 5 per cent increase in house price (Garrod and Willis, 1991). However, in general the HPM has been applied to the evaluation of environmental costs rather than benefits. For example, just as the water resource raised the value of local housing, so the noise from an airport may lower local house prices. Box 8.5 illustrates a HPM study applied to the valuation of traffic noise.

In order to apply the HPM, the analyst must first collect information regarding all the factors which go to determine house price. This can be a long and arduous task, for example, making measurements of the distance from each house to local facilities such as shops, workplace, etc. Up until recently, these measurements were generally calculated by hand from maps. However, analysts can now massively speed up this operation by the use of a Geographical Information System (or GIS) which holds digitized maps from which such distance calculations can be made by the computer. These GIS computer programs also have a specific application when we are trying to use HPM to place money prices upon landscape values; at the moment most HPM studies just use distance to local landscape features to assess whether they affect house prices. However, GIS maps can include the contours of hills and valleys so that the computer can calculate exactly what is the specific environment affecting each particular house, e.g. whether a house is directly exposed to the noise of a nearby major road or whether it is shielded by other houses. Wider use of this innovative technology may significantly increase the accuracy of HPM studies.

While the HPM approach does appear to be reasonably robust, it nevertheless does have some problems.

1. *User unfriendly*. Estimating the relationship between house price and environmental quality requires a high degree of statistical skill to separate

Box 8.5 Valuing road noise using the hedonic pricing method

The HPM relies upon the assumption that, among other factors (number of rooms, accessibility to shops, workplaces, etc.), the local environmental quality (or lack of it) will determine the price of a house, i.e.:

HOUSE PRICE = *f* (ROOMS, ACCESS, ENVIRONMENT)

The equation states that house price is a function of (*f*) the number of rooms in the house (ROOMS), the distance in miles to local facilities from the house (ACCESS) and some measure of local environmental quality (ENVIRONMENT). Suppose we were interested in valuing the environmental impact of local traffic noise, then we could measure this in terms of decibels of traffic noise inside the houses in question.

We then need to measure each of the items HOUSE PRICE, ROOMS, ACCESS and ENVIRONMENT for a large number of houses so that we can begin to see how, on average, house price changes when each of the influencing factors change. We would expect house price to rise as the number of rooms increase; that house price would fall as the distance to local facilities rises, and finally, for house price to fall as the traffic noise increased, i.e. a typical demand curve relationship. This is indeed the result obtained in an American study of road noise. Table 1 shows the average percentage fall in house price which corresponded to a one unit increase in traffic noise in a number of US areas.

So if a new road scheme was likely to raise traffic noise by one unit in, say, Washington DC, then a monetary value for this increased noise pollution could be found by finding 0.88 per cent of average house prices in the affected area.

Table 1. The impact of traffic noise upon house prices

Area of USA	% fall in house price due to a one unit increase in noise
North Virginia	0.15
Tidewater	0.14
North Springfield	0.18–0.50
Towson	0.54
Washington DC	0.88
Kingsgate	0.48
North King County	0.40
Spokane	0.08
Chicago	0.65

Note: Traffic noise was measured as the equivalent continuous sound level (measured in decibels) which would have the same sound energy over a given period as the actual fluctuation sound level measured at houses in the study.
Source: Nelson (1982)

out the other influences upon house price such as house size, accessibility, etc.

2. *The property market.* The method relies upon the assumption that people have the opportunity to select the combination of house features (size, accessibility, environmental quality) which they most prefer given the constraints of their income.

However, the housing market can often be affected by outside influences, for example, the Government may have a large influence over house prices because of changing tax concessions or interest rates, etc. Similarly, if the HPM is carried out over a large area (e.g. highly urban to very rural) then there may be a cut-off distance at which those who are employed in cities are unable to move further into the country. It may even be that there are different perceptions of landscape in such rural areas. In effect then, the demand curve for houses with different environmental characteristics may be significantly constrained by the supply curve so that the market does not operate freely. In such cases both the demand for and supply of houses will have to be taken into consideration considerably complicating the analysis.

Expressed preference methods: The contingent valuation method (CVM)

Both the TCM and HPM methods have, in some way, relied upon individual valuations of environmental goods as revealed in their purchases of market priced goods (petrol, houses, etc.) which are associated with the consumption of those environmental goods (recreation, peace and quiet, etc.). The contingent valuation method (CVM) bypasses the need to refer to market prices by asking individuals explicitly to place values upon environmental assets. Because of this, the CVM is often referred to as an expressed preference method. Although there are variants of the technique, the most commonly applied approach is to interview households either at the site of an environmental asset, or at their homes, and ask them what they are willing to pay (WTP) towards the preservation of that asset. Analysts can then calculate the average WTP of respondents and multiply this by the total number of people who enjoy the environmental site or asset in question to obtain an estimate of the total value which people have for that asset. Box 8.6 discusses a CVM study estimating the value of improving water quality in a river in Pennsylvania.

An interesting advantage of the CVM approach is that it can, in theory, be used to evaluate resources, the continued existence of which people value, but which they never personally visit. An example of such an asset is Antarctica which people are WTP to preserve but would not in general ever want to visit. Closer to home another example of these 'non-use' values was

provided when a UK forestry firm announced its intention to drain and plant the Flow Country, an important natural wildlife habitat and wetland area in Northern Scotland. Although few people actually visit the area, a CVM study (this time conducted through a postal survey of households) found that individuals were WTP a far higher sum to preserve the area than could ever be produced from growing timber there.

Compared to the methods previously discussed, the CVM approach may appear comparatively straightforward. However, there are a number of potential problems facing the unwary analyst, a few of which we discuss here.

1. *Understating WTP.* The central assumption of the CVM technique is that the WTP sums stated by respondents correspond to their valuation of the assets in question. Critics have questioned the validity of such an assumption claiming that the hypothetical nature of CVM scenarios make individuals' responses to them poor approximations of true value. However, in a series of experiments where hypothetical WTP questions have been followed up by actual requests for money payments, it was found that the sums which people stated they would be WTP were between 70–90 per cent of the amounts they eventually did pay. This indicates that people 'free-ride', i.e. tend to understate what they would really pay in an attempt to reduce any subsequent actual payments. However, as the magnitude of this understatement is relatively small this may not be too serious a problem.

2. *WTP v. WTA.* In theory the payment question can either be phrased as the conventional 'What are you willing to pay (WTP) to receive this environmental asset?' or in the less usual form 'What are you willing to accept (WTA) in compensation for giving up this environmental asset?' When comparisons of the two formats have been carried out analysts have noticed that WTA very significantly exceeds WTP, a result which critics have claimed invalidates the CVM approach showing responses to be expressions of what individuals would like to have happen rather than true valuations. However, recent work has shown that there are good psychological and economic reasons to indicate that individuals feel the cost of a loss (WTA compensation format) more intensely than the benefit of a gain (WTP format). If true, the observed WTA/WTP divergence then actually supports the validity of the CVM. However, it is also true to say that respondents will be far less familiar with the notion of receiving compensation for losing something than they will be with the notion of paying for something, a concept we all meet every day. This is likely to cause far greater uncertainty and variability in answers to WTA questions than occurs with WTP questions. Therefore, the former are to be avoided in favour of the latter. This, in turn, has consequences for the applicability of CVM to certain situations. We can obviously ask people their WTP for an environmental gain (e.g. to set up a new parkland), but in cases of environmental loss we must ask people their WTP to prevent that loss occurring (e.g. to fund flood defences to preserve marshland areas from

Box 8.6 Valuing river-water quality improvements using the contingent valuation method

The Monongahela River is a major river flowing through Pennsylvania, USA. Analysts asked a representative sample of households from the local area what they would be willing to pay in extra taxes in order to maintain or increase the water quality in the river. The analysts conducted several variants of the CVM survey. In one variant households were presented with three possible water quality scenarios and simply asked how much they were willing to pay for each.

Scenario 1: Maintain current river quality (suitable for boating only) rather than allow it to decline to a level unsuitable for any activity including boating.
Scenario 2: Improve the water quality from boatable to a level where fishing could take place.
Scenario 3: Further improve water quality from fishable to swimmable.

Among the households surveyed some used the Monongahela river for recreation while others did not. The analysts therefore could look at how much the users were willing to pay compared to the responses of non-users. Results for the sample as a whole were also calculated. Table 1 presents the willingness to pay, of users, non-users and the whole sample for each proposed river quality change scenario.

A number of very interesting conclusions can be drawn from these results. Considering the results for the whole sample, we can see that the stated willingness to pay sums draw out a conventional demand curve for water quality, i.e. people are prepared to pay a relatively high amount for an initial basic level of quality. However, they are prepared to pay progressively less for higher levels of water quality. Panel (a) draws out the demand curve

Table 1 Willingness to pay (WTP) for river quality scenarios

Water quality scenario	Averge WTP of whole sample ($)	Average WTP of users group ($)	Average WTP of non-users group ($)
Maintain boatable river quality	24.50	45.30	14.20
Improve from boatable to fishable quality	17.60	31.30	10.80
Improve from fishable to swimmable quality	12.40	20.20	8.50

Panel (a) Demand curve for water quality.

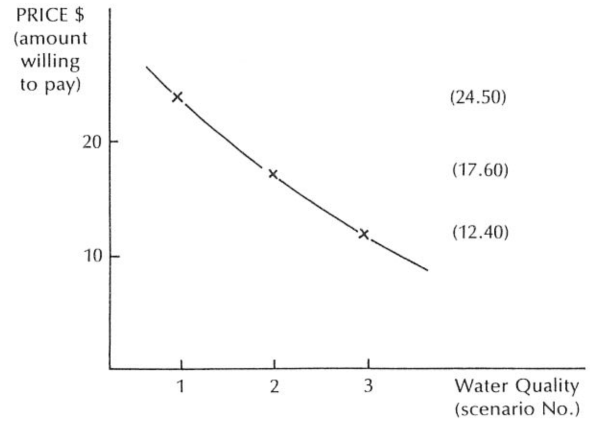

indicated by the results for the whole survey, i.e. for the average household.

From this demand curve we could attempt to calculate the total value of environmental quality at the river. More importantly, we could find the value gain experienced by the average household when a water quality improvement is achieved. The total benefit value of a specific improvement could then be estimated by multiplying this average household value by the number of households which it is thought would be affected by such an improvement. This benefit can then be compared against the cost of achieving such a quality improvement to see if it was worthwhile.

Turning to results for the users and non-users group, we can see that both map out conventional downward sloping demand curves. Furthermore, as we would expect, at every quality level the willingness to pay of the users group exceeds that of the non-users.

Finally, notice that the willingness to pay of non-users is not zero. This is due to the fact that such households, while not personally wishing to visit the river, nevertheless do value its continued existence and even upgrading so that others can enjoy its benefits. This non-use, 'existence' value derives from people's altruistic 'public preferences' showing that the concentration upon people's 'private preferences' as demonstrated by the market prices of marketed goods does not always fully capture the entire range of values which people have for things.

Source: Desvousges *et al.* (1987)

seawater flooding). Nevertheless, it may be that the WTP sum obtained does not reflect what people would consider adequate compensation (their WTA compensation) for losing the environmental asset; indeed the WTP sum may significantly underestimate the true WTA compensation. This problem is a focus of ongoing research.

3. *Part-whole bias.* Critics of the CVM have noted that if people are first asked their WTP for one part of an environmental asset (e.g. one lake in an entire system of lakes) and then asked to value the whole asset (e.g. the whole lake system) the amounts stated may be similar. Why is this? The reason appears to lie in how people commonly allocate their spending; first, dividing their available income up into several broad budget categories (e.g. housing, food, car, recreation) and then subdividing this between the actual items purchased. So for recreation, the first stage is to define the total budget which the individual has available for recreation and then subdivide this into how much they are willing to spend on each site they wish to visit. One approach to this problem is to first ask respondents to work out their overall recreational budget and then ask for their WTP for the environmental asset in question, reminding respondents of their limited recreational budget and that any money they allocate to this asset cannot be spent elsewhere. A second approach is to restrict the use of CVM to the evaluation of broad groups of environmental goods (i.e. wholes rather than parts), again reminding respondents of their limited recreational budget constraints. This restriction, if necessary, would considerably constrain the wide-scale application of the CVM and may itself raise further problems regarding respondents' ability to comprehend such broad amalgams of goods.

4. *Vehicle bias.* When asking a WTP question analysts must specify a realistic route by which such a payment could be made (the 'payment vehicle'). However, respondents may alter their WTP statements according to the specific payment vehicle chosen. For example, in a recent experiment regarding WTP for recreation in the Norfolk Broads, WTP via a charitable trust was noticeably lower than WTP via tax. In this case respondents stated that they doubted the ability of charitable funds to protect the environment and, while they did not like paying taxes, they did feel this was more likely to ensure effective environmental protection. It also compelled a wider group of people to contribute than would have if payment had been via charitable donation. Such results clearly tell us probably as much about the payment vehicles chosen as about the value of the asset in question. An obvious solution to such problems is therefore to use whichever payment vehicle is most likely to be used in reality.

5. *Starting point bias.* Many early studies attempted to prompt respondents by suggesting a starting bid and then increasing or decreasing this bid based upon whether the respondent agreed or refused to pay a such sum. However, it has been shown that the choice of starting bid affects respondents' final WTP sum. As an ancillary part of the water quality study

discussed in Box 8.5, a separate group of respondents were offered a starting bid of \$25 for the first scenario which produced a final average bid in that group of \$27.50. However, when another group were offered a starting bid of \$125 for the same scenario they gave a final average bid of \$94.70, i.e. starting point significantly affected the final bid. In the light of these findings more recent studies have abandoned the use of starting bids.

This chapter has examined a variety of methods which attempt to extend CBA appraisals by expressing the value of non-market environmental goods and services in money units. While many of these methods represent a considerable improvement upon previous practice, we have attempted to indicate that none of them are panaceas and each method has limitations. There is no simple answer to these problems which remain a focus of considerable ongoing research. However, we remain convinced that such valuation methods have an important role to play and, if carefully applied, provide valid and reliable value estimates. Explicit valuation via these methods is preferable to implicit valuations where the link to individuals and their preferences is unclear or non-existent.

Further reading

For a non-technical survey see:

D. W. Pearce, *Economic Values and the Natural World*, Earthscan, London, 1993.

D. W. Pearce and R. K. Turner, *Benefits Estimates and Environmental Decision-Making*, OECD, Paris, 1992.

R. K. Turner and I. Bateman, 'A Critical Review of Monetary Assessment Methods and Techniques', Environmental Appraisal Group Report, University of East Anglia, Norwich, 1990.

R. K. Turner and T. Jones (eds), *Wetlands: Market and Intervention Failures*, Earthscan, London, 1991.

J. Winpenny, *Values for the Environment*, HMSO, London, 1991.

The more technical journal articles referred to were:

J. F. Benson and K. G. Willis, 'The Aggregate Value of Non-Priced Recreation Benefits of the Forestry Commission Estate', Report to the Forestry Commission, Department of Town and Country Planning, University of Newcastle upon Tyne, 1990.

W. H. Desvousges, V. K. Smith and A. Fisher, 'Option price estimates for water quality improvements: a contingent valuation study of the Monongahela River', *Journal of Environmental Economics and Management* **14**: 248–67, 1987.

G. D. Garrod and K. G. Willis, 'The Environmental Economic Impact of Woodland: A Two Stage Hedonic Price Model of the Amenity Value of Forestry in Britain', Countryside Changing Unit, University of Newcastle upon Tyne, 1991.

N. D. Hanley and S. Craig, 'Wilderness development decisions and the Krutilla–Fisher model: The case of Scotland's Flow Country', *Ecological Economics* **4**, 145–64, 1991.

J. P. Nelson, 'Highway noise and property values: a survey of recent evidence', *Journal of Transport Economics and Policy* **XIC**: 37–52, 1982.
R. K. Turner, I. J. Bateman and D. W. Pearce, 'United Kingdom', in S. Navrud (ed.), *Valuing the Environment: The European Experience*, Scandinavian University Press, Oslo, 1992.

CHAPTER 9

Coping with uncertainty

Introduction

Risk and uncertainty are facts of life, and nowhere is this more true than in the environmental context. The reality is that *we often do not know* what the environmental consequences will be of undertaking a particular policy or project. In large part this uncertainty arises because we do not fully understand how ecological systems function, and because we do not know how man-made substances – or increased quantities of 'natural' substances – will interact with the environment. If we did know, then chlorofluorocarbons (CFCs) would probably not have been introduced. We know now that CFCs damage the ozone layer, and that the ozone layer serves valuable protective functions for life on earth (see Chapter 20). Uncertainties of this kind are pervasive. We cannot be sure what is happening with the increasing release of micropollutants into the environment, for example. Sometimes the result of undertaking action without knowing for certain what will happen is that we create *irreversible consequences* such as the elimination of a species. Once gone, we cannot recreate them. In turn, we cannot be sure what will happen if continued species elimination occurs. It may not matter much (from the human standpoint, that is) if species are lost, but we cannot tell. The *scale* of outcome could be large, or small. So, the context of much environmental policy is characterized by:

- uncertainty about the effect
- irreversibility of some effects
- uncertainty about the scale of the effect.

The issue is, then, how do we handle uncertainty? It turns out that there is no easy answer. In what follows we caution readers that the terminology of risk and uncertainty, and the approaches to them are not widely agreed. We have tried to provide a framework for understanding what is a very large and often very confusing literature.