Chapter

7

Social Welfare and Economic Evaluation

The interest of the community then is—what? The sum of the several members who compose it.

Jeremy Bentham

The Nature of Public Choices ♦ Opportunity Sets and Social Welfare Functions ♦ Utility and Income ♦ Utilitarian Social Welfare Functions ♦ Working with Income ♦ Social Welfare and Cost-Benefit Analysis

In the last chapter we discussed how to value individual preferences. In this chapter we discuss how to combine these preferences (utilities) into a measure of collective welfare, which we call social welfare. This measure of social welfare is intended to be a criterion for determining public policy. The greater the social welfare, the more desirable is the relevant policy.

Following the early utilitarian, Jeremy Bentham (1789), social welfare would be viewed as the simple addition of individual valuations of utility. However, individual utilities are not readily comparable or additive. Of special importance, society may want to give greater weight and assistance to those members with less welfare (utility). Moreover, if our measure(s) of social welfare are to guide policy making, they must be practical. In this chapter we confront the conceptual difficulties of estimating social welfare and seek to generate practical measures of welfare for use in policy making.

In the first two sections we discuss the nature of public choices, opportunity sets that represent these choices and measures of social welfare (social welfare functions). The third and fourth sections discuss the relationships between individual utility and income and between individual utility and social welfare respectively. The last part of the chapter discusses how economists work with incomes or dollar values, typically using cost–benefit analysis, to evaluate public policy and how these approaches relate to the concept of maximising social welfare.

The Nature of Public Choices

It might be thought that choice is easy in win–win situations when all parties gain and almost as easy when one party gains and no one loses. However, even in win–win cases, someone may feel relatively deprived by the outcome. Several parties may gain from an inheritance but argue vehemently about the perceived unfairness of the distribution! Or suppose that government decides to cut taxes. While the cuts may benefit most individuals and disadvantage no one, the distribution of the cuts may be strenuously debated. Even in win–

win situations, the preferred allocation of the gains must be determined. This almost always involves determining not only the economic impacts of the allocations but also the relative value of the benefits to the individuals concerned.

It might be supposed likewise that cost-effectiveness decisions (i.e. decisions about how to achieve a given target at least cost) are technical decisions that do not involve normative judgements. Again, this is rarely so. Take, for example, such routine decisions as how to service public hospitals or provide IT services, which may be done by in-house labour or outsourced. This often involves gains or losses of economic rents for some workers so there are winners and losers from any such decision and government may not adopt the least-cost option. Or suppose that government wishes to find the least-cost way to achieve an environmental target such as a maximum level of carbon monoxide or nitrogen oxides in cities. The alternatives will impose various benefits and costs on city residents, commuters, tourists, businesses and so on.

The existence of winners and losers is even more evident when government allocates expenditures. Government may fund hospitals rather than schools or fund both by raising taxes and so reducing private consumption. It may build hospitals or schools in one location rather than another. In almost all social assistance, there are gainers and losers. Also, income redistribution generally means both higher taxes and lower output as some taxpayers and income recipients substitute leisure for work.

It follows that the fundamental problem of public choice is to determine how to assess, compare and aggregate these gains and losses. The assessment is based on valuations of individual preferences discussed in the last chapter. Comparison and aggregation of the costs and benefits borne by individuals are discussed in this one.

Opportunity Sets and Social Welfare Functions

Typically, a policy or project evaluation sets out the options, assesses the positive and negative impacts on individuals and determines which option gives the maximum surplus of gains over losses. Here we set out a more formal evaluation procedure, based on the notion of maximising social welfare. While this process is somewhat abstract, it provides a foundation for other more practical evaluation methods.

In this more formal process, the set of options is described as an **opportunity set**. An opportunity set shows the feasible choices available. The concave curves in Figure 7.1 show three opportunity sets. These are opportunity sets for quantities of hospital and housing services, for total income represented by GDP and equality of income (which may be measured, for example, by the Gini coefficient),[1] and for the welfare (utility) of two persons, Amy and Ben. The slope of the opportunity set at any point shows the marginal trade-off available in each case. In practice policy makers are confronted with the kind of choices shown in panels (a) and (b) rather than with a utility map as in panel (c). However, because the end objective of policy making is generally to maximise the welfare of individuals rather than health care or houses or income per se, much of the fundamental analysis of welfare economics is conducted in terms of individual utilities.

To choose the preferred point on any opportunity set, a ranking criterion is needed. This criterion must have two attributes. First, it must be expressed in the same units as the opportunity set. For example, if the opportunity set is composed of various levels and distributions of income, the ranking criterion must reflect these variables. Second, the ranking should reflect social welfare. A higher ranking should indicate greater social welfare and a socially preferred outcome.

[1] The Gini coefficient is described in Chapter 20. A coefficient of 0 is perfect equality and a score of 1 is complete inequality.

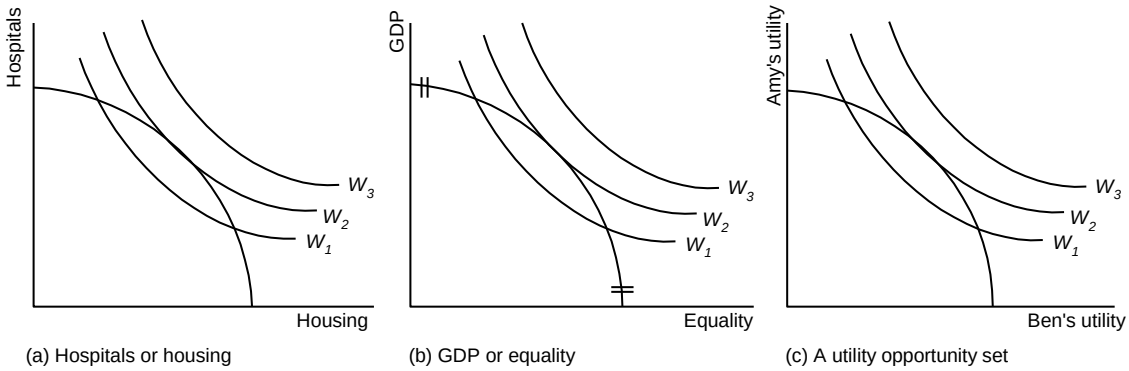

(a) Hospitals or housing (b) GDP or equality (c) A utility opportunity set

Figure 7.1 Opportunity sets and social welfare functions

In economics, the ranking criterion is usually called a social welfare function. A **social welfare function** (SWF) is any function that defines social welfare by its components. These components are typically individual utilities, but these utilities may be represented by economic or social variables that drive utility. An SWF provides a means to rank social states: the more social welfare the better. Ideally the SWF would rank all possible social states.

In this chapter we focus mainly on utilitarian SWFs with the general form:

$$W = f(u_i) = f(u_1, u_2, \ldots u_n) \tag{7.1}$$

where W is social welfare, u is individual utility (or welfare) and there are $i = 1 \ldots n$ individuals in society. In a utilitarian SWF, social welfare is a function of the utility of the members of society.[2] We focus on utilitarian SWFs because the fundamental aim of policy is the collective welfare of individuals. However, for practical purposes we will also need to consider other ways to rank policies.

Formally, the aim is to maximise social welfare (W) subject to the constraints of the opportunity set. In each panel of Figure 7.1, there are three convex iso-welfare curves (W_1, W_2, W_3), with W_3 representing the highest level of welfare. An iso-welfare curve shows points of equal social welfare (between which society is indifferent). The slope of the iso-welfare curve at any point shows the marginal rate at which society is willing to trade off the components that make up welfare. Thus, in Figure 7.1c, the slope shows the marginal rate at which society is willing to substitute Amy's utility for Ben's utility (the social marginal rate of substitution, MRS). Welfare is maximised when the iso-welfare curve (W_2) is at a tangent to the opportunity set. At this point, the slopes of the iso-welfare curve and the opportunity set are equal. The social MRS of Amy's utility for Ben's is equal to the marginal rate at which Ben's utility can be transformed into Amy's on the utility possibilities frontier (UPF).

In Figure 7.1b, the opportunity set is various combinations of GDP and equality, and there are again three iso-welfare curves. Welfare is maximised (again on the W_2 curve) when the marginal rate at which society is willing to give up income in return for more equality is equal to the marginal rate at which income has to be sacrificed to achieve extra equality.

Utility and Income

A major problem for policy evaluation is that we want to improve the welfare (utility) of individuals but our main measure of utility is income. Therefore, we need to understand the relationship between utility and income, which we introduced in the last chapter (Figure 6.8).

[2] This form of social welfare function is often called a Bergson–Samuelson SWF.

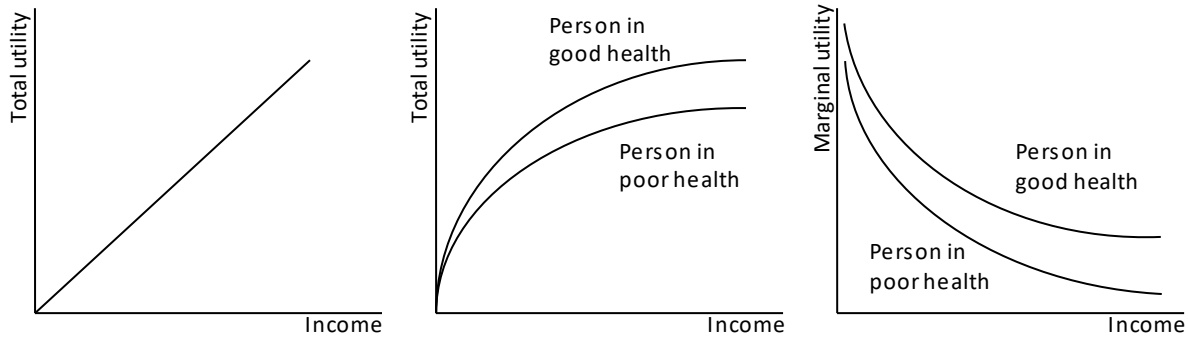

(a) Linear utility function with constant marginal utility

(b) Total utility of income

(c) Marginal utility of income

Figure 7.2 Utility and income

Figure 7.2a shows a linear utility function where the marginal utility of income is constant. If everyone were to have the same linear utility function, there would be a simple translation of income into utility. A dollar increase in income would have the same marginal value to all individuals regardless of who gained or lost the marginal dollar.

However, there are significant problems with this simple equation between income and utility. First, for most people income has decreasing marginal utility. That is, each extra dollar has a positive but decreasing value of utility. This is shown in Figures 7.2b and 7.2c. It follows that, if everyone has a similar utility function, a transfer of a dollar from a rich person to a poor one would always raise collective utility (assuming no loss of output due to the redistribution). If the redistribution involves an economic loss, we would also need to know the shape of the utility function to determine the optimal rate of dollar transfer between any two individuals.

Second, utility does not depend only on income. It may reflect health, family circumstances or a myriad of factors. Taking only differences in health, for any given level of income and personality, someone in good health is likely to have a higher level of utility than someone with poor health. Also, as shown in Figure 7.2c, an individual in good health may gain more marginal utility from additional income than does a person in poor health. Differences in utility functions complicate policy making. In this case, aggregate utility could be increased by transferring a dollar from the sick person to the healthy one. But this may not reflect social preference which could support more support for the sick person who has a lower absolute level of utility as well as lower marginal utility.

More generally, social valuations of relative need may differ from relative individual utility. Society may take the view that the happier that individual A is relative to B, the less socially valuable is an addition to A's happiness compared with an addition to B's happiness. Suppose that A and B have similar utility functions but that A has more income than B and gets x per cent of B's marginal utility from an extra dollar (where x is less than 100). Any policy that increases B's income by one dollar and reduces A's income by less than $1 / (1 - x/100)$ dollars increases their joint utility. But society may judge that more, or less, weight should be placed on the relative value of a dollar going to B (the less well-off person) than the value associated with their relative marginal utilities. Thus, society may consider that the lower the initial *level* of happiness, the more value should be given to additional units of happiness.

Measurement and interpersonal comparisons

Given that individual utility is a critical component of most SWFs, ideally utility functions would be measurable and comparable. These two requirements are described as a need for cardinality and interpersonal comparability respectively.

Of course, individual utility cannot be measured in the same scientific way that temperature or blood pressure can be measured. The alleged lack of both strict cardinality and a scientific basis for making interpersonal comparisons of utility has led some economists (most famously Robbins, 1938) to conclude that viewing social welfare as an aggregate function of individual utilities is unscientific and not useful. This warning is important, but the conclusion seems exaggerated. As Little (1957) pointed out, people commonly compare the happiness and quality of life of different individuals and conclude that one person is happier or has a higher quality of life than another, and we base our actions on these views. The fact that such comparisons cannot be precise does not render them meaningless or useless.

More recently, economists have undertaken several studies attempting to relate happiness (utility) to income (see Layard, 2006; Layard *et al*., 2008). To estimate the utility function it is typically assumed that the elasticity of marginal utility with respect to income is constant. This elasticity (ρ) is the percentage change in marginal utility associated with a one per cent change in income. This assumption implies that utility (u) is given by:

$$u = \begin{cases} \dfrac{y^{1-\rho}}{1-\rho} & \rho \neq 1 \\ \log(y) & \rho = 1 \end{cases} \qquad (7.2)$$

where y is income.[3] It follows that the ratio of the marginal utilities of two people is given by:

$$\frac{\partial u^B / \partial y}{\partial u^A / \partial y} = \left(\frac{y^A}{y^B} \right)^{\rho} \qquad (7.3)$$

If A has twice the income of B and p = 1, then A obtains half the marginal utility that B obtains from each extra dollar of income. Drawing on four large cross-sectional surveys of subjective happiness and two panel surveys, with data from over 50 countries and time periods from 1972 to 2005, Layard *et al.* (2008) estimated that the value of p varied from 1.19 to 1.36, with a best mean estimate of 1.26.[4] This implies that as income rises by 1 per cent, the marginal utility of income falls by 1.26 per cent. The authors found similar elasticities for sub-groups of the population.

These findings suggest that interpersonal comparisons of utility can be meaningful and that social welfare can be viewed as a function of individual utilities. However, this does not obviate the need for making critical ethical decisions about the nature of social welfare. It is still necessary to determine how changes to relative utilities should be weighted in an SWF, which is a normative decision.

---

[3] Equation 7.2 is known as the Constant Relative Risk Aversion function. Other common and simpler utility function specifications include $U = \log(y)$ and $U = y^p$, where $0 < \rho < 1$.

[4] Layard (2008) used a slightly modified version of the CRRA function to estimate a more realistic model whereby people need a minimum amount of income before they start receiving positive utility.

Utilitarian Social Welfare Functions

Utilitarian SWFs generally have four features. First, they depend only on the utility of individuals. Such a SWF is sometimes described as "welfarist". Rights, such as the right to freedom of speech or worship, matter only in so far as they affect individual utilities. Second, social welfare is assumed to increase with each individual's utility level, other things being equal. Third, iso-welfare curves are generally assumed to be strictly convex. This reflects the egalitarian ethic that inequality in utilities among individuals is socially undesirable. Fourth, it is generally assumed that everyone counts equally—it does not matter who enjoys a high or low level of utility. This is the principle of anonymity or symmetry.

In the following discussion, we identify various forms of utilitarian SWFs and a general SWF function that embraces all forms, depending on the parameter values applied.[5]

The additive utilitarian social welfare function

The most commonly discussed utilitarian SWF is the simple additive utilitarian SWF. The idea underlying this SWF is that society should aim to maximise the total happiness of the community.[6] In this formulation, social welfare is the unweighted sum of the utilities of all individuals:

$$W = u_1 + u_2 + \ldots u_n = \sum u_i \tag{7.4}$$

Note that individual utility is usually assumed to be a function not only of income (or goods consumed, x) but also of leisure (l):

$$u_i = f(y_i, l_i) \text{ or } = f(x_i, l_i) \tag{7.5}$$

Also, in most policy or project evaluations, goods include market and non-market goods, such as health and environmental goods.

Figure 7.3a shows the implied social indifference (iso-welfare) curves. With an unweighted additive utilitarian SWF, iso-welfare curves are linear with a 45° angle to each axis. A marginal (or unit) increase in an individual's utility always has the same social value,

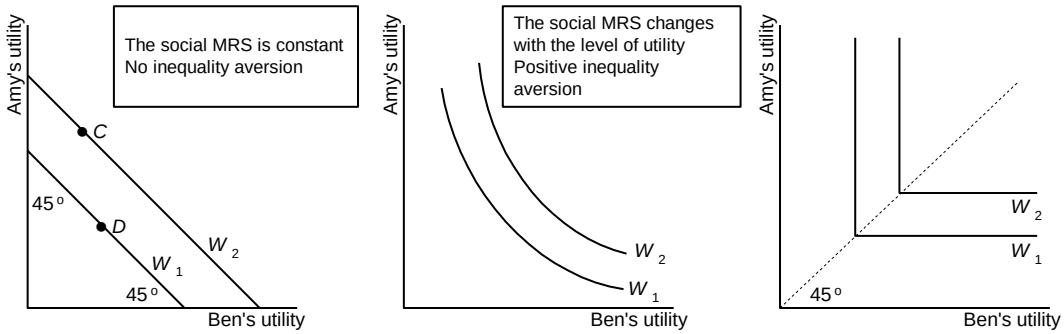

(a) Additive utilitarian SWF (b) Weighted additive utilities SWF (c) Rawlsian SWF

Figure 7.3 Social indifference (iso-welfare) curves

[5] In Chapter 9 we discuss another way of ranking social states based on individual preferences over states, drawing typically on voting methods, instead of on individual utilities.

[6] Traditionally, the additive utilitarian SWF was described as 'the utilitarian SWF'. It is also called the classical or Benthamite SWF, named after the 19th century school of utilitarian philosophers founded by Jeremy Bentham. Bentham argued that 'the greatest happiness of the greatest number is the foundation of morals and legislation'.

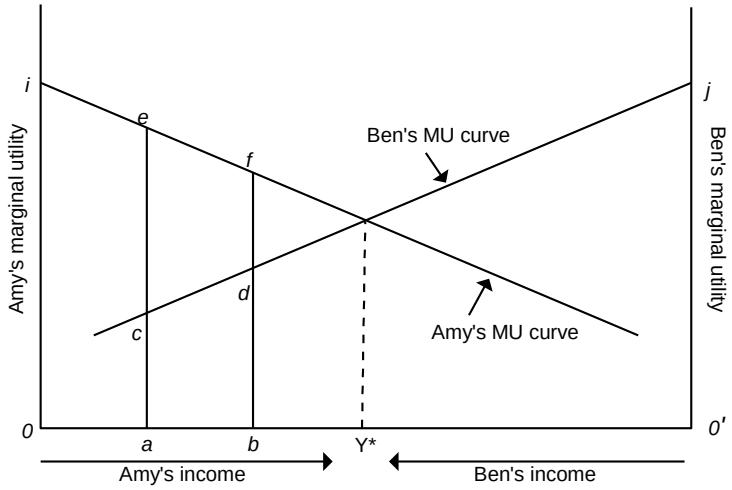

Figure 7.4 How equality of income can maximise social welfare

regardless of the level of their utility. However, this iso-welfare curve does not imply that a dollar to each person has the same social value. If Amy and Ben have a similar decreasing marginal utility of income function, a marginal dollar would provide more utility to whoever has less income. In *income* space, the iso-welfare curves would be convex.

The unweighted additive utilitarian SWF has some attractive features. It is easy to understand. It is democratic in that a marginal increase in the utility of each individual has equal weight. Also, it is often viewed as fair because it justifies redistributing income from high-income to low-income individuals (who usually have greater needs and a higher marginal utility of income). Indeed, under certain assumptions (as discussed below) it justifies complete equality of income distribution.

This is shown graphically in Figure 7.4. Here there is a fixed amount of income (OO') to be divided between Amy and Ben, who are assumed to have similar utility functions. Initially, say Amy has $O'a$ income and Ben's income is $O'a$. Then according to the utilitarianism calculus, total welfare is $Oiea + O'jca$. If ab income is transferred from Ben (the richer) to Amy (the poorer), total welfare will be $Oifb + O'jdb$, which increases welfare in the society by $cefd$. When Amy has OY^* amount of income and Ben has $O'Y^*$ income, their marginal utilities are equal and their total utility is maximised. Because Amy and Ben have the same utility function, total utility is maximised when total income is divided equally between them.

Nevertheless, the unweighted additive utilitarian SWF is often criticised for being concerned only with total welfare and not with its distribution. In Figure 7.3a, C is on a higher social welfare curve than D and, with an additive utilitarian SWF, would be preferred to D. But many people would disagree and prefer D to C. They may argue that income should be transferred to individuals with lower absolute levels of utility, without regard to the marginal utility that they obtain from income or to the possible loss of aggregate utility. With a simple additive utilitarian function, a transfer of income from a poor unhealthy individual who has little enjoyment in life to a rich healthy person could increase total utility. But this may not be regarded as desirable. An example of such a scenario is given in Box 7.1 overleaf.

Introducing equity explicitly into social welfare functions

Equity can be introduced explicitly into a SWF in various ways, for example by a weighted additive function. A multiplicative form of SWF also allows for the distribution of utility. In

Box 7.1 An additive utilitarian social welfare function may not be desirable

Say in a two-person (Ben and Amy) economy, total fixed income of \$100 is split between them. For Ben, the marginal utility of income is: $MU_B = 400 - 2y_B$ while for Amy the marginal utility is $MU_A = 400 - 6y_A$, where y_B and y_A are amounts of income to Ben and Amy respectively. If the welfare function is simple additive, to maximise W set marginal utilities equal subject to the constraint that $y_B + y_A = 100$.

So, $400 - 2y_B = 400 - 6y_A$.

Substituting $y_A = 100 - y_B$

$2y_B = 6(100 - y_B)$.

Therefore, $y_B = 75$, $y_A = 25$.

This shows that if, for the same amount of income, Ben gets more marginal utility than Amy, the welfare-maximising utilitarian solution is to provide Ben with more income. The additive utilitarian SWF, which aims to maximise the sum of individual utilities, may not be an attractive ethical objective.

such a SWF, welfare is a function of the product of individual utilities. This can be expressed in logarithmic form with social welfare a function of the sum of the natural log of individual utilities.

More generally, we may adopt a weighted multiplicative utilitarian SWF:

$$W = u_1^{a_1} \times u_2^{a_2} \times \ldots \times u_n^{a_n} = \prod u_i^{a_i} \tag{7.6}$$

where the a_i are weights between 0 and 1 assigned to each individual's utility and Π represents multiplication. Equation 7.6 allows us to give greater weight to increases in the utility of people with greater needs.

Equation 7.6 is represented by a convex iso-welfare curve (see Figure 7.3b). Convexity implies that society prefers an even distribution of individual utilities to an uneven one. As inequality increases, society is willing to accept a decrease in the utility of the poor only if there is a larger increase in the utility of the rich. The acceptable trade-off depends on the degree of inequality.

Another SWF that allows even more strongly for welfare distribution is the Rawlsian maximin SWF. Rawls (1971) proposed that social welfare should depend only on the welfare of the least well-off individual (or group of individuals) in society. This can be interpreted as:

$$W = \min (u_1, u_2, \ldots, u_n) \tag{7.7}$$

This implies an L-shaped iso-welfare curve (Figure 7.3c). An increase in the utility of anyone other than the least well-off person makes no difference to the welfare of society.

Rawls argued that people could choose a fair society only if they were unaware of their position, income-earning capacity and so on in society (sometimes described as a veil of ignorance). He held that individuals who did not know their position, being risk averse, would choose a society in which the welfare of the least well-off person was as high as possible. This approach attempts to generate an unbiased ethical ranking of social states. However, it has been questioned on the ground that the selfish views of individuals in an *original state*[7] do not necessarily have ethical content. Nor is there any evidence that individuals so placed would be so risk averse as to be concerned only with the welfare of the least well- off person.[8] The exclusive emphasis on the welfare of the least well off individual is limiting. The Rawlsian SWF is indifferent to the welfare of anyone else.

[7] The position was original because no social or political system existed, and therefore no individual in the imagined position had a known place in the society. Rawls described this as the principle of justice.

[8] Harsanyi (1955) argued that if individuals in their original position were to maximise their expected utility, they would choose an additive utilitarian SWF.

A general social welfare function

All the SWFs described above can be viewed as special cases of a general SWF.

$$W = \frac{1}{1-\varepsilon}\sum_i \left(u^i\right)^{1-\varepsilon} \tag{7.8}$$

where ε is a parameter that reflects social concern for equality of welfare.[9] With this formulation, the elasticity of substitution along each iso-welfare curve is constant and given by 1/ε. This means that the ratio of the proportional change in marginal utilities to the proportional change in absolute utilities is constant.

When $\varepsilon = 0$, there is no explicit concern for equality. Each marginal unit of utility has equal value regardless of the individual's level of utility. Equation 7.8 then reduces to the simple unweighted additive utilitarian SWF (Equation 7.4) and the iso-welfare curve in Figure 7.3a. When ε is positive, increases in individual utility are transformed into less than proportional increases of

$$\frac{1}{1-\varepsilon}\left(u^i\right)^{1-\varepsilon}$$

This implies that more weight is attached to a marginal increase in utility for someone with low utility than for someone with high utility. As ε increases, the weight for equality increases. In the limit $\varepsilon \to \infty$ and Equation 7.8 reduces to the Rawlsian SWF, Equation 7.7.

Table 7.1 shows how the social valuation of marginal utility may vary with differences in utility levels and social values of ε. If Ben's level of utility is M times that of Amy's, the weight attached to an additional unit of utility for Amy is M^{ε}. Suppose that $M = 2$ (Amy has half the absolute level of utility of Ben). If $\varepsilon = 0$, a unit increase in Amy's utility has the same social value as a unit increase in utility for Ben. If $\varepsilon = 2$, the social value of marginal unit of utility for Amy is four times the social value of a marginal unit of utility for Ben.

Note that these **social** valuations of utility are *additional* to any differences in the marginal private utilities of consumption. Suppose that a dollar provides Amy with one extra unit of utility and Ben with only 0.5 extra units of utility. And suppose that $M = 2$ and $\varepsilon = 2$. The net social value of a transfer of a dollar from Ben to Amy is then 3.5 units, because Amy's one unit attracts a social value of 4 units, whereas Ben forgoes 0.5 units of utility.

From utilitarian social welfare functions to public policy

SWFs are intended to assist in ranking social states and thus in determining public policies. However, utilitarian SWFs are a rather abstract concept with few direct practical applications.

Table 7.1 Valuations of marginal utility of individuals with differing social weights

Values of social weights (ε)	*Individual utility levels relative to average utility*				
	2/1	*1/1*	*½*	*1/3*	*1/4*
0	1	1	1	1	1
0.5	0.71	1	1.41	1.73	2
1	0.5	1	2	3	4
2	0.25	1	4	9	16
4	0.0625	1	16	81	256

[9] When $\varepsilon = 1$, Equation 7.8 is indeterminate and is replaced by $W = \sum_i \log(u^i)$. It may also be noted that this SWF is ordinal and that the value of W may be negative, but it rises with increases in any individual utility.

One utilitarian SWF with a possible practical implication is the Rawlsian SWF. This implies that income should be redistributed to the least well-off person or group of persons in society. Of course, individual welfare must be defined, for example in terms of income, health, dependants and so on.

The simple additive SWF was also regarded traditionally as implying a strong case for income redistribution. Indeed, under the following assumptions social welfare is maximised by complete equality of income:

1. Individuals have similar utility functions (they gain equal utility from any given amount of income).
2. Marginal utility of income falls as income increases.
3. The total amount of income available is fixed and independent of its distribution.

If these assumptions hold, any transfer of income from a richer person to a poorer one increases the welfare of society, because a marginal dollar gives the poorer person more utility. This was shown in Figure 7.4. Moreover, the larger the transfer the greater is the increase in social welfare.

A related policy question is: how should government raise a given level of tax revenue with least welfare cost? If the marginal value of a dollar falls with income, a dollar tax on a high-income person results in less loss of total utility than does a dollar tax on a low-income person. To achieve a tax revenue target with least loss of welfare, tax should be levied at a 100 per cent rate on the highest incomes downward until the target is reached. For the marginal dollars of tax paid, all those paying tax would have an equal sacrifice. But this marginal sacrifice is lower than the sacrifice that would occur with a tax on lower incomes. People with low incomes would pay no tax.

However, these conclusions depend on the assumptions. If the assumptions do not hold, neither do the conclusions. Assumption (1) is questionable because utility depends on many factors besides income. Assumption (2) is more plausible. However, assumption (3) is very doubtful because income redistribution almost always reduces the amount of income available for distribution. Therefore, the conclusions that total utility is maximised by complete equality of income are questionable.

These examples show that maximising a utilitarian SWF subject to constraints may produce significant policy guidelines, but not operational outcomes. We have also seen that it is possible to estimate utility functions which allow us to convert income into utility equivalents. However, for practical purposes, it is generally necessary to work directly with economic variables such as income or consumption along with leisure. For example, macroeconomic states may be ranked as a function of GDP and some measure of income equality. To assess particular policies or projects, money measures of benefits and costs are generally required. We review these approaches below and their implications for evaluating social welfare.

Utilitarian social welfare functions: other considerations

Other important SWF issues are the components of utility, the consequentialist nature of SWFs and the related treatment of equity.

The components of welfare. In the eloquent words of the American Declaration of Independence (1776), all men "are endowed … with certain inalienable rights … among these are life, liberty and the pursuit of happiness". Whether or not one agrees with the notion of natural rights, most people would consider that their utility depends not only on income and leisure but also on their rights to vote and elect governments, to freedom of speech and congregation and freedom from arbitrary arrest. These rights are hard to measure.

Outcomes and process. SWFs are often treated as consequentialist: welfare depends on outcome, not on process. In the consequentialist view of social welfare, government should distribute output according to need and so as to maximise the total welfare of society, subject to the constraint that redistribution may reduce the output available for distribution. In effect, all output belongs to the state. No individual has a right to their output. .

However, taking process into account, there is another equity principle—the principle of just reward. According to this principle, people should be compensated for both labour and savings. Compensation for labour can be represented formally by including leisure in the utility function as in Equation 7.5. Other things being equal, individuals with more leisure have more utility. Those who work longer hours may be compensated by higher money income.

But what is a fair return to labour? Nozick (1974) argues that individuals have a right to all legally acquired holdings, acquired either through earnings or through inheritance of wealth providing this was justly acquired. This is an extreme view. A more moderate view is that individuals are entitled to keep what they can earn in a competitive market, given equal positions of opportunity and ability. Individuals should not be able to appropriate surplus income, above their opportunity cost, when this is due to market power in imperfectly competitive markets. Nor should they be able to appropriate economic rent due to superior innate capacities. This rent should be shared with less gifted individuals.

Some writers, such as Holcombe (1998), argue that fair process is as important as fair shares of output. If the process is fair (e.g. there is equality of opportunity and markets are competitive) the economic outcome may be viewed as fair. If a football game is fair, the outcome is generally viewed as fair. However, given differences in individual abilities and in resources available to individuals, and the random nature of misfortune, society can scarcely be indifferent to variations in economic outcomes even in competitive markets.

Conclusions. Most utilitarian SWFs imply that resources should be shifted from individuals with high utility to those with less utility. But utility (or need) are not easily defined concepts. On the other hand, the just reward approach argues that individuals should be entitled to their earnings provided there is equality of opportunity and markets are competitive. Such issues arise at many points in the text. For example, in our analysis of public finance below, we will discuss ability to pay versus user benefit principles of equity and the meaning of vertical and horizontal equity. In our discussion of social assistance, we will meet a variety of concepts of need especially in relation to differences between households.

Working with Income

For most practical purposes, economists work with incomes or willingness-to-pay dollar amounts rather than with utilities. At the macroeconomic level, we may trade-off total income (GDP) against equality of distribution. Thus a SWF may be represented as:

$$W = W(S, \theta) \tag{7.9}$$

where S is represents total income, which captures the efficiency aspect, and θ represents the inequality of income, which captures the equity aspect. This SWF is increasing in S and decreasing with respect to θ. We have seen previously that a utilitarian SWF is generally increasing with individual utility and that individual utility increases with income. To satisfy this feature the above SWF has to satisfy the following principle:

$$\frac{\partial W}{\partial S}\frac{\partial S}{\partial y_i} > \left|\frac{\partial W}{\partial \theta}\frac{\partial \theta}{\partial y_i}\right| \tag{7.10}$$

An increase in someone's income affects the SWF in two ways: it increases the total income of the society and changes the inequality. The above condition states that an increase in someone's income has a greater effect on total welfare than does an increase/decrease in inequality. Suppose that a policy change increases the income of one person (or group) in the society. This increases S, but this may increase or decrease inequality depending on whose income has increased. If inequality increases its effect on SWF is negative. However, with the above SWF, whoever may be the beneficiary of a policy change, social welfare will increase. The effect of an increase in income on S always exceeds the effect of this increase in income on θ. From an axiomatic perspective, the above SWF can be derived as:

$$W = \text{GDP per capita} \times (1 - \text{GC}) \tag{7.11}$$

where GC is the Gini coefficient measure of income inequality.

However, the core components of social welfare are *individual* utilities and where possible economists estimate the benefits and costs of economic changes to individuals. As economists have stressed for a long time and as Stiglitz *et al* (2009) discussed in a major OECD report, GDP is not a good measure of welfare even when adjusted for distributional effects because it does not include any value for non-marketed goods, health, externalities, consumer surpluses or leisure. A full accounting of welfare effects for individuals includes values for all these elements of welfare drawing on the valuation principles discussed in Chapter 6.

Weighting incomes. Turning to individuals, the core problem is the relationship, or the lack of it, between income and utility. Even if individual utility can estimated as a function of income, social valuations of marginal utility may differ from individual valuations.

Suppose that the SWF embodies the idea that the social value of an additional unit of income is some function of an individual's income. In essence, we can replicate the welfare function shown in Equation 7.8 but substitute income (y) for utility and some form of social weighting factor (w) for equity. Thus, suppose that social welfare can be re-expressed as a function of incomes:

$$W(u_1, u_2 \ldots, u_n) = \sum_i (w_i y_i) = \sum_i y_i^{(1-w)} / (1-w) \tag{7.12}$$

where w represents the elasticity of the social valuation of the marginal utility of income for each individual (or the degree of aversion to income inequality).

Table 7.2 provides summary values of the *social* value for an additional dollar for three individuals depending on the choice of social weight. This is of course simply *a reduced version of Table 7.1* with income differences instead of utility differences. If $w = 0$, the social value of extra income is the same for everyone. Society would be indifferent to distributional changes in income. If $w = 1$, the social value of marginal income is inversely proportional to income. If Ben has twice the income of Amy, the social value of an extra dollar to Amy is twice the social value of an extra dollar to Ben. The higher the value of w, the higher is the social value of marginal income of low-income individuals. In the extreme case, if $w = \infty$,

Table 7.2 Values of marginal income as function of individual income and social weights

Social weight (w)	*Individual income relative to mean income*		
	2/1	*1/1*	*0.5/1*
0.0	1.00	1.0	1.00
0.5	0.71	1.0	1.41
1.0	0.50	1.0	2.00
2.0	0.25	1.0	4.00

all weight would be attached to the income of the least advantaged person. The weighting approach is attractive in that it recognises welfare differences between individuals and attempts to provide an operational procedure for dealing with these differences. Policy makers must make such judgements in many situations.

However, determination of any such weights requires value judgements. While such valuations might be elicited from social surveys and econometric analysis, as in Layard et al. (2008), weighting dollars according to the estimated average marginal utility of income (w = 1.26) would itself represent a value judgement. Therefore, most economists are cautious about employing such weights and treat all dollar values equally in an evaluation procedure like cost-benefit analysis. Of course, this approach does imply a social weight (i.e. $w = 0$). However, this is essentially an attempt to separate efficiency and equity considerations. It does *not* imply that equity considerations are irrelevant to public policy. Rather, as discussed below, they need to be considered alongside expected efficiency outcomes.

Social Welfare and Cost-Benefit Analysis

In practice, decision makers deal mainly with specific policies and projects rather than with overall design of the economy. Typically, we want to know whether a policy or project will increase social welfare. To answer this question, we generally focus on the change in social welfare (ΔW). Adopting utility units, we want to know whether:

$$\Delta W = \sum_i w_i \Delta u_i > 0 \tag{7.13}$$

where w_i are social weights where applicable.

However, given the difficulty of estimating u_i and w_i, as a practical procedure, economists generally estimate whether the benefits exceed the costs by drawing on unweighted income-based valuation measures described in Chapter 6. Therefore, we estimate whether:

$$\Delta W = \sum_i (b_i - c_i) = \sum y_i > 0 \tag{7.14}$$

where b_i and c_i are estimated benefits and costs for individuals (or other economic agencies) and y_i are income-based measures of changes in individual welfare, which may be positive or negative. If the estimated total benefit exceeds total cost, there is said to be a net social benefit (NSB > 0). Not surprisingly, this approach is called cost–benefit analysis (CBA).

The monetary valuations of goods usually reflect changes in consumer and producer surpluses based on ordinary demand schedules, but they may reflect estimates of compensating or equivalent variations when appropriate. As we saw in Chapter 6, the differences between these valuation measures are generally small. When they are not, the choice of measure can have significant implications. However, the aggregation of individual values into an overall measure of social welfare is generally more contentious than the valuations of goods to individuals or other economic agents. This takes us back to the issue of how to treat equity if not by weighting valuations.

Aggregating income-based measures of welfare changes

Evidently when a project or policy produces winners and losers and when a dollar has a different value to different individuals, NSB = $\Sigma y_i > 0$ may not represent a net social welfare improvement if the distributional effects are regressive.

However, $\Sigma y_I > 0$ does represent a *potential* net social welfare improvement (it is a potential Pareto improvement). If estimated benefits exceed costs, the gainers from a change can fully compensate the losers. If compensation is made, someone is better off from the change and no one is worse off. The resource reallocation would then be an actual Pareto improvement. This approach is formalised in three related compensation criteria:

1. **Kaldor criterion:** a change from *A* to *B* is desirable if those who gain from *B* could compensate the losers and still be better off at *B* than at *A*.
2. **Hicks criterion:** a change from *A* to *B* is desirable if the losers from *B* cannot bribe the gainers into not wanting the change to *B*.
3. **Scitovsky criterion** (sometimes described as the double criterion): a change from *A* to *B* is desirable if those who gain from *B* could bribe the losers to accept the change and the losers could not bribe the gainers into not making the change.[10]

There is relationship between these compensation criteria and the concepts of compensating and equivalent variation (CV and EV) described in Chapter 6. The Kaldor test requires that the sum of CVs is positive. The Hicks test requires that the sum of EVs is positive.

Scitovsky (1941) pointed out that the Kaldor and Hicks criteria could create inconsistent outcomes and that strictly both criteria would need to be met. The Scitovsky reversal paradox is illustrated in Figure 7.5. Suppose that the initial utility point is *A*. Government then decides to build a new road that provides benefits only to *Y*, but for which both *X* and *Y* pay. The new utility levels are given by point *B*. However, *Y* can compensate *X* and both can be better off at point *C*. But if compensation is *not* paid, the starting point on the utility map with the road is point *B* and the comparator (without the new road) is point *A*. Given that *X* prefers non-road goods, she is better off at *A* with no road. She is also willing to bribe *Y* into having other goods and not having the road at point *D*, which both prefer to point *B*. Therefore, when compensation is not made and the distribution of income changes, the Kaldor and Hicks criteria can justify both a change and the reverse decision. Fortunately, this inconsistency is rare and for most purposes the Kaldor criterion is adopted.

Kaldor and Hicks designed their criteria in response to the Robbins (1938) critique that individual utilities cannot be compared and to avoid the need for a social welfare function. If losers are compensated, changes that satisfy their criteria are Pareto efficient. There would be no need for interpersonal comparisons of utility or indeed for any SWF. Moreover, such an approach may be justified pragmatically. Because governments undertake many projects, adopting projects with a positive NSB (or where there are alternatives, options with highest NSB) will raise the welfare of most people more than would projects with lower NSBs,

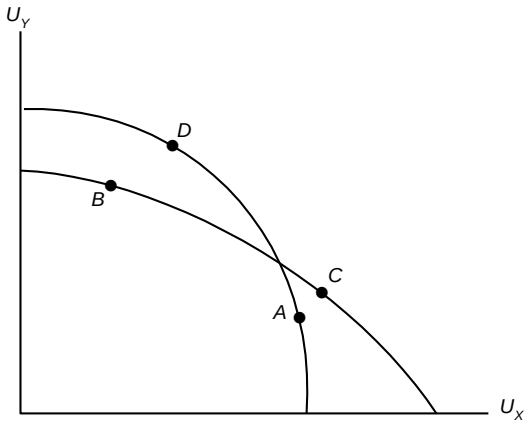

Figure 7.5 The Scitovsky reversal paradox

[10] See Kaldor (1939), Hicks (1939) and Scitovsky (1941).

including projects with negative net benefits. However, the compensation principle is based on the notion of hypothetical compensation. It does not require that compensation is paid.

If compensation is not paid and someone loses from a policy change, the argument that compensation could be paid is of questionable relevance. The NSB criterion (Equation 7.14) gives equal weight to all dollar valuations of benefits and costs regardless of their distribution. Of course, some projects are redistributive from rich to poor as well as efficient. However, if high income individuals gain from a project and low-income people lose, a project may satisfy the net social benefit criterion but worsen the distribution of income.

Evaluation options

There are three main ways to evaluate policies or projects. All are based on estimates of individual valuations of benefits and costs. The differences relate to the treatment of aggregate net benefits.

First, we may evaluate policies with a social welfare function that applies social (equity) weights to WTP values for different groups of individuals. This approach aims to fully integrate efficiency and equity outcomes. Equity-weighted measures of NSB are used occasionally, for example in evaluations of economy-wide tax and income transfer systems (see optimal taxation literature in Chapter 28). However, they are used here mainly for academic demonstration purposes because there is no objective basis for establishing equity weights for groups of individuals. This approach is rarely used to evaluate specific policies or projects because establishing equity weights for individuals for individual cases is even more problematic. Moreover, an estimated equity-weighted NSB is an unclear mixture of efficiency and equity impacts. It does not have a clear welfare meaning. A policy could show a positive NSB even if the distributional impacts were regressive.

Second, policies may be evaluated simply with an efficiency criterion such as unweighted NSB. Government would adopt policies that pass the NSB test, because this makes most efficient use of resources. Government would deal with equity separately as needed through income transfers or other compensation programs. This could include indirect compensation whereby government would compensate losers via fiscal policies to ensure that overall changes do not disadvantage less well-off individuals. This approach may achieve high efficiency outcomes and reasonable equity at low cost because fiscal instruments are generally better instruments for redistribution than are projects. However, acceptance of it requires faith in the political process. When a policy disadvantages some individuals, they may not take the sanguine view that they will receive appropriate compensation via other policy changes.

The third approach to evaluation accepts that equity may be important at the individual policy or project level. However, it adopts a dual track consideration of efficiency and equity rather than attempt to integrate efficiency and equity into a single overall estimate of social welfare (as in the first approach). The economist would present decision makers with both (1) the estimated aggregate (unweighted) NSB of a policy and (2) a description of the distributional impacts, possibly including potential compensation packages. This is a common approach.

The estimated NSB, based on cost–benefit analysis, indicates the potential efficiency gains from a policy. But the economist generally does not attempt to convert this into an aggregate measure of social welfare. Rather the economist provides separate information on the efficiency and equity consequences of any policy. Government then makes the policy determination.

Conclusions

Welfare economics provides a rigorous framework for determining the allocation of resources and related public policy decisions. Valuations of individual benefit and cost are based on individual preferences. The principles of compensating and equivalent variation, or of consumer or producer surpluses, provide exact monetary measures of benefit and cost. Thus, welfare economics provides a means to quantify costs and benefits.

However, technical analysis cannot resolve all issues. A dollar may be more valuable to one person than to another. In any case, social valuations of individual needs may differ from private valuations of the marginal utility of income. Thus, it is often necessary to make trade-offs when someone gains from a policy and another person loses. Unless full compensation is possible, policy making requires a trade-off between the efficiency gain of a policy change and the distribution of the gains and losses.

It is sometimes feasible to treat efficiency and equity separately. At the macroeconomic level, policy makers would aim first to maximise the total value of output by the most efficient use of resources and then redistribute the income earned in an equitable way. At the microeconomic level, government would adopt efficient policies that provide the highest net social benefit (those policies with greatest excess of total benefit over cost) and deal with distributional issues separately, where appropriate by compensating losers. Of course, a social welfare function, reflecting values judgements, is still required (at least implicitly) to determine the optimal redistribution.

A rationale for treating efficiency and equity as separate issues is that efficiency of resource use is a technical issue and distribution of output a normative matter. Also, when policies have minor equity implications, the separation of efficiency and equity may be practical and reasonable. However, in the absence of individualised lump sum transfers, it is impossible completely to separate equity from efficiency—one will affect the other. It is then impossible to judge between alternative economic states without an ethical social welfare function or a political judgement.

Summary

- The fundamental problem of public choice is to determine how to assess, compare and aggregate gains and losses from the use of resources.
- The assessment is based on valuations of individual preferences. Comparison and aggregation of the valuations raise greater problems as the marginal utility of income varies across individuals.
- Also social valuations of marginal income may vary according to the level of wellbeing of individuals in society.
- Ideally policies or projects would maximise social welfare, using some agreed measure of social welfare, that is, a social welfare function (SWF).
- The SWF describes the relationship between social and individual welfare. However, individual utilities can be combined into an overall measure of social welfare in many ways and this aggregation requires value judgements.
- Most public policy involves possible changes from the present rather than design of a whole social state. In this case we want to measure changes in the welfare of individuals (i.e. their benefits and costs) and to find a way to sum these changes.
- If the estimated total benefits of a policy exceed the costs, the policy is described as efficient in the sense that the gainers could compensate the losers. If compensation is made directly or indirectly in the longer run, someone would gain, and no one would lose, from the policy.
- If compensation is not made and a project or policy has some expected adverse effects, policy makers must decide whether the efficiency benefits justify the distributional impacts.
- There is no ready technical basis for weighting the benefits and costs into an overall measure of social welfare.

Questions

1. What is a social welfare function?
2. Is it possible to compare individual welfare levels? If such comparison is not possible, is it possible to make judgements about social welfare?
3. Does a Pareto-efficient outcome avoid the need for a social welfare function?
4. Can a utilitarian social welfare function be equitable?
5. Assume an economy with only two people (Amy and Ben).
   i. Let the social welfare function be $W_1 = u_a + u_b$ where u_a and u_b are the utilities of Amy and Ben. Graph the iso-welfare curves (the social indifference curves). What importance, if any, is attached to their relative wellbeing?
   ii. Repeat for $W_2 = 2u_a + u_b$. What are the implications?
   iii. Draw a utility possibilities curve and show how the SWFs in i and ii affect the desired outcome.
   iv. What are the main weaknesses of the additive utilitarian social welfare function?
6. Should government be concerned with the distribution of utility or the distribution of income?
7. Is it possible to estimate the marginal utility of income?
8. Explain why social valuations of marginal needs may differ from individual valuations of the marginal utilities of income. Is it possible to measure these social valuations?
9. When does cost–benefit analysis maximise a utilitarian social welfare function?
10. Is maximising an additive utilitarian social welfare function always inequitable?
11. Define the compensation principle. What is the relationship between this principle and a potential Pareto improvement? What value judgements underlie the compensation principle?
12. When may the Kaldor and Hicks compensation tests give different policy answers?
13. Amy's utility is u_A and her income is y_A while Ben's utility is u_B and income is y_B. Suppose: $u_A = 10y_A^{0.5}$ and $u_B = 10y_B + 0.8u_B$. Suppose initially that both of them have \$100. If you consider a (simple) additive SWF, what will happen to social welfare if \$36 is taken away from Ben and given to Amy?
14. How can cost–benefit analysis deal with distributional impacts of policies?

---

15. Society contains three individuals (groups) with the following income, health and perceived utility.

	1	*2*	*3*
Income (\$ per annum)	20 000	40 000	80 000
Health	Good	Poor	Good
Utility	100	120	200

The government has six policy options (A to F) which give forecast income changes as shown below.

Option	*1*	*2*	*3*	*Total*	*Rank*
A	+1000	0	-3000	-2000	
B	+500	+500	-2500	-1500	
C	+1500	-500	-500	+500	
D	+500	0	+500	+1000	
E	-100	+1100	0	+1000	
F	+1000	-250	+500	+1250	

i. How would you rank the options?
ii. What does this imply about a social welfare function?
iii. What does it imply about compensation schemes?

Further Reading

Buchanan, J.M. (1959) 'Positive economics, welfare economics, and political economy', *Journal of Law and Economics*, 2, 124–138.

Kaldor, N. (1939) 'Welfare propositions of economics and interpersonal comparisons of utility', *Economic Journal*, 49, 549–552.

Layard, R. (2006) 'Happiness and public policy: a challenge to the profession', *Economic Journal*, 116, C24–33.

Layard, R., Nickell, S. and Mayraz, G. (2008) 'The marginal utility of income', *Journal of Public Economics, Special Issue: Happiness and Public Economics* (eds) T. Besley and E. Saez, 92, Nos 8–9.

Little, I.M.D. (1957) *A Critique of Welfare Economics*, Chapters I and XV, Oxford University Press, Oxford.

Mueller, D.C. (2003*) Public Choice III*, Chapter 23, Cambridge University Press, Cambridge.

Ng, Y-K. (2003), *Welfare Economics*, Macmillan, Basingstoke, UK.

Pearce, D.W. (1983) *Cost–Benefit Analysis*, Chapters 1–3, Macmillan, London.

Stiglitz, J.E., Sen, A. and Fitoussi, J-E. (2009) *Report by the Commission on the Measurement of Economic Performance and Social Progress*, OECD, Paris.

Tresch, R.W. (2002) *Public Finance: A Normative Theory*, Chapter 4, Academic Press, New York.