

From Economic ‘Efficiency’ to Economic Welfare

I COST-BENEFIT ANALYSIS IN WELFARE ECONOMICS

The basic theory of welfare economics enables one to identify how far the market is operating ‘efficiently’. This is taken to mean how far the market allocates resources in such a way as to make the maximum contribution – given resources and technical knowledge – to society’s economic welfare. Welfare economics provides the criteria necessary for this purpose. This includes the well-known rules that any student of economics would know, such as the relationship between prices and social costs. Satisfaction of these rules is supposed to ensure that society is operating at a point where no reallocation of resources can raise one person’s utility without reducing somebody else’s utility. Such a point is known as a ‘Pareto optimum’ point, in honour of the great nineteenth-century Italian sociologist and economist, Vilfredo Pareto.

Before proceeding to examine the criteria of an efficient allocation of resources, however, it should be noted that, as the Nobel Laureate Joseph Stiglitz has put it, ‘Unemployment – the inability of the market to generate jobs for so many citizens – is the worst failure of the market, the greatest source of inefficiency. . . .’¹ Since recurrent periods of very high unemployment have characterised developed economies for over a century, it might appear that analysing the ‘efficiency’ with which resources are allocated, when so many are just not allocated anywhere at all, is a flagrant example of the old warning against swallowing a camel while straining at a gnat. Nevertheless, given the level of resource used at any

time – which is a matter of macro-economic policy – there are still innumerable problems of resource allocation to be faced. There is no point in making what may be a bad situation worse by misallocating such resources as are employed. So this book is about the allocation of resources in the market at any point of time.

In practice, of course, there is not just one market. The economy is made up of innumerable markets. There are labour markets, financial markets, second-hand markets, online markets, covered markets, commodity markets, housing markets, and so on. But do these markets all operate in a way that will maximise their potential contribution to society's economic welfare? Most people would say, quite reasonably, that one does not need any fancy theory to conclude that financial markets have been behaving disastrously over the ages. But here we are concerned with other, less 'special' markets. And it is the theory of welfare economics that shows why most markets will generally fail to make the maximum contribution to economic welfare. It is welfare economics that demonstrates that there are likely to be instances of what are known in the jargon as 'market failure' – that is, a failure of markets to make their full potential contribution to society's economic welfare.

There are various well-known forms of market failure. The most widely known include what are known as 'negative externalities': aircraft and trucks can produce a lot of noise, some factories emit poisonous emissions from their chimneys, and so on. Other goods or services may not be produced up to the socially desirable level on account of some monopolistic restriction, or a serious imbalance between the information available to buyers and sellers, or taxes or tariffs. Welfare economics identifies these 'market failures' and the criteria by which the socially 'efficient' level of such activities can be determined.

Thus the basic theory of welfare economics has much to contribute to the promotion of society's economic welfare. When any particular instance of 'market failure' is identified there is often a *prima facie* case for some policies to rectify the matter. Sometimes this may indicate a *prima facie* case for government intervention. But 'government failure' is sometimes likely to be as bad as market failure, if not worse, as when policies are introduced as a result of the influence of particular pressure groups or administrators. Some changes of policy concerning the subsidisation of offshore wind farms, or 'fracking', or investment in 'prestige projects' are modern examples.

And, as everybody will have noticed, governments play a big part in the operation of most markets. All markets are affected by government taxes

and subsidies. In the labour market many countries have legal minimum wages or restrictions on the age at which people may be employed. In financial markets, banks are usually subject to certain regulations, though these do not seem to have been very effective lately. Government subsidies to financial institutions by bailing them out when they would otherwise have collapsed may well have helped sustain the instability of financial institutions. In many countries governments are also the main agents in the market for certain goods or services, such as medical services, education, national defence, infrastructure such as transport, and what passes for law and order in our cities. Under what conditions are all such forms of government intervention justified? By what criteria can one judge when the market is not performing well?

Welfare economics helps identify these criteria. And it may also help to identify possible cures. For example, it enables one to identify what tax ought to be imposed on some polluting activity if that is to be the preferred way of ensuring that polluters take account of the full social costs of the activity in question.

2 COST-BENEFIT ANALYSIS AND 'FRANKLIN'S ALGEBRA'

In practice, what is generally known as 'cost-benefit analysis' (CBA) is an essential tool of welfare economics. For it enables one to establish whether the economic benefits of a project exceed the costs and hence whether the conditions of a Pareto-optimising move can be satisfied (leaving aside distributional considerations). In its most general form 'cost-benefit analysis' resembles what Cass Sunstein has called 'Franklin's Algebra', namely a list of all the pros and cons of any particular decision.² This decision tool, which was set out in a letter by Benjamin Franklin in 1772, is not necessarily confined to economic considerations or quantifiable inputs. But such economic considerations – including even those that are quantifiable – that enter into a comparison of the pros and cons of any choice are invariably influenced by ethical considerations in two main ways.

First, they influence what factors are included in the economic part of any more general comparisons of the pros and cons of any particular policy. Should the analysis be confined to the 'efficiency' with which resources are allocated in the economy or should account also be taken of the effect of any policy on the distribution of income or the welfare of particular groups within one's country or abroad? For example, consider

the provision of health care to poor people, or an appraisal of some trading arrangement such as the export of toxic waste or the protection of domestic agriculture, or industry. How far can beneficence towards poorer groups or countries be taken into account in the analysis?

Second, even when the scope of the CBA has been settled, ethical considerations also affect the prices attached to the relevant costs and benefits. The gap between people's revealed preferences and their welfares that has been discussed in [Chapter 5](#) clearly affects the welfare significance of the pattern of prices on the market. For if people often make choices that do not really reflect their welfares, the market prices may not always correspond to the welfare of the market participants in question. Should one adjust downwards the prices for some items that enter into a CBA if we believe that the goods and services in question are really bad for the people who buy them, or, conversely, are better for the consumers than they know so that they really ought to buy more of them? That is, if the prevailing market price does not adequately reflect the contribution made by the goods or services in question to their welfare.

Third, prices used in a CBA of some facilities can be distorted as a result of the 'commodification' of certain goods and services to which reference was made in [Chapter 5](#). Elizabeth Anderson points out that these prices are often based on surveys of how much people would be prepared to pay for the facilities in question, such as the provision of parks, or improvements in the local environment, or recreational facilities, or improvements in safety.³ But these prices will usually reflect – mainly or solely – people's desires for the goods and services for their own use, rather than their valuations of them as facilities that are shared by members of society, including perhaps goods that they do not even expect to use themselves. So a CBA based on these prices assumes that the goods in question fulfil only the same indiscriminate want-satisfaction function that is provided by market transactions in ordinary commodities. This also means that the preferences of richer people will have more weight than those of poor people, which may be undesirable for certain categories of goods and services that are publicly provided.

3 PARETO OPTIMALITY AND THE COMPENSATION TEST

As stated earlier, the whole point of an economic CBA is that it shows whether – abstracting from all non-economic considerations and distributional values – the adoption of the project in question would enable

somebody to be made economically better off without anybody else being worse off. For if the economic benefits exceed the costs it is *theoretically* possible for some of the benefits to be transferred from the potential beneficiaries of the project to the potential losers and still leave something over for the former. In other words, the beneficiaries could still have gained even if they have fully compensated the losers. Such a move, therefore, would enable the economy to move towards what has been defined previously as a Pareto optimum point. It is called a '*Pareto optimising move*'.

If compensation is actually carried out – that is, the move satisfies what is known as the 'Hicks/Kaldor compensation test' in honour of its originators, John Hicks and Nicholas Kaldor – the move can be judged to have increased the combined utility of all parties concerned without having to make interpersonal comparisons of utility. This is because there is no need to compare how much utility the gainers gain with how much utility the losers lose. For, after compensation, the losers will have lost nothing.

But an excess of benefits over costs only shows that it is *theoretically* possible for the gainers to compensate the losers while remaining better off. Whether the compensation is *actually* carried out or not is another matter. For this reason a distinction is made between a '*potential*' Pareto optimising move and an '*actual*' Pareto optimising move. And the distinction may often be very important, as illustrated in [Section 5](#).

4 PRACTICAL LIMITATIONS ON THE COMPENSATION TEST

To begin with there are various practical difficulties involved in making the transfer from the beneficiaries to the losers. First, the mere act of making the transfers from the gainers to the losers may involve costs – that is, some new kind of losses. If the transfers take the form of taxes and benefits these could, in theory, distort resource allocation as well as incur administrative costs. For example, they may distort people's incentives to work or invest. So total output (and hence prosperity) may be reduced more than is gained by the improved resource allocation.

Second, it may often be impossible to identify who are the gainers and who are the losers from any project. Consider, for example, a CBA of the location of a new airport. In transport studies the name of the game is usually time – that is, time saved. Suppose that the new airport saves time for the airline passengers to reach their final destination. For business travellers this will be a gain to their companies, or their companies'

shareholders, who may include some insurance companies who have obligations to pay out pensions to retired people. How is one to track down all the individual pensioners and others who may gain indirectly through the project? Obviously impossible. So who would pick up the bill in the end to compensate the losers, such as the people who may have to put up with environmental damage? The taxpayer, perhaps? But this just creates a new class of uncompensated losers.

Third, the costs and benefits have to be ‘cleaned up’ to allow for market distortions, such as those mentioned earlier, namely the existence of taxes and benefits, or externalities, or imperfect competition. This adjustment from observable market prices to what are known as ‘shadow prices’ that are supposed to bear a closer relationship to real social costs and benefits will usually be a highly speculative operation, though in certain circumstances a fairly good approximation can be achieved. However, it does mean that some projects should only be assumed to be desirable if there is a substantial excess of benefits over costs. But, on account of certain theoretical considerations to which we shall now turn, even this is not always a compelling criterion.

5 PARETO OPTIMALITY AND THE DISTRIBUTION OF INCOMES

But the biggest limitation on the role of an economic CBA is its neglect of distributional considerations. It is true that, at first sight, a ‘Pareto-optimising’ move that can make somebody better off without anybody being worse off would seem to be ethically compelling. All that it requires is a little bit of the spirit of beneficence. The trouble is that there are an indeterminate number of Pareto optimal points, corresponding to very different distributions of utility. Consider two people, Smith and Jones, who have been washed up on a desert island. They find that they can manage to catch ten fishes every day (or ten fishes jump out of the sea every day). And let us assume that both like fish. But Smith is a big, tough and selfish person, without a shred of egalitarian or altruistic instincts in him. So every day he eats, with relish, nine of the fish and leaves only one for Jones. But the position is Pareto optimal, since any increase in Jones’s share of the fish must mean a reduction in Smith’s share of the fish. The same would apply if Smith appropriated eight fish every day, or any number of the fish.

In fact, in this example, there are eleven possible ways in which the fish are shared out, ranging from Smith eating all of them to his eating none of

them. All of them would be Pareto optimal points. And in a complex economy with hundreds and thousands of different resources and products there will be a vast indeterminate number of Pareto optimal points, each of which will represent a different distribution of utilities. It is for this reason that Amartya Sen has written that 'a situation may be Pareto optimal but be perfectly disgusting'.⁴ So Pareto optimality may not be such a big deal. It is compatible with any degree of inequality in the distribution of utility, which may, in practice, be closely related to the inequality in the distribution of incomes.

Figure 7.1 shows the relationship of the utility of Mrs A up the vertical axis to the maximum utility that can be enjoyed by Mrs B along the horizontal axis, given the economy's resources. It corresponds to the 'budget line' that plays a crucial part in the elementary theory of how a consumer distributes her expenditures between two goods, given their relative prices and the constraint on her total expenditures. In a similar manner, this Utility Possibility Frontier (UPF) indicates how society could distribute utility between individuals or groups A and B, given its resources and technical knowledge.

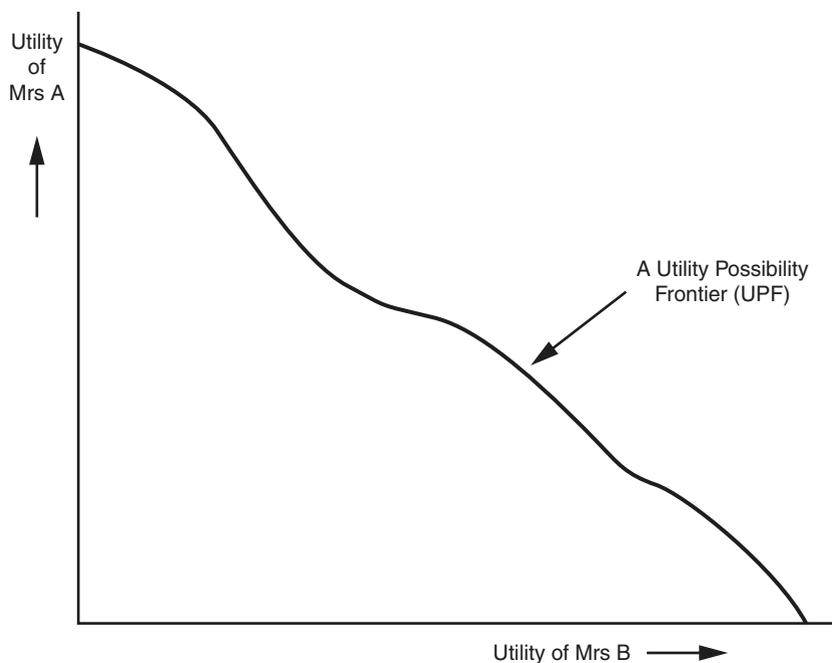


Fig. 7.1 A 'Utility Possibility Frontier'

In the previous diagram, starting from any point on the Utility Possibility Frontier, no move could be justified in terms of the Pareto criterion however much it might mean moving to what society would deem to be a more equitable income distribution. For, by definition of a UPF, somebody must lose by the move. So a point on the UPF at which one agent is very poor and another is very rich is just as Pareto optimal as one where the utilities are shared out more equally. How could society identify which particular point on its UPF corresponds to its optimal choice of how utilities ought to be distributed in society? A method of doing so is analogous to the method used in elementary economic theory to show how an individual consumer chooses her optimal combination of the goods available to her given their prices and her income.

In this theory the consumer selects the optimum distribution of her consumption between various goods (given their prices and her income) in the light of her relative preferences between the goods in question. Diagrammatically, this is shown in a two-good figure as the point at which the budget line representing her constraint is tangential to the highest indifference curve she can reach, given her relative preferences between different combinations of the goods in question. At the level of society's choice between different feasible distributions of utility, what is needed then is to add some function that represents society's preferences between different distributions of the utilities accruing to the various members of society. This is provided by the concept of a 'social welfare function', which will be explained in the next section.

6 INTRODUCING THE 'SOCIAL WELFARE FUNCTION'

The introduction of the concept of a social welfare function seems to provide an escape from the distributional neutrality of the Pareto optimum. It provides a useful analytical tool for clarifying the differences between various value judgements concerning equality, though it is not necessarily confined to this purpose. It can also be invoked in making decisions about any project.

The standard concept of a social welfare function in which the focal variable is consumption has been expressed in the recent *Stern Report on The Economics of Climate Change* as follows:

The objective of policy is taken to be the maximisation of the sum across individuals of *social* utilities of consumption. . . . In particular, we consider consumption as involving a broad range of goods and services that includes

education, health and the environment. The relationship between the measure of social wellbeing – the sum of *social* utilities in this argument – and the goods and services consumed by each household, on which it depends, is called the social welfare function . . . (Stern, 2006, p. 30, Box 2.1; my italics)

In order to bring the social welfare function into relation with the two-dimensional utility possibility frontier in a two-person (or group) figure it needs to be in terms of utilities. Thus one way of interpreting a social welfare function is that it indicates the *social* value that one places on the utility, or welfare, accruing to any individual. These values are sometimes referred to as the 'social utilities' of the individuals' consumption. Hence, economists have generally defined social welfare as some function of individuals' utilities, and abbreviated versions of social welfare functions are generally written as

$$SW = W(U_1, U_2, \dots U_n) \quad (1)$$

where W represents society's total (economic) welfare and U_i represents the social value of the expected utility of the i th individual. As indicated earlier, this is not restricted to the utility an individual derives from his income or his consumption of personal goods, but can include other features of the state of affairs such as the degree of equality. Arrow suggested that an individual's preferences between his own consumption bundles reflect his *tastes*, whereas his preferences between other features of a social state, such as the degree of inequality therein reflect his *values*, though he pointed out that this distinction is by no means clear cut [*loc.cit.* p.18].⁵

The 'social values' attached to individuals' utilities in a social welfare function represent mainly distributional values, though one may not attach a very high social value to some of the ways in which some people derive utility. But for present purposes, we shall focus on the distributional values reflected in a social welfare function. For example, one might attach less *social* value to a marginal increase in the consumption of somebody who started off at a high level of utility than to an equal marginal increase in the consumption of somebody with a relatively low level of utility. In other words, we shall assume that a social welfare function indicates primarily to what extent – other things being equal – one prefers social states that are more equal than others in terms of utilities.

There is, of course, no such thing as 'the' social welfare function. Different people will have different views as to what are the most

important variables that characterise any particular social state as well as the degree of inequality in society that should be promoted. For practical policy purposes, the social welfare function that matters will be the social welfare function of the decision makers. If they did not care about the way utilities are distributed and simply wanted to maximise society's total utility they would prefer the following utilitarian social welfare function, in which W is simply the arithmetic sum of individual utilities and the distribution of utilities does not come into it.

$$SW = U_1 + U_2 + \cdots + U_n \quad (2)$$

But most people do care about the way utilities are distributed among the population. Most people would tend to attach more weight to a unit of utility of a poorer person than to a richer person. This would be the case if one subscribed to some form of egalitarianism (see [Chapter 16](#)). In that case they may decide that social welfare has increased even if an increase in the utility of the poor person is accompanied by an equal, or greater, decrease in the utility of the rich person. An egalitarian social welfare function would thus be convex to the origin, rather like an ordinary consumer's indifference curve relating his preferences between, say, apples and pears. Social welfare would depend on the *relative* utilities of people, as well as their *absolute* utilities. [Figure 7.2](#) compares a utilitarian and a mildly egalitarian social welfare function.

In [Fig. 7.2](#), the point at which society reaches its highest social welfare subject to the constraint given by the UPF will, of course, be that point at which the UPF is tangential to the highest attainable SWF. Only at such a point will it appear that we have satisfied the *necessary and sufficient conditions* for maximising social welfare, namely that we are on a UPF – that is, a Pareto optimal position – but also reaching the highest possible SWF. As indicated earlier, this procedure is analogous to the representation of the optimality of consumers' choice in terms of the tangency between his budget line (the counterpart of the UPF) and his highest possible indifference curves (the counterpart of the SWF). It can be seen from [Fig. 7.2](#) that a mildly egalitarian would prefer a point on the UPF at which the utilities were fairly equally distributed between the two people, A and B. A utilitarian, however, would maximise social welfare at point Y, where there is a very unequal distribution of utility among the agents in question.

For practical policy purposes, however, it is more important to know that a very strong egalitarian might even prefer some points inside a UPF to

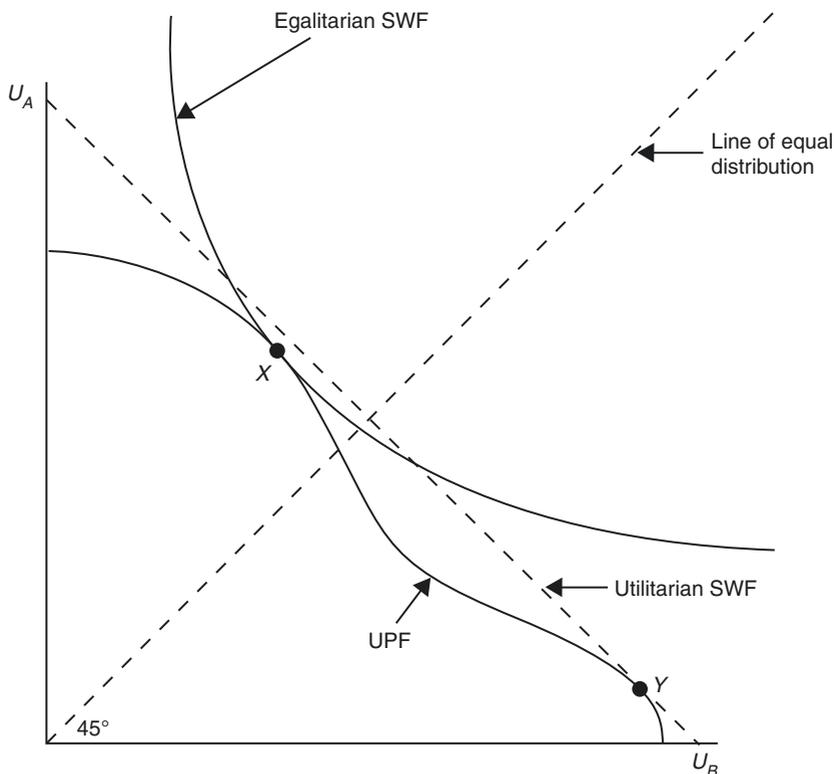


Fig. 7.2 How to reconcile efficiency with social welfare

many points *on* the UPF. Hence, even a Pareto-optimising move (i.e. one that involves no losers) does not necessarily lead to an increase in social welfare. This would be the case, for example, with a move from point I to point II in Fig. 7.3. For in even though B finishes up with slightly more utility than before, A has received far more additional utility. This means that inequality has increased. And given society's strong aversion to inequality – as illustrated by the shape of its SWF – it has actually moved to a lower SWF.

This seems to be precisely what has happened in the Western world over the last two decades or so. As explained in Chapter 15, although even the incomes of the worst off have risen slightly, there has been a much greater increase in the incomes of the richest group in society. The

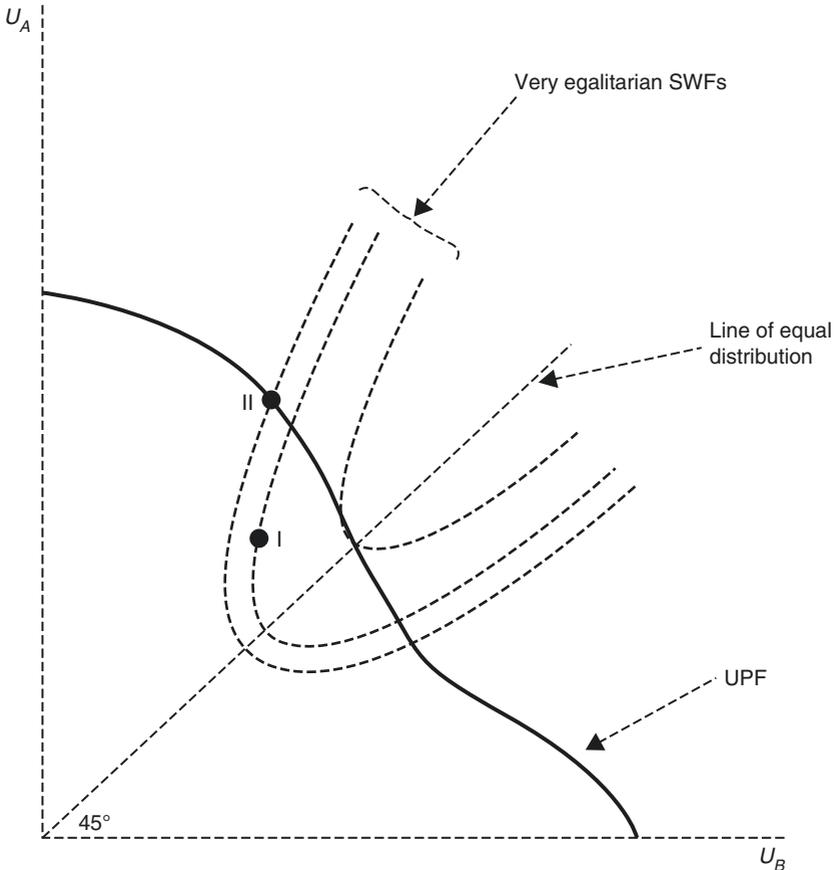


Fig. 7.3 Economic efficiency versus strong egalitarianism

resulting increases in inequality could well have contributed to greater friction and resentment in society.

And conversely, a move from II to I leads to a higher SWF even though it may well fail to satisfy the Pareto optimisation criterion. Thus, for example, a project that led to a slight improvement in the welfare of, say, poor workers (along the horizontal axis) but at considerable loss to rich consumers of their products (along the vertical axis) might be deemed desirable. This might be the case if, say, if poor workers gained a little

welfare through more safety in their working environment, but rich consumers in wealthy countries are estimated to lose more through having to pay higher prices for the goods produced as a result of the better safety regulation. In this case the distribution of income has become more equal even after compensating the losers. So, depending on the importance attached to distributional considerations the Pareto optimality criteria is not a sufficient condition for carrying out the project.

7 CONCLUSIONS

Overall, therefore, the introduction of the social welfare function enables one to see clearly why, if distributional considerations are taken into account, Pareto optimality is neither a necessary nor a sufficient condition for an improvement in social welfare. In simple terms this means that some egalitarians might be willing to trade off a fall in national income in the interests of greater equality. After all, this is a judgement that most of us are prepared to make in many circumstances and reflects the relative intrinsic value we attach to equality and prosperity.

Of course one should not exaggerate the distributional effects of any project. Quite often these will be negligible as far as the whole society is concerned. But it may often be important in a local context. For example some roads may simply serve to help rich landowners or commuters get back home at night, whereas other roads may help poor farmers get their produce to market. There may well be cases where such distributional considerations will outweigh a narrow comparison between social costs and benefits that took no account of them. What the CBA does is to show roughly how much it will cost in certain circumstances to pursue greater equality. It is then up to society (or the decision makers) to decide whether or not it is acceptable.

Furthermore, in a society in which numerous projects are carried out, it may well be that what some people lose on the swings they can more than gain on the roundabout. This is more likely to be the case if the projects are calculated carefully and correspond correctly to the economic concepts of social costs and benefits.

Finally, it is always open to governments to influence the final distribution of incomes in a more or less equal manner, as it wishes, through the mechanism of fiscal policy. Adjusting project evaluation as a means of achieving one's distributional objectives is not the only means available for pursuing this objective. On the other hand, it cannot always be

assumed that any prevailing regime of taxes and benefits is optimal given society's egalitarian values.

However, although the concept of a social welfare function is a valuable tool in many areas of public policy, it is also, in turn, subject to its own limitations. These are discussed in the next chapter.

NOTES

1. Stiglitz, 2012.
2. Sunstein, 2014.
3. Anderson, 1999.
4. Sen, 1970:22.
5. In the Atkinson measure of the equality of income distribution, the utilities of the individuals are replaced by their incomes so that social utility is a function of these incomes only.