

9 What's the Good of Equality?*

JOHN BROOME

9.1 INTRODUCTION

A country's income distribution can have many different shapes. It can be spread out or sharply peaked, skewed or symmetrical, perhaps even bimodal between the sexes. It is obviously impossible to capture all this variety in a single statistic such as the variance.

But this is *not* the difficulty in measuring inequality. The purpose of an inequality measure is not to compress a lot of information into one number. There is often no need to do that anyway. If someone wants to know, say, whether income is more equally distributed now than it was ten years ago, there may be no difficulty about presenting her with the whole distribution for each date. The difficulty is that, when she has all this information, she may *still* not know whether the distribution has become more or less unequal.

What more does she really want to know? Well, she could simply be interested in how she should use the term 'unequal' and apply it to this case. But if so the answer to her question may simply be that our notion of inequality is a pretty vague one, and she can use it how she wishes within broad limits. It is more likely that she wants to know whether inequality has grown better or worse. She is likely to be interested in the badness of the inequality rather than straightforwardly in the amount of it. And that is what I am going to be concerned with in this chapter. There is, of course, a distinction between the amount of inequality and the badness of it,

and it is a good idea to settle from the start which we shall be concerned with. Authors have not always made it clear which their 'measures of inequality' are supposed to be measuring,¹ and the ambiguity can be damaging.²

The question is, then: when is the inequality in one distribution worse than in another? But there is a wider question: when is one distribution worse than another? To ask the more specific question about inequality presupposes that the wider question can be split into two, one question about the relative goodness of the total income, and one about the relative badness of the inequality in its distribution. Is this a reasonable presupposition? Section 9.2 is about that.

9.2 TOTAL INCOME AND ITS DISTRIBUTION

Let me announce from the start that I am not going to try and deal with differences in population. That is a subject fraught with too many difficulties to broach in this paper.³ I shall only try to compare distributions with the same population n . The distribution of income is the vector $y = (y_1, y_2, \dots, y_n)$, specifying each person's income. There is a relation, 'at least as good as' between different distributions. I shall call it, inaccurately, the 'betterness' relation. It is, as a matter of logic,⁴ transitive and reflexive. It is commonly assumed to be a complete preorder. Then, provided it is continuous, it can be represented by a function $G(y)$. $G(y)$ says, up to an increasing transformation, how good the distribution y is. I shall call it a 'goodness function'. It only represents the betterness relation *ordinally*: $G(x) \geq G(y)$ if and only if x is at least as good as y . I doubt that such a function exists. As Sen (1973, 1978) tirelessly reminds us, it is likely that some pairs of distributions will not be comparable in terms of their goodness: neither better than the other, nor equally good either. If so, there will be no goodness function. Nevertheless, most authors have discussed inequality in terms of functions rather than incomplete relations, and for the sake of convenience and simplicity I am going to follow that practice.

The betterness relation will have to be indifferent to permutations of the distribution. Although we are comparing only populations with the same *number* of people, it would be an absurd

restriction if we could only compare populations consisting of the very same people. This means we shall have to assume that the identities of the people do not matter. So it must be the case that $G(x) = G(y)$ whenever x is a permutation of y ; $G(y)$ is symmetrical in the components of y .

This is forced on us, but there is a major assumption implicit in it. People differ, and some may make better use of income than others. In applying the betterness relation to *income* distributions, and then insisting on symmetry, I have made it impossible to take account of these differences. Everyone is being assumed to be the same apart from their income. To decide whether or not one distribution is at least as good as another, we really need to know more about the people receiving the incomes. A more thorough analysis would take account of people's characteristics (e.g. Fine, 1985). It might turn out that it is better to deal, not with income distributions but with distributions of some other variable that is itself a function of income and characteristics. Sen (1985) for instance, is interested in the distribution of 'capabilities', by which he means what people are able to do with their income. A characteristic that certainly makes a difference is how much income a person receives at other times in her life than the one at which we are looking now. Inequality of income is presumably less bad if the people at the bottom are not always at the bottom. So mobility is important, and it might be better to deal in lifetime income rather than current income.⁵ Still, for better or worse, it is distributions of current income with which this chapter is concerned, and the anonymity assumption is unavoidable.⁶

We are interested in whether the betterness relation can usefully be split into something to do with total income and something to do with inequality in the distribution of income. One routine procedure for making a split of this sort has been popularised by Atkinson (1970). For a distribution y , take the *equal* distribution that is equally good: the distribution $\hat{\mu}_e$ such that $G(\hat{\mu}_e) = G(y)$ where e is the unit vector $(1, 1, \dots, 1)$. Let $\hat{M} = n\hat{\mu}$. $\hat{M}$ is an index of the distribution's goodness; better distributions will have a higher $\hat{M}$. It is, in fact, just another representation, like G , of the betterness relation. The difference between $\hat{M}$ and the total income M shows how much income is being wasted, as it were, as a result of the inequality; things could be just as good with that much less income, if only it were equally distributed.

The badness of the inequality may be measured as $I_a = (M - \hat{M})$. This is a measure in money terms. An alternative dimensionless measure is to take the ratio of I_a to M : $I_r = (1 - \hat{M}/M)$. Let us call I_a an 'absolute measure' of inequality and I_r a 'relative measure' – relative, that is, to the total.⁷ Then

$$\hat{M} = M - I_a \quad (9.1)$$

$$\hat{M} = M(1 - I_r). \quad (9.2)$$

Since $\hat{M}$ measures the goodness of the distribution, this says that you can think of the distribution's goodness as made up of M measuring the goodness of the total, less I_a measuring the badness of the inequality. Or else you can think of it as M reduced by the fraction I_r . In this way the badness of inequality has been split off as a consideration on its own. The split may not be a very clean one, though. Nothing in the way that I_a and I_r are constructed guarantees that either of them will itself be in any sense independent of M .

Is I_a or I_r the better measure of inequality? This is a pretty sterile question. Suppose there has been a change that has made I_a go up and I_r go down. (There must have been an increase in M .) Has inequality got better or worse? Well, it has got worse in money terms and better relative to the total. There is not much more to be said. The notion of better or worse *simpliciter* applies to the distribution, not to its inequality. The notion applied to inequality depends on the context. It depends on how the badness of inequality is going to be put together with the other relevant consideration, the goodness of the total, in judging the overall goodness of the distribution. If they are to be put together in an additive way, then I_a is the right measure of badness; if in the fractional way shown in (9.2), I_r . Perhaps it is more natural to put goods and bads together additively. So the absolute measure may be a bit more natural.

There is, however, a substantive question that affects which measure is the more convenient to use (see Kolm, 1976). It can be put like this: should we expect either measure to be independent of total income, and if so which?

Take any distribution y and compare it with αy for some scalar $\alpha > 1$, a distribution where everyone's income is higher in the same

proportion. Which distribution has the worse inequality? It has sometimes been suggested that the inequality is equally bad.⁸ What people mean when they say this is that the relative measure I_r is the same in each. I_r , in fact, is homogeneous of degree zero in y . In this sense it is independent of total income. I_r may consequently be a convenient measure to use. But it would be a mistake to suggest that, if I_r has this property of homogeneity, a proportional increase in everyone's income leaves inequality *really* neither better nor worse. In absolute terms, measured by I_a , it makes it worse.

Now compare y with $(y + \alpha e)$ for some scalar $\alpha > 0$, a distribution where everyone's income is higher by the same amount. It might be suggested that the inequality in *these* two distributions is equally bad.⁹ The suggestion here is that the absolute measure I_a is the same in each. I_a is then a nought-translatable function (i.e. has the property that $I_a(y + \alpha e) = I_a(y)$ for any $\alpha > 0$). In *this* sense I_a is independent of total income. This may make it a convenient measure to use. But again, it would be a mistake to suggest that a uniform increase in everyone's income leaves inequality *really* neither better nor worse. In relative terms, measured by I_r , it makes it better.

The conclusion I draw is that to argue in the abstract about the badness of different patterns of inequality is pointless.¹⁰ The argument must be set in the context of a particular method of putting the badness of inequality into an assessment of the goodness of the distribution as a whole. What we are really concerned with is the form of the goodness function G . I shall concentrate on that, and not on measures of inequality. Once the form of G is settled, a measure of inequality can be split off from it in whatever way suits the context.

However, my discussion of the form of G will quickly reveal that I think inequality is in an important way a consideration on its own, distinct from other considerations that help to determine G . That is the main point of this chapter. To clear the way for it, I first need to raise an objection to an assumption that is commonly made about G .

9.3 SEPARABILITY OF THE BETTERNESS FUNCTION

The assumption to which I am going to object is that G is strongly

separable. First I shall explain what this means. G is a function of the components of the vector y :

$$G(y) = G(y_1, y_2, \dots, y_n)$$

Pick out some collection of these components, say the first k : $y_1, y_2, \dots, y_k$. This collection is said to be *separable* in G if and only if the function G can be written in the form:

$$G(y) = \hat{G}[H(y_1, y_2, \dots, y_k), y_{k+1}, y_{k+2}, \dots, y_n]$$

This says that there is an ordering over the collection, represented by the function H , that is independent of the level of any of the other components of y . G is *strongly separable* if and only if every collection of components of y , not just the first k , is separable in G .

It turns out that G is strongly separable if and only if it is additively separable (Debreu, 1960). For G to be additively separable means that it has the form:

$$G(y) = T[h_1(y_1) + h_2(y_2) + \dots + h_n(y_n)]$$

where T is a positive transformation. Since G is only an ordinal representation of the betterness relation, we may as well divide out the transformation and take G itself to be

$$G(y) = h_1(y_1) + h_2(y_2) + \dots + h_n(y_n) \quad (9.3)$$

Why should we think G is strongly separable? Authors who have adopted this assumption have not often tried to defend it. Dalton (1920) simply announces that it 'is reasonable in a preliminary discussion'. Atkinson (1970) says only that he is following Dalton. Rothschild and Stiglitz (1973) offer a half-hearted defence, concluding only that it is defensible but not compelling. They do little more than state the conditions for strong separability in two different forms, in the hope that one or the other will seem attractive. Strong separability seems to be recognised as an assumption that should be treated with caution. Nevertheless, it is still used (e.g. Fine, 1985; Kanbur and Stromberg, forthcoming). I want to establish that it has no place at all in the theory of inequality.

One thing that has *motivated* this assumption – however it might be justified – is an analogy with expected utility theory. For a moment reinterpret the distribution vector y , as a single person facing some uncertainty about what her income will be. Instead of n people, let there be n states of nature. Let each component of y stand for what the single person's income will be in a particular state of nature. A vector y then stands for a particular 'income prospect'. In expected utility theory the person is supposed to have preferences between prospects like this, and her preferences are supposed to conform to some axioms. Now look at strong separability as I defined it above and give it this new interpretation. Instead of the betterness relation think of the person's preference relation, and think of G as a utility function that represents it. Interpreted this way, strong separability is nothing other than the strong independence axiom or sure-thing principle, one of expected utility theory's fundamental axioms.

If we assume strong separability for income distributions then expected utility theory will be closely analogous to the theory of inequality. Risk-aversion will be analogous to inequality-aversion. Well-known facts about the former can be applied to the latter (see Atkinson, 1970). This makes strong separability an attractive assumption for income distributions. And Rothschild and Stiglitz's (1973) cautious defence of it uses arguments gleaned from expected utility theory.

But comparing it with the analogous assumption in expected utility theory also shows that, applied to income distributions, strong separability is very dubious indeed. In expected utility theory, this assumption has been under heavy fire since it was first developed. It can be defended against these attacks. But the analogous defence is not available at all in the theory of inequality.

Look at this example, which is a version of Allais's famous paradox.¹¹ Imagine one ticket out of a hundred is to be drawn at random. Consider four possible gambles, each determining your annual income according to which ticket is drawn. The prizes are shown in Table 9.1 (the states of nature are ticket numbers). It turns out that a lot of people, when asked to choose between these gambles, prefer w to x , and also z to y . But this is not consistent with strong separability. To see this, look at equation (9.3) and substitute into it the components of w , x , y and z . It will become obvious that if $G(w)$ is greater than $G(x)$ then $G(y)$ has to be greater than

TABLE 9.1 States of nature or people

		1 (£)	2–10 (£)	11–100 (£)
	w	50 000	50 000	50 000
Prospects or	x	1 000	150 000	50 000
distributions	y	50 000	50 000	1 000
	z	1 000	150 000	1 000

$G(z)$. Allais offers this explanation. Strong separability assumes that what happens in one state of nature can be valued independently of what happens in other states. Equation (9.3) makes this clear. But actually, says Allais, there may be interactions or complementarity between states. In the example, gamble w makes it *certain* that your income will be £50 000. It eliminates uncertainty, and there is a special value attached to this. It gives w an advantage over x that y does not have over z . But this fact of certainty is only visible when all the states of nature are looked at together. If you try to value the prize in each state taken separately, it does not appear. According to Allais, then, strong separability is a bad assumption because it ignores interactions between what happens in different states.

But there is a defence against this argument. (Whether it succeeds or fails is irrelevant. See Broome, 1988b.) It has to be conceded that in practice people's preferences are not always strongly separable; Allais's example (and many others) show this. But strong separability can be defended as a condition of *rationality* for preferences. A rational person should realise that there is really no room for interactions between different states of nature. What happens in one state should not affect the value of what happens in another, because if one state comes about the other does not. As Samuelson (1952) says: 'Within the stochastic realm, independence has a legitimacy that it does not have in the nonstochastic realm. Why? Because either heads *or* tails must come up: if one comes up the other cannot.'

Now look again at the example in Table 9.1 but interpret w , x , y and z now as alternative distributions of income for a hundred

people. It is not implausible to think the distribution w is better than x , whereas z is better than y . This conflicts with strong separability. Strong separability requires that distributions should be valued one person at a time,¹² just as in the alternative interpretation it required that prospects should be valued one state at a time. But in support of the view that w is better than x and z better than y , one might point out that w has income equally distributed, and this is a merit that is not shared by y . The equality in w will not emerge if you try to value it one person at a time. Because equality is valuable, there is an interaction between what happens to different people. So strong separability must be rejected.

All this is analogous to Allais's argument. But in the context of inequality the reply I described is not available. Whether or not it was successful, it could obviously only get off the ground because we were dealing with states of nature, and only one of them would actually happen. With an income distribution, each of the different incomes will be received by someone, and all these people will live in one society. We cannot possibly rule out interactions between the incomes of different people.

9.4 UTILITARIANISM AND EGALITARIANISM

Indeed, I think we can go further. It is not simply that we cannot insist on strong separability. To assume strong separability is actually to repudiate a concern, of an appropriate sort, for equality. This claim needs some explaining, since people have often tried to capture the value of equality by means of a strongly separable goodness function. Two related but distinct theories support them in this. In this section and section 9.5 I shall explain why I think both are unsuccessful. And in section 9.5 also I shall explain why I think, quite generally, that true egalitarianism is incompatible with strong separability.¹³

The first theory is straightforward utilitarianism. Suppose that $g_i(y_i)$ is the amount of good the i th person derives from her income y_i . Utilitarianism identifies the overall goodness of a distribution with the total of people's good.¹⁴ So we can write:

$$G = g_1(y_1) + g_2(y_2) + \dots + g_n(y_n)$$

Suppose now that all the functions g_i are the same, and also strictly concave. Then:

$$G = g(y_1) + g(y_2) + \dots + g(y_n) \quad (9.4)$$

where g is strictly concave. This is a strongly separable goodness function that has some egalitarian tendency. For any given total of income, it says it is better for it to be equally distributed rather than unequally. Utilitarianism is against inequality here because it is inefficient in a sense. Units of income are more efficiently used to generate good if they go to the poor rather than the rich.

However, it is well known that utilitarianism has only a very weak tendency to favour equality.¹⁵ It happens to do so if everybody's function g_i is the same and strictly concave, but otherwise it may not. That the function is strictly concave is traditional and plausible wisdom; income has diminishing marginal benefit. That everybody's function is the same is not plausible at all.¹⁶ It is much more plausible to think that some people, the handicapped and people with expensive tastes for instance, derive less benefit from a given amount of income than others. If there is a given amount of income to distribute, the utilitarian formula will then recommend that these people should get less of it than others. Handicapped people, it will say, who need more income than others to achieve the same level of good, should actually receive less. This is not egalitarian. That utilitarianism has egalitarian implications in equation (9.4) is really no more than an accident. And even when its implications are egalitarian, utilitarianism does not capture the spirit of egalitarianism at all. An egalitarian believes that inequality is unfair or unjust. But according to utilitarianism the most that can be wrong with it is that it is inefficient.¹⁷ 'As a framework for judging inequality' says Sen (1973), 'utilitarianism is indeed a non-starter.' I do not need to dwell on this point.

It is possible to move away from utilitarianism in an egalitarian direction, whilst keeping the goodness function strongly separable. This leads to the second theory I mentioned above. I shall call it 'modified utilitarianism'. It starts out from the thought that the reason utilitarianism is not fundamentally egalitarian is because it does not care about equality in the distribution of *good*. To put that right it transforms each person's good by a strictly concave function before adding up:

$$G = t(g_1) + t(g_2) + \dots + t(g_n) \quad (9.5)$$

where g_i is i 's good and t is strictly concave. If there is a given amount of good to be distributed, this formula will favour distributing it equally.

If each person's good is a function $g_i(y_i)$ of her income, then

$$G = \hat{g}_1(y_1) + \hat{g}_2(y_2) + \dots + \hat{g}_n(y_n)$$

where $\hat{g}_i(y_i)$ is $t(g_i(y_i))$. If each function g_i is the same:

$$G = \hat{g}(y_1) + \hat{g}(y_2) + \dots + \hat{g}(y_n) \quad (9.6)$$

where $\hat{g}(y_i)$ is $t(g(y_i))$. Provided g is concave, or not convex enough to outweigh the strict concavity of t , $\hat{g}$ will be strictly concave. Equation (9.6) will be indistinguishable in practice from (9.4).¹⁸ And modified utilitarianism need be no more egalitarian in its practical implications than ordinary utilitarianism. It may, for instance, still recommend giving handicapped people less income than others. Nevertheless, modified utilitarianism does seem more genuinely egalitarian in spirit than the unmodified version. It does, for one thing, care directly about the distribution of good. And second, the function $\hat{g}_i$ depends, not just on g_i , which shows how much good i derives from income, but also on the transformation t . If this is very concave, the differential effect of different g_i functions will be small, and the egalitarian tendency will in practice be strong.

All the same, I do not think that modified utilitarianism represents adequately the value of equality. I shall spell out my reason in section 9.5. But before that I also want to mention a doubt I have about the intelligibility of the theory itself (see Broome, 1987).

In equation (9.5), g_i stands for i 's good. But $t(g_i)$ rather than g_i is the quantity that is added to others to make up overall good. What does $t(g_i)$ stand for? Well, evidently it is the amount that i 's good counts in judging overall good. The theory relies on a distinction between the quantity of i 's good and how much her good counts in overall judgments. I am not convinced that this distinction can really be maintained. When it comes to a choice between two actions, the theory will sometimes say that one of them does more good but the other brings about the best result. This is *prima facie* puzzling.

To be sure, a distinction between good and how much it counts very often makes clear sense. My family's good, for instance, may count for more than other people's in determining what is right for me to do. But this is one specific context. Modified utilitarianism, on the other hand, makes the very general claim that *whenever* it comes to judging overall good, every individual's good must always be transformed by the function t before being added to those of others. In aggregating across people, a quantity of good never appears in the calculation as its own naked self. This makes me wonder what can be meant by the quantity of a person's good as distinct from the transform of it that appears in the calculations. If the transform is the only thing that ever counts, I suspect it must really be the good itself, and that the untransformed version is a sham.

Perhaps modified utilitarianism has been deceived by a sham. But I am not sure. What a modified utilitarian needs to do is produce a way of giving the quantity of a person's good a meaning independently of how much it counts in aggregations across people. Perhaps this can be done. The way to do it is to find a different context – not the aggregation of good across people – where quantities of good are put together *without* being transformed. This could be the context where the notion of the quantity of good gets its meaning. Such a context might, for instance, be the intertemporal aggregation of the good that comes to a person at different times of her life. But there is a difficulty, about finding another context to serve this purpose (see Broome, 1987). Until this is overcome, I shall continue to doubt that modified utilitarianism is really intelligible.

I find it hard to understand for another reason too (or it may be another aspect of the first reason). Suppose we have succeeded in making the necessary distinction between good and how much good counts. It is then a puzzle to know *why* one person's good should count differently from another's. With ordinary utilitarianism it is easy to understand why one person's income should count differently from another's. That is because it does that person a different amount of good. With modified utilitarianism, though, good itself counts differently. Why? The answer can only be that we care about equality, and this is the way that equality is given a value in modified utilitarianism. We give less weight to the good of better-off people. But this is a feeble answer. If equality is valuable,

we ought to be able to explain why it is valuable, and then express its value in a way that properly represents the reason. Certainly the reason has something to do with fairness, so an account of fairness will be required. Modified utilitarianism just takes for granted that equality is valuable and finds an *ad hoc* way of fitting its value into a formula. It chooses this particular way, presumably, because it makes the smallest departure from utilitarianism. Nothing in it suggests to me that it is even trying to capture the idea of fairness. It seems to be activated by a fond attachment to strong separability.

9.5 INEQUALITY AND INCOME COMPARISONS

I do not believe the value of equality can be properly captured by any additively separable function of incomes, modified utilitarianism's or any other. What is wrong with inequality (if anything) is a comparative matter: some people have less income than others. To assess the badness of inequality you must therefore compare together different people's incomes. It is hopeless to try and represent its badness by a function that explicitly eschews such comparisons. I shall elaborate this point.

Imagine a society containing just two people. Compare two alternative prospects for the society, shown in Table 9.2. The brackets show the incomes of the two people. Prospect *U* gives each person an equal chance of getting one unit of income or two units. So does prospect *V*. So both people's prospects of income are the same under *U* and *V*. Under *U*, however, it is certain that income will be distributed unequally. Under *V* it is certain that it will be distributed equally. True egalitarianism therefore implies that *V* is better than *U*.

What does modified utilitarianism have to say about this example? From equation (9.6):

$$G(1,2) = (G(2,1) = (G(1,1) + G(2,2))/2.$$

So the expectation of *G* is the same under both *U* and *V*. It may look as if this immediately commits a modified utilitarian to thinking that these prospects are equally good. But that is not so. The function *G* only represents the betterness relation ordinally.

TABLE 9.2 Income distributions

	States of nature (equally probable)	
	1	2
Prospects <i>U</i>	(1,2)	(2,1)
<i>V</i>	(1,1)	(2,2)

Nothing says it can be put straight into an expected utility calculation. It is open to a modified utilitarian to be risk-loving about *G*. That way prospect *V* would come out better than *U*. And, given equation (9.6), this is the only way it is going to come out better. Since egalitarianism implies *V* really is better than *U*, this is the only way modified utilitarianism can hope to accommodate egalitarianism.

But this is a desperate measure. It is obvious that to take on a risk-loving attitude about *G* is not really going to achieve what is needed. It happens that in the particular example of Table 9.2, equality in income goes along with risk to *G*, so a risk-loving attitude happens to lead to equality. But there is no reason why it should do so in general.¹⁹

The real trouble is that for *V* to be better than *U* is inconsistent with something I shall call *weak prospect separability*. Each prospect such as *U* or *V* gives each person an income prospect. *U*, for instance, gives the first person the prospect (1,2). For generality, write the first person's prospect *Y*₁, the second person's *Y*₂, and so on. Suppose there are *n* people. Let *Y* be the vector (*Y*₁, *Y*₂, ... *Y*_{*n*}); this is a way of writing prospects such as *U* and *V*. *U* is, in fact, ((1,2), (2,1)). *V* is ((1,2), (1,2)). Each component of *Y* is itself a vector. These prospects *Y* will be related by a betterness relation, and if that relation is complete it can be represented by a goodness function *F*(*Y*). Weak prospect separability says that each component *Y*_{*i*} is separable in this function. That is to say, for every *i*

$$F(Y) = \hat{F}(Y_1, Y_2, \dots, Y_{i-1}, f_i(Y_i), Y_{i+1}, \dots, Y_n)$$

This immediately implies that F has the form

$$F(y) = \hat{F}(f_1(Y_1), f_2(Y_2), \dots, f_n(Y_n)).$$

F , in fact, depends only on each vector Y_i taken as a whole. So there is no interaction amongst the components, taken separately, of different vectors.

Weak prospect separability, then, implies that the goodness of U and V in particular will be given by $\hat{F}(f_1(1,2), f_2(2,1))$ and $\hat{F}(f_1(1,2), f_2(1,2))$. These amounts could only be different if $f_2(2,1)$ were different from $f_2(1,2)$. But that is not possible because both the prospects (2,1) and (1,2) mean that the second person has an equal chance of one unit of income or two units. These prospects must therefore have an equal value. So weak prospect separability implies that U and V are equally good.

But according to egalitarianism they are not equally good. Egalitarianism is therefore incompatible with weak prospect separability. Could it still be compatible with strong separability? I think not. In one way weak prospect separability is a weaker condition than strong separability. It only requires the individuals to be separable one at a time, whereas strong separability requires any arbitrary group of them to be separable. In another way, though, weak prospect separability is stronger, because it applies to uncertain prospects rather than just to incomes received for certain. There is no doubt that extending separability to prospects is a big step (see Broome, 1987). But I should be surprised if a believer in strong separability could hold back from taking this step. Since she believes income distributions should be judged taking one person at a time, it would be hard for her to deny that prospects should be judged that way too. Besides, her only escape from the example of Table 9.2, even if she is willing to abandon weak prospect separability, is to take on the bizarre attitude to risk which I described. I think, therefore, that this example is enough to rule out strong separability as an acceptable bedfellow for egalitarianism. If equality is valuable, strong separability has to go.²⁰

Sen's comment about utilitarianism needs to be extended to modified utilitarianism too: as a framework for judging inequality it is a non-starter. I agree wholeheartedly with Temkin (forthcoming):

The problem with this view is clear. It is not concerned with

equality. Equality describes a relation obtaining between people that is essentially comparative. People are more or less equal relative to one another. The view in question is concerned with how people fare, but not with how they fare relative to each other.

9.6 FAIRNESS AND THE VALUE OF EQUALITY

If we are to accommodate the value of equality properly in a goodness function, the first step should be to decide just what, if anything, is valuable about equality in the first place. Too much work on inequality consists of trying out various axioms or various functions to see where they lead, without giving much attention to what could justify them. So I am next going to offer a theory about the value of equality. It will be a theory of fairness, since fairness is the central notion of egalitarianism.

The principal evidence I have to offer in support of my theory is that it is the only adequate way of explaining the fairness of random selection. It often happens that there are several candidates to receive an indivisible good, but not enough is available to go round them all. Haemodialysis is an example. On some occasions like this, it seems that the best way to choose between the candidates is by a random lottery. The advantage of a lottery must be that it is fair way of making the choice. My account of fairness explains why this is so, and I can find no other adequate explanation. So I think that understanding the fairness of a lottery is a useful way of coming to understand fairness in general (see Broome, 1988a).

Suppose some good, divisible or not, is to be distributed amongst a number of candidates. It might, for instance, be income to be distributed amongst a population. For each candidate there will be reasons why she should have the good, or some of the good. One reason why a person should have some income, for instance, might be that she would derive some benefit from it. This is a straightforward utilitarian reason. It would need to be spelled out more fully by specifying her good as a function of the income she gets. Another reason might be that she is entitled to income because she has worked for it. All these reasons have some part to play in determining who should have how much of the good.

But what part? One might simply weigh up the reasons against

each other. Take the good to be distributed one unit at a time. (If it is divisible take very small units, if not take its natural units.) For each unit weigh each person's reasons why she should have it against those of other people. Award the unit to the person with the strongest reasons on balance. Then do the same for the next unit. A utilitarian, for instance, would award each small unit of income to the person who would derive most good from it. The effect is to maximise good overall. Equation (9.4) embodies the results of this process. A modified utilitarian would do much the same, but weight reasons differently. The effect is still to end up maximising some function of the distribution; equation (9.6) specifies it. Generally, if reasons work in a decently coherent fashion, this process will end up maximising some function of the distribution. The process might be said to achieve the maximum satisfaction of reasons.

Weighing reasons, or maximising the satisfaction of reasons, is undoubtedly one consideration that comes into determining the best distribution of the good. I shall call it the 'aggregative' consideration. But it ignores fairness.

To take account of fairness we must start by dividing the reasons why a person should get a good into two classes: 'claims' and other reasons. By a claim to the good I mean a duty owed to the candidate herself that she should have it. Many reasons are not claims. Imagine, for instance, that someone has to be sent on an unpleasant and very dangerous – probably fatal – mission. One person out of a group must be chosen to go. And one of them has special talents that make her more likely than the others to accomplish the mission well. So there are stronger reasons why she, rather than one of the others, should go. The good to be distributed here is the good of being left behind. There are stronger reasons for giving this good to the untalented candidates. But their lack of talent does not give them a stronger *claim* to this good. It is not *owed* to them that they should be left behind.

Claims, and not other reasons, are the object of fairness. Fairness is concerned only with mediating between the claims of different people. If there are reasons why a person should have a good, but she does not get it, no unfairness is done her unless she has a claim to it.

But when it mediates between people's claims, what exactly does fairness require? Does it require simply that claims should get their

proper weight when they come to be weighed against other reasons and the claims of other people? This cannot be enough because *all* reasons should get their proper weight. So does it require, perhaps, that claims should be given extra heavy weight? This would be inadequate too. In the example of the dangerous mission, suppose everybody has the same claim to the good of being excused. So claims are exactly matched. But there is a separate reason why the talented person should not get this good: she would accomplish the mission better. So if claims are simply weighed up, however much weight they are given, the result will be that the talented person gets sent. But it is unfair that she should be required to endanger her life just because of her special talents. She might make this plausible complaint. She has as strong a claim to the good of being excused as anybody else. But because of her talents, the weighing-up of reasons simply amounts to overriding her claim. It is never even on the cards that she might get what she has a claim to. This is not giving her claim the recognition it deserves.

When claims conflict, I believe that what fairness requires is not that they be weighed against each other and other reasons, but that they actually be satisfied in proportion to their strength.

This formula, to be honest, has more precision than I really intend. The essential point is that fairness prescribes how far each person's claim should be satisfied *relative* to the satisfaction of other people's claims. Stronger claims require more satisfaction and equal claims require equal satisfaction. Also, weaker claims cannot simply be overridden by stronger ones: if a stronger claim is satisfied to some extent, then so should a weaker one be to a lesser extent. Fairness is a relative matter. It is not at all concerned with the total satisfaction of claims. If everyone has an equal claim to some good, they are all treated perfectly fairly so long as they all get the same amount of it, even if the amount is small or even none at all. Of course, the more they get the *better*, but not the *fairer*. There is at least one reason for each of them to have the good; this is the reason that constitutes a claim. On *aggregative* grounds, therefore, it is better for them to get more. But *fairness* does not require this.

In the example of the dangerous mission claims are equal. Fairness therefore requires that either everyone goes or no one. Assume that neither of these alternatives is possible. Then the requirement of fairness cannot be satisfied. Some unfairness is

inevitable because someone has to go when others, who have no stronger claim to be excused, stay behind. I argued in my other paper (Broome, 1988a) that the unfairness might be mitigated by holding a lottery. Or it may be right simply to send the talented candidate. This will be unfair to her, but fairness is not an overriding end. It might be right to sacrifice it for other goals on this occasion. One more point. Fairness is a personal good, and unfairness is a personal harm. If a distribution is not fair, that is because someone's claim has not been satisfied in proportion to its strength. This is a harm done specifically to her. Fairness is not some sort of suprapersonal good.

To summarise: There are two sorts of consideration that go into judging the distribution of a good, the aggregative consideration and fairness. The former has to do with maximising the satisfaction of the reasons why people should have the good. The latter is concerned only with some of these reasons, namely claims. And it is not concerned with the aggregate satisfaction of claims, but only with proportionality in their satisfaction. Now to apply this to judging the distribution of income. It will become obvious immediately that what I have been saying leaves some very large issues still to be settled. One of the issues that remains to be settled is what reasons are claims, and that forms the title of Sen's (1980) paper 'Equality of What?' What is it that we should be aiming to distribute equally? This has been much argued over,²¹ and implicitly the argument has been about claims.

I cannot pursue this argument. Instead, for the sake of illustration only, I shall adopt an assumption that I hope is not ridiculous. I shall assume that everyone has an equal claim to income. This is only a vague statement. It does not say, for instance, what happens when people start with unequal amounts of income: does a richer person have no claim until all the poorer people have caught up? But this vague statement is enough to get us going. It is meant to be a gesture in the direction of the view that people have a claim to resources rather than to the good or welfare they derive from resources.²² Income is standing in as a surrogate for resources. Plainly I am ignoring differences in the amounts people work. It is also meant to be a gesture towards the view that because claims are part of the overt commerce of the world, people's claims must be to overt things like income rather than private things like welfare.

Fairness, then, will require that income should be equally

distributed. If it is not, then some unfairness is done. How much? This is another issue which my theory leaves open. This chapter began with the question of how bad is the inequality in a particular distribution of income. We have now arrived at the question of how much unfairness does it do, and we have no answer to that yet. This may not look like progress. But actually we have learned something. First, this unfairness is definitely not a matter of inefficiency or a failure to maximise good; all of that has gone into what I called the aggregative consideration. Second, the unfairness is relative, a matter of how each person fares relative to others. Third, it is a personal harm, suffered specifically by individuals.

This last point gives us a start. We can consider how much unfairness is done to each person. I like the term introduced by Temkin (forthcoming) who calls the amount of unfairness suffered by person i her 'complaint' c_i . A person's complaint will be a function $c_i(y)$ of the distribution. More specifically (the second point above), it will be determined by how much her income is less than other people's. She has the same claim as other people, so unfairness is done her if she ends up with less. But this still leaves open many questions. Is unfairness done to a person whenever she gets less than someone else? Or is she only treated unfairly when she gets less than the mean, which is what she would have got in a perfectly fair distribution? Questions like these and their implications for measuring inequality have been investigated very thoroughly by Temkin (forthcoming). I have nothing to add to his discussion. I shall just mention a couple of examples. Person i 's complaint might simply be how far she falls short of the mean μ :

$$c_i(y) = \max\{0, \mu - y_i\} \quad (9.7)$$

Or it might be the total of the differences between her income and the income of each person who is better off:

$$c_i(y) = \sum_j \max\{0, y_j - y_i\} \quad (9.8)$$

A person's complaint need not represent how badly she *feels* about unfairness; nor need she actually complain. Unfairness is a harm that is done her, but not all harms are necessarily bad feelings. Unfairness is not necessarily connected with envy.

Overall unfairness will be some sort of aggregate $C(c_1, c_2, \dots, c_n)$

of the complaints. To determine the overall goodness of the distribution, this unfairness has to be put together with the other consideration, the aggregative one. The latter will value the distribution according to some function $F(y_1, y_2, \dots, y_n)$, so overall goodness will be:

$$G = \hat{G}(F(y_1, y_2, \dots, y_n), C(c_1, c_2, \dots, c_n)).$$

What this says is that the goodness function $G(y)$ may be written

$$G = \tilde{G}(y_1, y_2, \dots, y_n, c_1, c_2, \dots, c_n)$$

and in this function $(y_1, y_2, \dots, y_n)$ and $(c_1, c_2, \dots, c_n)$ are both separable in $\tilde{G}$.

Now I am going to assume that each pair (y_i, c_i) is also separable in $\tilde{G}$. My grounds are these. First, I take it that person i 's good is entirely determined by y_i and c_i . So it can be written $g_i(y_i, c_i)$. This amounts to assuming that any interaction or complementarity in determining i 's good between i 's income and that of other people has been taken up into c_i ; there is no other source of interaction besides fairness and c_i accounts for that. Second, I take it that the overall goodness of a distribution is entirely determined by the good of the people. I call this the 'principle of personal good', and I have argued for it elsewhere (Broome, 1987). It means that G is a function of individual good:

$$G(y) = \hat{G}[g_1(y_1, c_1), g_2(y_2, c_2), \dots, g_n(y_n, c_n)]$$

And this says that each (y_i, c_i) is separable in G .

The principle of personal good is often called 'individualism' (though other principles are often called 'individualism' too). Earlier I said that egalitarianism is incompatible with the view that overall good is separable between individuals' *incomes* or prospects of income. But it is compatible with individualism. It can accept that overall good is separable between individuals' *good* and prospects of good.

By a theorem of Gorman (1968) all this separability implies that G can be transformed into this additively separable form:

$$G = \bar{g}_1(y_1) + \dots + \bar{g}_n(y_n) - w_1(c_1) - \dots - w_n(c_n)$$

Symmetry requires all the $\bar{g}_i$ functions and all the w_i functions to be the same:

$$G = \Sigma_i \bar{g}(y_i) - \Sigma_i w(c_i) = \Sigma_i (\bar{g}(y_i) - w(c_i)) \quad (9.9)$$

In equation (9.9) the amount $\bar{g}(y_i) - w(c_i)$ is the value given to i 's good in the overall evaluation of the distribution. It would be technically possible to suppose this amount is some transform of i 's good, but I think it is most reasonably taken to be her good itself. The reason is the one I mentioned in section (9.4). This amount is definitely how much i 's good *counts* in the evaluation, and it is hard to see how an intelligible distinction can be made here between how much her good counts and her good itself. Furthermore, there is no point in maintaining a distinction between i 's complaint and how much that complaint counts. We may *define* her complaint as just the amount that the unfairness done her should count (negatively) in evaluating the distribution. So $w(c_i)$ is just c_i . Person i 's good is $\bar{g}(y_i) - c_i$; $\bar{g}(y_i)$ is her good apart from the matter of fairness, and from that you have to subtract the harm done her by unfairness. So

$$G = \Sigma_i \bar{g}(y_i) - J = \Sigma_i (\bar{g}(y_i) - c_i) \quad (9.10)$$

where J is total unfairness $\Sigma_i c_i$.

G , which represents the betterness relation,²³ turns out in equations (9.9) and (9.10) to be the total of people's good. So these equations are utilitarian in a broad sense. Evidently Sen's stricture that utilitarianism is a non-starter as a framework for judging inequality does not apply if utilitarianism is interpreted broadly enough. Utilitarians only need to recognise the specific sort of harm that inequality does: namely unfairness.

However, my theory of fairness is decidedly non-utilitarian in another way. It agrees that the best distribution is the one with the greatest total of good, taking account of unfairness. However, it denies that when there is a choice to be made between alternative distributions, the right one to choose is necessarily the best. Suppose the distribution is at present unequal, and a little more income becomes available. The best result would be achieved by letting this income go to the poorest person. And this would be the best way of reducing the total of unfairness J . But, nevertheless, this may not be what fairness requires. The poorest person has the

strongest claim to the extra income. But other people may have claims too. And if so, these claims should be satisfied in proportion to their strength. At least they should be satisfied to some extent. So other people should get some share of the extra income. The theory I have given, then, must not be understood as just another maximising theory, with a different conception of good to be maximised. But it does *have* a different conception of good, and that is what this chapter is about. The fact that it need not always be in favour of maximising good does not concern us now.

9.7 A MEASURE OF UNFAIRNESS

What will the unfairness measure J be like? This depends on what determines each person's complaint. I mentioned some examples. If c_i is given by equation (9.7), then

$$J = \sum_i c_i = \sum_i \max\{0, \mu - y_i\} = (1/2) \sum_i |\mu - y_i|$$

This is half the absolute mean deviation. If c_i is given by equation (9.8), then

$$J = \sum_i c_i = \sum_i \max\{0, y_j - y_i\} = (1/2) \sum_i |y_j - y_i|$$

This is the 'absolute' Gini coefficient, the common or relative Gini coefficient multiplied by the mean μ (Donaldson and Weymark, 1980).

These two are amongst the many measures of inequality that have been endlessly argued over. Naturally, objections have been made to them. The main objection to the absolute mean deviation is that it does not satisfy Dalton's (1920) principle of transfers. This principle says that a transfer of income from a richer person to a poorer makes the inequality less bad. But it does not increase the absolute mean deviation unless these people are on opposite sides of the mean. The main objection to the Gini coefficient is, I think, best explained by Blackorby and Donaldson (1978). This measure is always in favour of a transfer from a richer person to a poorer, but the value it attaches to such a transfer is independent of the inequality of the distribution. It is plausible perhaps (I am not sure

myself) that the worse the inequality to start off with, the more valuable is a transfer.

These objections are to the consequences of the measures. Until some ground has been given for them, there is no way to judge the measures except by their consequences. But I have suggested a way they might be grounded. I hope this may open up a new approach to judging them.

In any case, I suggested these measures only as examples. They derive from very specific assumptions chosen to illustrate the theory. There are many other possibilities. For instance a person's complaint might be a non-linear function of the difference between her income and the mean. And different people's complaints might not be aggregated simply by adding. More broadly, it may be that people do not have equal claims to income in the first place. We might simply not have any claim to income, in which case an issue of fairness would not arise in the distribution of income. The only issue would be inefficiency. Or our claims might be determined by our labour. And so on. (A most important possibility is mentioned in section 9.10).

Evidently the theory I have described allows a lot of freedom. But there are some things it does not allow. For instance, if the goodness of an income distribution has anything to do with fairness, it is not going to allow the goodness function to be strongly separable. Another example: in recent years various authors (e.g. Mehran, 1976; Donaldson and Weymark, 1980; Yaari, 1987) have proposed to generalise the Gini coefficient. Now, evidently some generalisations of the Gini coefficient would fit nicely into the framework of claims and complaints. But so far as I can tell, the ones that have been proposed are incompatible with it. I have offered a theory of inequality that could underlie the ungeneralised Gini. What theory, I wonder, could underlie these generalisations?

Equation (9.10) shows that unfairness need not be whole badness of inequality. Diminishing marginal benefit of income is still a plausible story, and that will appear as strict concavity in the function $\bar{g}(y_i)$. Inequality will then do two sorts of harm. First it will create unfairness. Second it will cause the total of good apart from fairness to be less than it would have been if the same income had been equally distributed. It will be inefficient in the sense I

mentioned in section (9.4). So inequality can be both inefficient and unfair.

It is possible to quantify the inefficiency on a scale comparable to the other harm J . Let

$$I_d = G(\mu e) - G(y) = n\bar{g}(\mu) - \{\sum_i \bar{g}(y_i) - J\}$$

where μ is the mean of y . I_d is the difference between what G would be if income were equally distributed and what it actually is. G now measures total good. So I_d is a measure of the badness of the inequality. It is an absolute measure, like I_a described in section 9.2, but in units of good rather than money. It is actually the absolute version of the inequality measure originally proposed by Dalton (1920).²⁴ Now let

$$I = I_d - J = n\bar{g}(\mu) - \sum_i \bar{g}(y_i)$$

I is the difference between what G would be if income were equally distributed, and what it actually is, ignoring unfairness. So it is a measure of inefficiency.

$$G(y) = G(\mu e) - I_d = G(\mu e) - I - J$$

Actual good is equally distributed good, less inefficiency and unfairness.

I hope that the separation of inefficiency and unfairness may be useful. Atkinson's (1970) measure, which is derived from either (9.4) or (9.6) above, is evidently aimed at inefficiency. The Gini coefficient is aimed at unfairness.

9.8 FAIRNESS, INEQUALITY AND POVERTY

I shall close this paper with some inchoate remarks about poverty. In section 9.7 I examined the implications of supposing that everyone has an equal claim to income. Now suppose alternatively that we all have an equal claim to enough income to satisfy our basic needs. Needs are, indeed, one of the most plausible sources of claims.²⁵ Perhaps this is the only claim we have on income,²⁶ perhaps it is a particularly strong claim and we also have a weaker

claim to higher levels of income. Let us call the level of income needed to satisfy a person's basic needs the 'poverty line'.

Almost all the analysis in section 9.7 is still applicable. If everyone has the same income, either below or above the poverty line, no unfairness is done. If people have different incomes, but they are all below the poverty line, the analysis of section 9.7 is unaltered. If some people are below the line and some above, then new issues arise. But they will only require us to alter our formula for a person's complaint.

When someone is poor and other people have more income than her, she is being treated unfairly because they have no stronger claim to income than she has. If, in addition, some of them are above the poverty line, then the unfairness done her is much greater. Above the poverty line they have either no claim at all on income or a much weaker claim than hers. Yet they have income that she needs. Her complaint, then, depends very much on the number of people above the poverty line and how much they are getting. The tradition in the literature on measuring poverty, initiated by Sen (1976), has been to make the measures depend on the incomes of the poor only. This may be because the aim has been to produce a measure of the amount of poverty, rather than of its badness. But what I said in section 9.1 about inequality is also true of poverty: we are more likely to be interested in its badness than in its amount.²⁷ The unfairness of poverty is a part of its badness, and it depends strongly on the incomes of the nonpoor.

Unfairness is not the only thing wrong with poverty. Another is simply that living in poverty, without one's basic needs satisfied, is not a good life. Let us assume that a person's good falls precipitously if she drops below the poverty line, so the function $\bar{g}$ in equation (9.10) is very concave around that income. If everyone is below the poverty line, then things are bad because total income in the society is low, and if there is also inequality there will be unfairness and inefficiency of the sort described in section 9.7. If, on the other hand, some people are above the line and some below, there is much greater unfairness as I have explained. There will also be a lot of inefficiency because of the concavity of $\bar{g}$.²⁸

I think, then, that the badness of poverty can be accounted for within the framework I have already developed. The goodness of a distribution can be split into three components, one positive and two negative: the level of the mean or total income, the inefficiency

and the unfairness. Inefficiency and unfairness are the result of inequality. Poverty can add poignancy to all three components, and it can make inequality particularly bad. But the only analytic innovation it calls for is to recognise the especial claim of needs.

10 Institutional Economics

ROBERT A. SOLO

10.1 INTRODUCTION

This chapter is addressed to the student who has worked through this book, who has mastered and feels at home with the microeconomic model, and who would like know how 'institutional economics' relates to what he has already learned; to answer, in effect, what it is about institutional economics that might interest the fledgling economist flushed from a neoclassical nest.

There are today two dominant modes of thought about the economy: the neoclassical and the Marxist. Each of the two develops from a few basic assumptions, a closed model that would encompass and explain the operation of an entire economy. Some believers will claim universality for their model in the interpretation of behaviour and event. Institutional economics, on the other hand, is not an integral system. It offers no encompassing model that would explain the operation of an entire economy. It can be understood rather as clusters of thought that supplement, complement and challenge mainstream economics; clusters of thought that evolved piecemeal and pragmatically in response to perceived inadequacies of the neoclassical paradigm. Institutional economics, reform-oriented rather than revolutionary, reflecting experience where class has a quite different meaning from that in Europe, is perhaps the only significant non-Marxist critique of modern capitalism.

We will try here to comprehend and explain at least the character if not the substance of institutional economics as an offset and a challenge to perceived inadequacies of the neoclassical model. To