5

Approaches for Measuring Inequality

At the end of chapter 2, I observed that once we see what the notion of inequality involves, we may come to think it as either largely inconsistent and severely limited, or complex, multifaceted, and partially incomplete. Even if we ultimately decide that the notion of inequality is "merely" complex rather than inconsistent, as I have been assuming since chapter 2, its *practical* significance will be fairly minor unless we are generally able to compare situations regarding inequality. After all, it can be of little consequence that we regard inequality as bad, if we can't usually tell whether one situation's inequality is worse than another's. Eventually, therefore, the egalitarian would like to arrive at a measure that would enable her to accurately compare situations regarding inequality. In this chapter, I examine some measures of inequality economists have offered. In doing this, I note some significant parallels, as well as points of disagreement, between my work and theirs. I then offer my own suggestion for how one might best proceed if one wants to capture accurately a complex notion like inequality.

Four background remarks. First, though the economists' measures were offered as measures of *economic* inequality, I shall, except for Atkinson's measure, be considering them as they would apply to inequality of welfare. This does not affect my conclusions.

Second, this chapter does not offer (anything like) a comprehensive list of economists' measures of inequality. Nor is its aim to identify and critique the currently most popular measures. In fact, a fair amount of my discussion focuses on measures economists regard as passé. Still, as we will see, there is value in focusing on the measures I do; for they illuminate and support many of this book's claims and, importantly, vice versa.

Third, in this chapter and the following two I offer many criticisms of economists, but few of philosophers. Lest this be misinterpreted, let me point out that this is a case where criticism is a form of flattery. There is virtually nothing in the philosophical literature that is even relevant to the issues raised in these chapters, whereas economists have published work that is not only relevant but interesting and suggestive. Thus, my criticisms of economists' work testifies to the importance and relevance of that work to this book's concerns.

Finally, some readers unfamiliar with economic formulas may have difficulty with sections 5.1 and 5.3. Though I have tried to clearly explain the formulas presented, there is no reason to get bogged down in the mathematics. This chapter's central points should be accessible by simply skimming the formulas and reading the accompanying text. In any event, let me add that later chapters do not presuppose an understanding of the economists' measures.

5.1 The Statistical Measures

In this section, I shall consider the so-called statistical or summary measures. These include: the *range*, the *relative mean deviation*, the *variance*, the *coefficient of variation*, the *standard deviation of the logarithms*, and the *Gini coefficient*. These measures have been widely discussed in the economics literature, and there are well-known difficulties associated with each of them.[1] Still, despite their difficulties, each of these measures has, to a greater or lesser extent, been advocated as a plausible measure of inequality. In the ensuing discussion, I shall briefly examine each of these measures and suggest the source of their appeal. In doing this, important similarities will be noted between these measures and the theoretical results of chapter 2.

A preliminary comment. One criticism often made of the Gini coefficient is that it is implausible for comparing situations whose average levels differ greatly. A similar criticism has been made of the variance and, for reasons to be presented in chapter 6, I think each of the statistical measures lose plausibility when applied to such situations. In response to such criticism, advocates of the Gini coefficient have contended that even if the Gini coefficient is not acceptable for comparing all situations, it is still a useful and plausible measure for comparing situations whose average levels do not differ greatly. A similar claim might be made on behalf of the other statistical measures. Therefore, throughout the following discussion I shall assume that the statistical measures are to be employed for comparing situations that have the same number of people, the same total amount of welfare, and, hence, the same average levels of welfare. This assumption enables me to ignore μ and n in the statistical measures and, in so doing, ensures that I am comparing situations for which the statistical measures seem most plausible when I consider why they have been offered as measures of inequality. (The reason the statistical measures seem most plausible on the assumption in question will become clearer in chapters 6 and 7 where the roles of μ and n are explained and criticized. Briefly, it can be put as follows. The statistical measures contain certain features—μ and/or n—that yield controversial answers to the questions of whether, and to what extent, inequality is affected by a population's level or size. However, this is not intuitively evident except in cases where the population's level or size varies greatly. In fact, it turns out that the statistical measures' controversial features can be ignored in cases where the numbers of people and total amounts of welfare are the same, since in such cases the measures' orderings are not influenced by the features in question. Intuitively, then, certain shortcomings of the statistical measures will be most evi-

1. Among the many discussions of these measures in the economics literature, I have found A. K. Sen's to be particularly useful. In presenting these measures and the shortcomings they face, my discussion closely follows the one he provides in *On Economic Inequality* (Clarendon Press, 1973), chap. 2.

dent in cases where the levels or sizes of the populations vary; hence, the statistical measures will seem most plausible for comparing situations in which these factors are constant.)

So one way of looking at my task in this section is as follows. Statisticians have devised a large number of formulas to measure or reveal various features of a given population or curve. Of these, several have been offered by economists as plausible measures of inequality. I want to suggest what it is about the selected measures that, at least for a certain range of cases, gives them their plausibility as measures of inequality.

The Range

Let me begin my consideration of the statistical measures by looking first at the range. The range is the simplest of the statistical measures of inequality. Its formula is given by:

$$E = (\mathrm{Max}_i\, y_i - \mathrm{Min}_i\, y_i)/\mu$$

where $\mathrm{Max}_i\, y_i$ and $\mathrm{Min}_i\, y_i$ are the levels of welfare of the best- and worst-off individuals, respectively, and μ is the society's average level of welfare. Now clearly, for the sort of worlds I am presently interested in, that is, worlds where the total amount of welfare and the number of people are fixed, the average level of welfare, μ, will be constant. This means that for the purpose of ordering such worlds, the denominator of E may be ignored, as E's orderings will be identical to E*'s, where

$$E^* = (\mathrm{Max}_i\, y_i - \mathrm{Min}_i\, y_i)$$

According to the range, then, A's inequality will be worse than B's if and only if the gap between the best-off person and the worst-off person is larger in A than it is in B.

Serious criticisms have been leveled at the range as a measure of inequality. The most obvious is that it completely ignores the pattern of distribution between the extremes. Consider diagram 5.1.

All things considered, most would probably say that B's inequality was worse than A's, since in B one-half of the population is much worse off than the other, whereas in A the vast majority of the population is perfectly equal, though a few are fortunate to be better off, and a few are unfortunate to be worse off. The range, of course, focusing as it does on the extremes, yields the judgment that A's inequality is worse than B's. This sort of result has led many to reject the range as a plausible measure of inequality.

A second criticism of the range, closely connected with the first, is that it violates the Pigou-Dalton condition, according to which any (even) transfer from a worse-off to a better-off person worsens inequality.[2] Many have felt that any adequate measure of inequality must meet at least this minimal condition. Clearly,

2. See sections 3.4 and 3.6 for an extended discussion of the Pigou-Dalton condition and an explanation of the qualifying parenthetical "even."

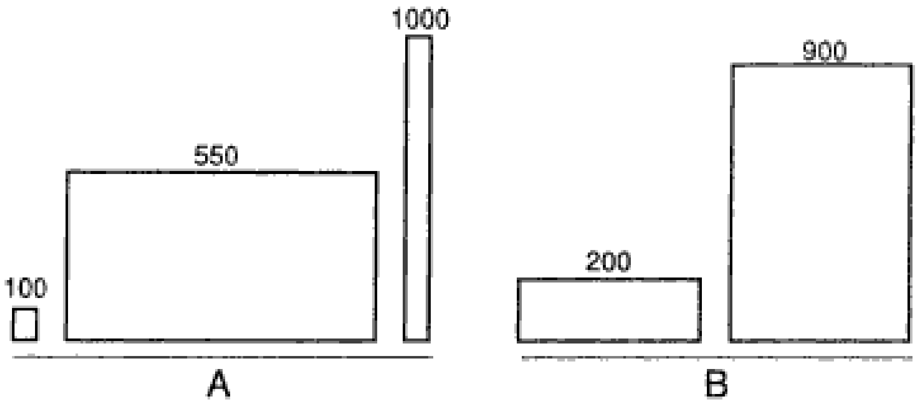

DIAGRAM 5.1

however, the range fails to do this, as it fails to reflect transfers between people who are not at the extremes.

To these criticisms, the general considerations of chapters 2 and 3 might be added. We have seen that our notion of inequality is complex and multifaceted. Correspondingly, I would argue that the range is inadequate as a measure of inequality because it does not recognize and capture the full complexity of that notion.

In light of such shortcomings one might wonder why the range was ever offered as a measure of inequality? To this question, I suggest the following answer. In chapter 2, I argued that in accordance with certain intuitions the size of someone's complaint regarding inequality would depend upon how that person fared relative to the best-off person. I also argued that in accordance with certain other intuitions—maximin intuitions—the worse of two worlds regarding inequality would be the one in which the worst-off person had the most about which to complain. Together, these intuitions combine to support the judgment that how bad a world is regarding inequality depends upon how badly the worst-off person fares relative to the best-off person. Now it is essentially *this* judgment that the range yields. I believe, therefore, that it is from these two sets of intuitions that the range derives its plausibility.

Insofar as we are caught in the grip of the intuitions underlying the relative to the best-off person view of complaints and the maximin principle of equality (see chapter 2) a measure like the range will have great plausibility. However, it must be recognized, as advocates of the range apparently did not, that the intuitions in question are not the *only* plausible intuitions underlying and influencing our egalitarian judgments. Thus, advocates of the range were mistaken in simply offering the range as a measure of inequality, rather than as a measure of one particular aspect of that notion. Similarly, however, opponents of the range were mistaken in simply focusing on its shortcomings as a measure of inequality and then dismissing it as implausible. Had they carefully considered the range's properties and the orderings it yielded, they might have seen that despite its faults it had a fair amount of intuitive appeal. This, in turn, may have led them to see that *certain* egalitarian intuitions would support a measure like the range.

The Relative Mean Deviation

In certain respects my discussion of the relative mean deviation, as well as the other statistical measures, closely resembles my discussion of the range.

The formula for the relative mean deviation is:

$$M = \sum_{i=1}^{n} |\mu - y_i| / n\mu$$

where, as applied to any society, S, n is the size of S's population, y represents, in turn, the level of each of S's members, and μ is, as before, S's average level. Once again, for worlds of the sort I am presently interested in—where n and μ will be constant—the denominator of M may be ignored. For such worlds, M's ordering will be identical to M*'s, where:

$$M^* = \sum_{i=1}^{n} |\mu - y|$$

M* sums up the absolute value of the differences between the average level and the level of each individual. In doing this, M* measures the total deviation from the average. So, according to the relative mean deviation, A's inequality will be worse than B's if and only if the total deviation from the average is larger in A than it is in B.

As with the range, the relative mean deviation has been criticized for violating the Pigou-Dalton condition, as it is insensitive to transfers within the same side of the mean. So, for example, while a transfer of 10 units of welfare from someone 1,000 units below the mean to someone 20 units below the mean would worsen inequality according to the Pigou-Dalton condition, such a transfer would have no effect according to the relative mean deviation as the total amount of deviation from the mean would remain unchanged.

Given that the relative mean deviation is insensitive to transfers *within* the same side of the mean, it directly follows that it is insensitive between transfers of any given size *to* the same side of the mean. As we saw in chapter 3, this is another feature many would find objectionable. All things considered, most egalitarians believe it would be better for an above average person to transfer n units of welfare to someone way below the average than to someone only a little below the average—that, for instance, it would be better to transfer 10 units of welfare to someone 1,000 units below the average than to someone 20 units below the average. However, according to the relative mean deviation there would be no reason to prefer the one transfer to the other, as both would reduce the total deviation from the mean by the same amount. (The reader will notice that a similar criticism could have been made against the range.)

Again, to these criticisms the considerations of chapters 2 and 3 might be added. Those considerations indicate that the relative mean deviation is inadequate as a measure of inequality because it fails to capture fully the complexity of that notion.

Like the range, the relative mean deviation has evident difficulties. So, as with the range, we may wonder why the relative mean deviation was offered as a measure of inequality. Here, too, I think our answer lies in the fact that while the relative mean deviation may not capture the *whole* of our notion of inequality, to a large extent it captures one very important *aspect* of that notion. Specifically, it reflects the combination of those intuitions underlying the relative to the average view of complaints and the additive principle of equality; for, as I pointed out in chapter 2, together those intuitions support the judgment that the worse of two worlds is the one in which the deviation from the average is greatest.

It is not surprising, then, that the relative mean deviation should have been offered as a measure of inequality. Because it appears to capture certain of our egalitarian intuitions, it will seem plausible insofar as one is caught in the grip of those intuitions. These considerations suggest that, as with the range, it would be a mistake to simply regard the relative mean deviation as a measure of inequality and dismiss it as inadequate. To do this would be to miss the source of its appeal—to fail to realize that it may capture some, though not all, of our egalitarian intuitions.

The Variance

Let us next consider the variance, whose formula is:

$$V = \sum_{i=1}^{n} (\mu - y_i)^2/n$$

For worlds of the sort we have been considering, V's ordering will be the same as V*'s, where:

$$V^* = \sum_{i=1}^{n} (\mu - y_i)^2$$

This means that for such worlds the only significant difference between the variance and the relative mean deviation is that the variance first squares the differences between the average level of welfare and the level of each individual before adding them together. (This is the only relevant difference between V* and M*.) The effect of this squaring feature is to give larger differences from the mean extra weight. For example, being 10 units worse off than the average would count for more than simply twice being 5 units worse off than the average, as the former would increase the variance by 100, whereas the latter would increase it by 25.

Because of its squaring feature the variance avoids two of the criticisms raised against the relative mean deviation. It does not violate the Pigou-Dalton condition and is not indifferent between transfers to the same side of the mean. (So, if p is worse off than q, it will always be better to raise p, n units, rather than q.) Still, the variance is not without problems as a measure of inequality. From a moral perspective, the squaring feature is an arbitrary way of reflecting the view that large deviations should be given extra weight. This renders the accuracy of the variance open to doubt, thereby lessening its plausibility. In addition, the variance is liable to the familiar general criticism that it does not fully capture and reflect inequality's complexity. (As noted earlier, the variance is often criticized as implausible when comparing situations whose average levels vary greatly. As this criticism does not apply to the kinds of situations I am presently interested in, I shall not discuss it here. However, this is an important line of criticism I shall return to in chapter 6.)

In chapter 2, I argued that the intuitions underlying the relative to the average view of complaints could combine with the intuitions underlying the weighted additive principle to influence our egalitarian judgments. Moreover, I pointed out that according to these intuitions, inequality might be measured by, first, comparing the level of each individual to the average, second, attaching weight to the figures thus arrived at so as to give more weight to large deviations, and, third, summing up the weighted figures. Now the variance measures inequality in just such a manner. I believe, therefore, that it is from the intuitions underlying the weighted additive principle and the relative to the average view of complaints that the variance derives its plausibility.

This conclusion should not be surprising given that (1) for worlds of the sort being considered the main difference between the variance and the relative mean deviation is that the variance attaches greater weight to large differences from the mean, (2) the relative mean deviation reflects the intuitions underlying the additive principle and the relative to the average view of complaints, and (3) the difference between the additive principle and the weighted additive principle is that the latter sums up complaints only after first weighting them so as to give extra weight to large complaints.

It appears, then, that like the range and the relative mean deviation, the variance reflects some, but not all, of our egalitarian intuitions. This helps to explain both its plausibility and its limitations. To the extent the intuitions underlying the weighted additive principle and the relative to the average view of complaints influence our egalitarian judgments, the variance will seem plausible. Since, however, the intuitions in question are not the only ones influencing our egalitarian judgments, the variance is best regarded as a measure of one aspect of inequality, rather than as a measure of inequality itself.[3]

The Coefficient of Variation

My discussion of the coefficient of variation will be brief, since for worlds of the sort being considered there is no significant difference between the coefficient of variation and the variance. The coefficient of variation is simply the square root of the variance divided by the mean income level. Thus its formula is:

3. Note that in explaining the variance's plausibility I say it is best to regard it as a measure of one aspect of inequality. I have not claimed it is the most plausible measure of that aspect. For example, the morally arbitrary squaring feature is no more plausible in a measure of an aspect of inequality than it is in a measure of inequality itself. More generally, as I shall stress later, I think the statistical measures are best regarded as reflecting inequality's aspects, but that even in that capacity they are inadequate and best regarded as first approximations. Better measures must still be developed to capture inequality's aspects fully and accurately.

$$C = \sqrt{V}/\mu \text{ or } C = \sqrt{\sum_{i=1}^{n}(\mu - y_i)^2/n}/\mu$$

Now for the sort of worlds under consideration C's ordering will be identical to C*'s, where

$$C^* = \sqrt{V}$$

Similarly, C*'s ordering will be identical to C**'s, where

$$C^{**} = V$$

For such worlds, then, C's ordering will be the same as V's, as in fact the relevant portion of both formulas for comparing inequality in such worlds is that corresponding to C***, where

$$C^{***} = V^* = \sum_{i=1}^{n} (\mu - y_i)^2$$

These considerations suggest that for the sort of worlds being considered the coefficient of variation will have the same strengths and weaknesses as the variance, and for just the same reasons. Like the variance, therefore, I believe the coefficient of variation is best regarded as a measure of one aspect of inequality. More specifically, I believe the coefficient of variation corresponds to those intuitions underlying the weighted additive principle and the relative to the average view of complaints.

The Standard Deviation of the Logarithm

Let us next consider the standard deviation of the logarithm. Its formula is:

$$H = \sqrt{\sum_{i=1}^{n} (\log \mu - \log y_i)^2/n}$$

For comparing worlds with the same number of people the relevant portion of H will be that corresponding to H*, where

$$H^* = \sum_{i=1}^{n} (\log \mu - \log y_i)^2$$

H* may look familiar. It bears a close resemblance to C*** and V*. The sole difference between C*** and H* is that H* weights the deviations from the average even before it squares them and adds them together. It does this by comparing the *logarithm* of each individual's level to the logarithm of the average, rather than by directly comparing each individual's *actual* level to the average. Now the logarithmic function attaches proportionately greater weight to small numbers than to large ones. As a result, H* attaches more weight to deviations below the average than to deviations above the average. It also attaches more weight to transfers occurring at the low end of the scale than to transfers occurring at the high end of the scale.

The foregoing considerations suggest the following. Like the variance and the coefficient of variation, the standard deviation of the logarithm corresponds to the relative to the average view of complaints and a weighted additive principle. However, the standard deviation of the logarithm differs from those measures in the *manner* in which it weights deviations from the mean. Whereas the variance and the coefficient of variation weight deviations above and below the mean similarly, the standard deviation of the logarithm weights deviations below the mean more than deviations above it. Still, as with the variance and the coefficient of variation, the standard deviation of the logarithm attaches extra weight to large deviations *below* the mean. So, for example, being ten units below the mean is more than twice as bad as being five units below the mean.

Various criticisms have been leveled at the standard deviation of the logarithm as a measure of inequality. Like the variance and the coefficient of variation, the standard deviation of the logarithm involves a morally arbitrary squaring feature. Similarly, the logarithmic transformation is a morally arbitrary way of capturing the view that extra importance should be attached to transfers at the low end of the scale. Additionally, it has been claimed that the standard deviation of the logarithm becomes *so* insensitive to transfers between people who are *very* well-off that it can end up violating the Pigou-Dalton condition. But perhaps the main criticism needing to be made is that it is a mistake to consider the standard deviation of the logarithm as a measure of the *whole* of our notion of inequality, as surely it does not capture the full complexity of that notion. Still, as with the other measures, it is better to try to discern the source of its plausibility than simply to dismiss it. As indicated, I believe the standard deviation of the logarithm is best regarded as giving expression to a weighted additive principle and the relative to the average view of complaints.

Before going on, let me comment on the apparent tension between the last several measures discussed. I have suggested that the variance, the coefficient of variation, and the standard deviation of the logarithm each give expression to a weighted additive principle and the relative to the average view of complaints. I have also noted that while there is nothing to choose between V and C for worlds of the kind we are now discussing, there *is* a significant difference between those measures and H in the way they weight deviations from the mean. It might seem, therefore, that H is in competition with V and C, and that we will ultimately have to choose between them and decide which sort of measure best captures the aspect in question. However, in the spirit of the rest of this book, I would suggest that before rejecting any of the measures, one should be sure they are indeed measures of the *same* aspect of inequality. In fact, it is arguable that while certain of our intuitions support a measure like V or C, others support a measure like H.

More particularly, I would suggest the following. Given the kinds of considerations underlying the relative to the average view of complaints, it may seem natural and plausible to count deviations above and below the average similarly. This is because among equally deserving people, it may seem just as bad for someone to have *more* than his fair share as for someone to have *less* than his fair share. That is, it may seem just as unfair or unjust for Fate to discriminate in *favor* of someone (to the disadvantage of others) as for Fate to discriminate *against* someone (to the advantage of others). Nevertheless, despite this position's appeal, there is also appeal to the view that deviations below the average should count for *more* than similar deviations above the average. I suspect this may be connected with an asymmetry in our thinking about the lucky and the unlucky. While we empathize with the unlucky person who suffers misfortune, and tend to feel that his burdens should be shared evenly, we tend not to begrudge the lucky person his good fortune, and tend not to feel that his benefits should be shared evenly. Such feelings intuitively support the view that deviations below the average should count for more than deviations above the average. (See the connection, noted in chapter 2, between our notion of luck and the relative to the average view of complaints.)

Despite their similarities, then, H may not simply be a more or less accurate version of C and V. Different intuitions may be underlying the different ways they weight deviations from the mean. If this is so, then it is important to identify those intuitions, and to assess their plausibility. Upon reflection, we may decide there is yet another element of inequality's complexity, and that H is best regarded as expressing that element. Alternatively, if H ultimately rests on a bias in favor of the lucky, we may decide that, despite its appeal, it should be rejected. This is because it may seem that, insofar as inequality is one's concern, such a bias is irrational and not to be accorded any weight.

A word of caution. In discussing the statistical measures we have seen how one might adopt two different views insofar as one is influenced by the relative to the average view of complaints and the weighted additive principle, one of which would be roughly captured by C and V, and one of which would be roughly captured by H. It is perhaps worth emphasizing that our discussion does not rule out the possibility that on reflection one might find other views plausible as well, or perhaps instead. Diagram 5.2 illustrates some of the views one might consider.[4] Each graph represents a possible weighting curve with the origin representing the average level, the y axis representing deviations above and below the average, and the x axis representing the amount of weight attached to deviations such that greater weight is represented by those points that are further to the right along the x axis.

A is merely for illustrative purposes. It represents the straight—that is, nonweighted—additive principle against which various weighted principles might be compared. The vertical line runs through the x axis at point one implying that deviations above and below the average will each receive their "normal" or "full"—that is, nonadjusted—weight. On this view a deviation of n above the average would count exactly as much as a deviation of n below the average, and half as much as a deviation of 2n above or below the average.

B represents the view expressed by the Variance and the Coefficient of Variation. Deviations above and below the average count equally, with extra weight being attached to greater deviations.

C represents the view expressed by the Standard Deviation of the Logarithm. Deviations below the average count more than deviations above the average, with

4. The following diagrams, and their interpretations, were suggested by Shelly Kagan.

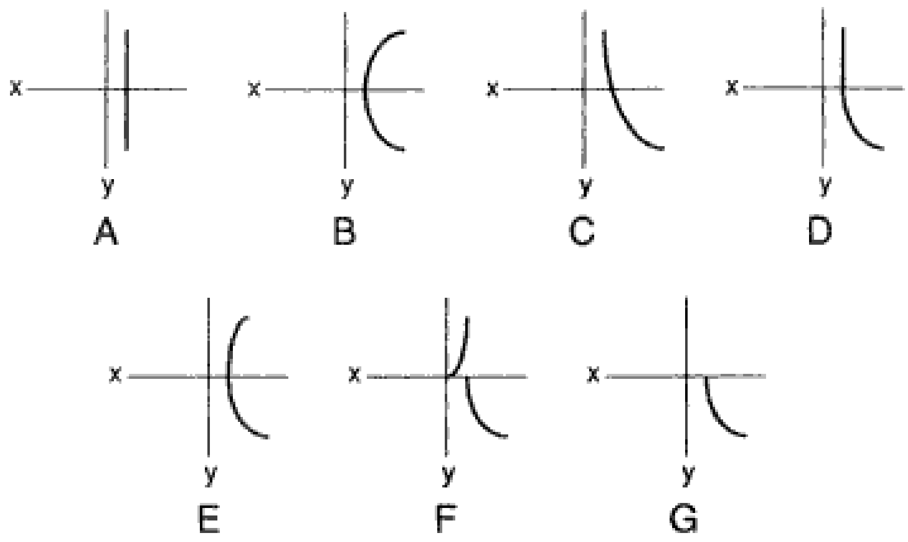

DIAGRAM 5.2

the upper half of the curve asymptotically approaching the line where deviations would only receive their "normal" or "full" (nonadjusted) weight. Notice, on this view, although large deviations above the average might actually count more in the end than small deviations about the average, the latter would be *weighted* more than the former.

D represents the view that while deviations above and below the average are both objectionable, only the latter are especially objectionable the larger they are. On this view extra weight is only attached to greater deviations below the average, with deviations above the average receiving their "normal" or "full" weight.

E represents the view that deviations above and below the average should both be weighted so that extra weight is attached to greater deviations, but that deviations below the average should count more than comparable deviations above the average.

F represents the view that extra weight should be attached to greater deviations below the average, that slight deviations above the average hardly matter at all (they approach zero weighting), and that greater deviations above the average become increasingly objectionable approaching the point where they deserve to be given their "full" or "normal" weight.

Finally, G represents the view that extra weight should be attached to greater deviations below the average, but that only deviations below the average matter—that is, we shouldn't care about deviations above the average *at all*.

I have not tried to motivate B–G, and shall not explore them further in this work (though I will confess I find E particularly appealing). I mention them because they remind us of how much work remains to be done in addressing the issues we have raised. In deciding how best to capture the positions underlying the relative to the average view of complaints and the weighted additive principle, a number of positions might be considered. These include, but are by no means limited to, C, V, and H.

One final comment regarding H. I have argued that H corresponds to a relative to the average view of complaints and a weighted additive principle. If this is so, then H's plausibility is not automatically enhanced because it captures the view that transfers between people above the average matter less than similar transfers below the average. Though very appealing, the plausibility of *this* view *may* depend on the intuitions underlying the relative to the best-off person and the relative to all those better-off views of complaints, both of which capture the view in question when combined with a weighted additive principle. The question is whether certain intuitions underlying the relative to the average view of complaints may also have this implication. Leaving this question open, let us turn to the final statistical measure to be considered.

The Gini Coefficient

At one time, the Gini coefficient was perhaps the economists' most popular measure of inequality. Probably the most common way of looking at the Gini coefficient is in terms of the Lorenz curve, which plots the percentage of a population from the least well-off to the most well-off along a horizontal axis, and the percentage of welfare that the bottom n percent of the population has along a vertical axis. Plotted on a graph like diagram 5.3, the Lorenz curve always runs from the lower left hand corner to the upper right hand corner, since 0 percent of the population always has 0 percent of the welfare, and 100 percent of the population always has 100 percent of the welfare.

In diagram 5.3, the diagonal, e, is the line of absolute equality. It would be the Lorenz curve of a situation where everyone was perfectly equal, where for all n, n percent of the population had n percent of the welfare. C, the lower and right-hand

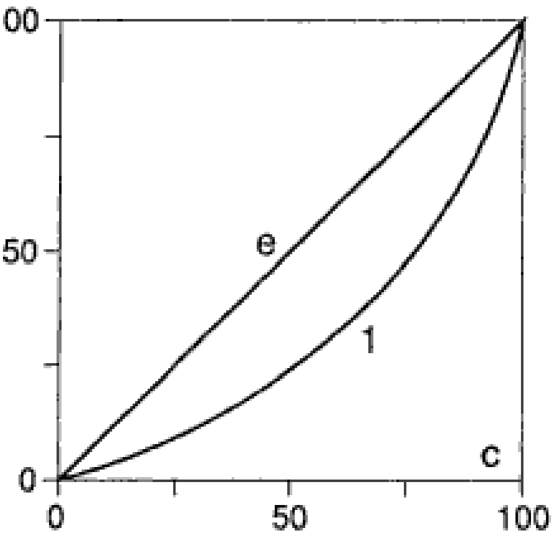

DIAGRAM 5.3

border of diagram 5.3, would be the Lorenz curve of a situation where one person (somehow) had all the welfare, since for all n < 100, the bottom n percent of the population would have 0 percent of the welfare, while, as noted already, 100 percent of the population would have 100 percent of the welfare. L would be the Lorenz curve of a fairly typical situation. Obviously, like l, the Lorenz curve for most situations will lie between e and c.

The Gini coefficient can now be defined. It is the ratio of the area between the line of absolute equality and the Lorenz curve to the triangular area underneath the line of absolute equality. Thus, the Gini coefficient of a world whose Lorenz curve is l, is given by the formula:

$$G = \frac{\text{the area between between e and l}}{\text{the area between e and c}}$$

Obviously, the Gini coefficient can range from zero, when the Lorenz curve corresponds to e, to one, when the Lorenz curve corresponds to c. Correspondingly, for advocates of the Gini coefficient, the better of two worlds regarding inequality will be the one whose Gini coefficient is smaller; intuitively, the one whose Lorenz curve lies closer to the line of absolute equality.

Various critics have shown that the Gini coefficient faces serious shortcomings as a measure of inequality.[5] Still, the Gini coefficient has great intuitive appeal. I think, therefore, that as with the other measures it would be a mistake simply to dismiss it. Instead, one should try to determine the source of its appeal.

Despite the popularity of looking at the Gini coefficient in terms of the Lorenz curve and graphs like diagram 5.3, perhaps a more useful way of looking at it (at least for our purposes) is in terms of the relative mean difference, which is the arithmetic average of the absolute value of the differences between all pairs of incomes. It turns out that the Gini coefficient is equal to one-half of the relative mean difference. Thus, an alternative formula for the Gini coefficient is:

$$G = (1/2) \sum_{i=1}^{n} \sum_{j=1}^{n} |y_i - y_j| / n^2\mu$$

Since for worlds of the sort we are now considering, n and μ are constant, G's ordering will be the same as G*'s, where

$$G^* = (1/2) \sum_{i=1}^{n} \sum_{j=1}^{n} |y_i - y_j|$$

But G*'s ordering will be identical to that yielded by a new measure, K,[6] where K

5. See, for example, Sen's *On Economic Inequality*, A. B. Atkinson's *The Economics of Inequality* (Clarendon Press, 1975), and David Newbury's "A Theorem on the Measurement of Inequality," *Journal of Economic Theory* 2 (1970): pp 264–66.

6. The reader may be puzzled by my introduction of K, since my argument could be made without reference to it. However, though K is not *necessary* for my argument, there are various respects in which dispensing with K would make my presentation awkward or slightly misleading. Thus, in this case I have opted for a smoother, hopefully clearer presentation to an equivalent shorter, but more cryptic one.

measures inequality by adding up the absolute values of the differences between each person and all those better off than he. To see this, consider the following.

G* measures inequality by adding up the absolute values of the differences between each person and all of the members of his world (including himself), and dividing the figure thus arrived at by two. Thus,

$$G^* = 1/2\ (a + b + c)$$

where a is the sum of the absolute values of the differences between each person, and all those better off than he (so a = K), b is the sum of the absolute values of the differences between each person and all those (including himself) as well-off as he, and c is the sum of the absolute values of the differences between each person and all those worse off than he. The value of b will of course be zero, so

$$G^* = 1/2(a + c)$$

Moreover, the value of a will be identical with the value of c, since for each component of a there will be an equivalent component of c. Suppose, for instance, that x who has n units of welfare, is better off than y, who has m units. Then one of a's elements will be $|m - n|$, and similarly, one of c's elements will be $|n - m|$. But obviously this will hold for any two people x and y, where x is better off than y. Therefore, since for all n and m, $|m - n| = |n - m|$, a = c. So

$$G^* = 1/2\ (a + a) = 1/2\ (2a) = a = K$$

Thus, for the sort of worlds we are now considering, the gini coefficient's ordering will be identical to K's.

K should seem familiar. In chapter 2, I argued that in accordance with certain intuitions the size of someone's complaint regarding inequality would depend upon how he fared relative to all those better off than he. I also argued that, in accordance with other intuitions, inequality could be measured by summing up individual complaints. The first intuitions supported the "relative to all those better off" view of complaints, the second, "the additive principle of equality." Together, I argued, these positions constitute a plausible aspect of inequality, one that underlies and influences our egalitarian judgments. But, of course, together these positions support a measure like K, which assesses inequality by adding up the absolute values of the differences between each person and all those better off than he.

We can now see, I think, why the Gini coefficient has the appeal it does. In accordance with certain powerful intuitions, the orderings yielded by a measure such as K will have great plausibility. But, for worlds of the sort we are now interested in, the orderings yielded by the Gini coefficient will be identical to those yielded by K. Thus, insofar as we are caught in the grip of the intuitions in question, the orderings yielded by the Gini coefficient will have great plausibility.

I might note that this way of looking at the Gini coefficient seems to be suggested by some remarks of Sen's. Speaking about income inequality, Sen writes that the following interpretation is suggested by the formula:

$$G = (1/2n^2\mu) \sum_{i=1}^{n} \sum_{j=1}^{n} |y_i - y_j|$$

"In any pair-wise comparison the man with the lower income can be thought to be suffering from some depression on finding his income to be lower. Let this depression be proportional to the difference in income. The sum total of all such depressions in all possible pair-wise comparisons takes us to the Gini coefficient."[7] Talking about welfare, rather than income, I would amend Sen's remarks as follows: in any pair-wise comparison the man with the lower level of welfare can be thought to have a "complaint" regarding inequality. This complaint may seem to be proportional to the difference between his level of welfare and that of the better-off person. How bad a world's inequality is may seem to depend upon its sum total of complaints (taking into consideration all pair-wise comparisons). It is these intuitions the Gini coefficient appears to capture.

5.2 The Statistical Measures in Perspective

I have pointed out important similarities between the statistical measures of inequality and the general theoretical considerations of chapter 2. Specifically, I have shown that, at least for certain sorts of cases, the range expresses the intuitions underlying the relative to the best-off person view of complaints and the maximin principle of equality; the relative mean deviation expresses the intuitions underlying the relative to the average view of complaints and the additive principle of equality; the Gini coefficient expresses the intuitions underlying the relative to all those better-off view of complaints and the additive principle; and the variance, the coefficient of variation, and the standard deviation of the logarithm each express intuitions underlying the relative to the average view of complaints and a weighted additive principle.

We see then that the statistical measures directly reflect four of chapter 2's twelve aspects of inequality. More generally, as diagram 5.4 illustrates, we see that between them the statistical measures reflect each of chapter 2's three ways of measuring individual complaints and each of the three principles of equality. Accordingly, on the view that each way of measuring complaints might plausibly combine with each principle of equality, one might hold that between them the statistical measures *indirectly* support nine of chapter 2's aspects. Only the aspects corresponding to gratuitousness, deviation from the median, and social inequality are not represented by (some combination of the positions underlying) the statistical measures.

The statistical measures were a major step toward elucidating the notion of inequality. Among other things, their plausibility and shortcomings led economists

7. *On Economic Inequality*, p. 33.

	Three Ways of Measuring Individual Complaints			Three Principles of Equality		
Statistical measures	AVE	BOP	ATBO	MP	AP	WAP
R:		X		X		
M:	X				X	
V:	X					X
C:	X					X
H:	X					X
G:			X		X	

AVE = Relative to the Average
BOP = Relative to the Best-Off Person
ATBO = Relative to All Those Better Off
MP = Maximin Principle
AP = Additive Principle
WAP = Weighted Additive Principle

R = The Range
M = The Relative Mean Deviation
V = The Variance
C = The Coefficient of Variation
H = The Standard Deviation of the Log
G = The Gini Coefficient

DIAGRAM 5.4

to appreciate both the importance and difficulty of measuring inequality. This, in turn, led economists to understand the complexity of inequality much better than most. Still, economists did not really pursue the question of *why* the statistical measures have the plausibility they do.[8] As indicated, I think the answer to this question can be found in chapter 2.

It was, I think, unfortunate that the sources of the statistical measures' plausibility were not pursued. I think this for several related reasons. First, given the complexity and confusion surrounding the topic of inequality, there is reason to try to ascertain whether the various measures are indeed grounded on firm intuitive foundations—to be certain, for instance, that they are not *ad hoc*, and that they are truly capturing *egalitarian* intuitions, rather than other moral sentiments, which often, perhaps, accompany such intuitions.

Second, if the source of their plausibility had been pursued, it might have been recognized that the various measures are best regarded as measures not of inequality itself, but of certain *aspects* of that notion. Thus, it might have been seen that *insofar* as these measures seem plausible, it is because they reflect certain powerful intuitions that do indeed underlie our egalitarian judgments, but which are not the only ones to do so. Importantly, this might have prevented numerous wrong turns and fruitless attempts to find the better, more sophisticated, measure of inequality

8. It might be claimed that economists did pursue this question, but that they did so inadequately. My own sense of the literature is that such a charge cannot be sustained. By and large I think the economists' main shortcoming on this point is one of omission rather than commission. It isn't so much that their attempts fail, but that they fail to attempt the task in question.

that would suffice by itself and enable one to simply dispense with the "troublesome" and often conflicting statistical measures. Although, as will be noted, I think our final measure of inequality will ultimately replace the statistical measures with even more plausible measures of the different aspects, one way or another it will have to accurately accommodate the central intuitions that give the statistical measures their plausibility.

Third, if the intuitions underlying the statistical measures had been uncovered, people might have seen that inequality is even *more* complex than had been realized. So, for instance, once people recognized that the range, the gini coefficient, and the variance were supported, respectively, by the intuitions underlying the relative to the best-off person view of complaints and the maximin principle of equality, the relative to all those better off view and the additive principle, and the relative to the average view and the weighted additive principle, they might also have recognized that these intuitions might plausibly combine in other ways to influence our egalitarian judgments. Specifically, they might have recognized that each of the ways of measuring complaints could plausibly be combined with each of the principles of equality.

Fourth, only by discovering the source of their appeal is the seriousness of the conflict between the different measures likely to be appreciated. As long as one focuses on the measures themselves, it is easy to believe that they merely represent different aspects of a complicated notion that happen, on some occasions, to conflict. However, once one explores the foundations of these measures, one sees that some conflict at a very deep level with underlying intuitions that may appear to be fundamentally opposed. This is important for, as noted in chapter 2, it raises serious questions about the scope and intelligibility of the notion of inequality. (These questions, as noted previously, I shall return to in chapter 10.)

Finally, for reasons I discuss in chapters 6 and 7, I think if the sources of the statistical measures' plausibility *had* been pursued, people might have recognized that the statistical measures are best regarded as first approximations, and that even *more* accurate measures of the different aspects might be found.

One point worth emphasizing is the manner in which the argument of section 5.1 supports the general argument of chapter 2. In chapter 2, I presented various considerations showing that the notion of inequality is surprisingly complex. I argued that different intuitions influence our egalitarian judgments, and claimed that even if they are not all *equally* appealing, many really do have the intuitions in question, they have significant plausibility, and they cannot easily be dismissed. Now clearly, many economists have felt about the statistical measures the way I feel about the intuitions presented in chapter 2. In fact, it is arguable that it is precisely *because* it seemed one couldn't simply choose one statistical measure and dismiss the others that many economists first recognized inequality's complexity. The problem was not simply that competing measures were offered, but that competing *plausible* measures were offered. Thus, Meade was not atypical when, in discussing the statistical measures he wrote, "changes . . . [that] have a marked effect in reducing the degree of inequality according to one *very reasonable* measure . . . at the same time . . . have a marked effect in increasing the degree of inequality if some alternative but *equally reasonable* method of measurement is employed" (emphasis added).[9] But then, if the argument in section 5.1 is correct, and there *is* the connection between the statistical measures and the aspects of inequality I have suggested, the fact that many economists have responded to the statistical measures in the manner indicated lends further independent support to the claims of chapter 2—that we *do* have the egalitarian intuitions in question, that they *are* plausible, and that they *cannot* easily be dismissed.

5.3 Atkinson's Measure

Let me next consider Atkinson's measure, one of the most widely accepted measures of inequality.[10] On the surface, at least, Atkinson's approach seems diametrically opposed to this book's. Whereas I want to compare situations *regarding inequality*, independently of how they compare all things considered, or with respect to other moral ideals, Atkinson seems to imply that the value of inequality cannot be understood in such a manner. Let me explore the relations between our positions.

Unlike with the previous measures, it will be most useful to consider Atkinson's measure as it was originally offered, that is, as a measure of *income* inequality. I begin with a detailed presentation of Atkinson's measure, relying heavily on direct quotations from Atkinson. Though lengthy, this presentation will greatly facilitate my discussion.

Atkinson's measure, or index of inequality, is given by the following formula:

$$I = 1 - \left(\sum_{i=1}^{n} \left(\frac{Y_i}{\mu} \right)^{1-\varepsilon} f_i \right)^{\frac{1}{1-\varepsilon}}$$

where "Y_i denotes the income of those in the *i*th income range (*n* ranges altogether), f_i denotes the proportion of the population with incomes in the *i*th range, and . . . [μ] denotes the mean income."[11] As Atkinson observes, the above formula "may look intimidating, but the measure has a very natural interpretation as the proportion of the present total income that would be required to achieve the same level of social welfare as at present if incomes were equally distributed."[12] Let me try to clarify this "natural interpretation."

In his classic article, "On the Measurement of Inequality,"[13] Atkinson assumes that concern for inequality should be expressed in our social welfare function for ranking outcomes. Following Dalton,[14] he suggests that we consider a social wel-

9. J. E. Meade, *The Just Economy* (Allen and Unwin, 1976), p. 113.

10. I am grateful to Peter Mieszkowski and John Broome for their comments regarding the relation between my work and Atkinson's measure. In addition, my presentation of Atkinson's measure is indebted to that of Sen's in *On Economic Inequality*, pp. 38–39.

11. Atkinson, *The Economics of Inequality*, p. 48.

12. Ibid, p. 48.

13. Atkinson, *Journal of Economic Theory* 2 (1970): 244–263.

14. Hugh Dalton, "The Measurement of the Inequality of Incomes," Economic Journal 30, 1920, pp. 348–361.

fare function that "would be an additively separable and symmetric function of individual incomes."[15] On this view,

$$W = \sum_{i=1}^{n} U(y_i) = U(y_1) + U(y_2) + \ldots + U(y_n)$$

where W is the social welfare function, yi is the income of the *i*th person, and U is the individual welfare function (specifically, U gives you an individual's welfare as a function of her income). Because, by hypothesis, U is the same for all individuals,[16] if one assumes, as Atkinson does, that the function U should be increasing and concave—so that the more income one has the less additional income increases one's welfare—then it is easy to see that W will be maximized when income is equally distributed.[17] Thus, the maximum value of W will be $nU(\mu)$, since if income were distributed equally everyone would have the average level of income.

Given these assumptions Atkinson introduces the concept of "the equally distributed equivalent level of income (y_{EDE}) or the level of income per head which if equally distributed would give the same level of social welfare as the present distribution."[18] It is this concept Atkinson's measure intuitively relies on. Specifically, it can be shown that the complicated formula given for Atkinson's measure, I, is equivalent to the following one:

$$I = 1 - \left(\frac{y_{EDE}}{\mu}\right)$$

"or 1 minus the ratio of the equally distributed equivalent level of income to the mean of the actual distribution."[19] As Atkinson observes, "if I falls, then the distribution has become more equal—we would require a higher level of equally distributed income (relative to the mean) to achieve the same level of social welfare as the actual distribution. [In addition] the measure I has . . . the convenient property of lying between 0 (complete equality) and 1 (complete inequality).[20]

Now according to Atkinson, the parameter ε plays a key role in his measure. Specifically, ε "represents the weight attached by society to inequality in the distribution. It ranges from zero, which means that society is indifferent about the distri-

15. "On the Measurement of Inequality," pp. 244–245. On the standard economic model, where social welfare is understood to be a function of individual welfares, basically, the symmetry requirement means that each individual has the same individual welfare function—so that if two people have the same income they have the same welfare—and the requirement of additive separability means that "individual components of social welfare . . . [are] judged without reference to the welfare components of others, and the social welfare components corresponding to different persons are eventually added up to arrive at an aggregate value of social welfare" (Sen, *On Economic Inequality*, p. 40).

16. This is entailed by W's being a symmetric function of individual incomes. See the previous note.

17. If income were not equally distributed one could always increase W's value by transferring income from somebody above the average, say, yj to somebody below the average, say, y_k. U's concavity insures that the decrease in the value of U(yj) would be smaller than the increase in the value of $U(y_k)$, and hence (given W's additive separability) that the value of W will increase.

18. "On the Measurement of Inequality," p. 250.

19. Ibid.; see also Sen's *On Economic Inequality*, p. 38.

20. "On the Measurement of Inequality," p. 250.

bution, to infinity, which means that society is only concerned with the position of the lowest income group. This latter position may be seen as corresponding to that developed by Rawls . . . where inequality is assessed in terms of the position of the least advantaged . . . . Where ε lies between these extremes depends on the importance attached to redistribution at the bottom."[21]

Atkinson offers an illuminating interpretation of ε worth citing in detail:

> Suppose there are two people, one with twice the income of the other (they are otherwise identical), and that we are considering taking 1 unit of income from the richer man and giving a proportion x to the poorer, (the remainder being lost in the process . . .). At what level of x do we cease to regard the redistribution as desirable?
>
> If . . . [one] is at all concerned about inequality, then $x = 1$ is . . . desirable. What is crucial is how far . . . to let x fall below 1 before calling for a stop . . . . The answer determines the implied value of ε from the formula $1/x = 2^{\varepsilon}$. For example, if . . . [one] stops at $x = 1/2$, this corresponds to $\varepsilon = 1$, but if . . . [one] is willing to go until only a quarter is transferred, then the implied ε equals 2.[22]

To the foregoing, Atkinson adds the following important note, "to calibrate ε fully we should have to consider different gaps between the rich and poor man; the [preceding] calculation . . . is only illustrative."

Once one has determined ε's value, calculating y_{EDE} and I is straightforward. Moreover, these values can be conjoined with intuitively appealing interpretations. For example, if the ratio of the equally distributed equivalent level of income to the actual distribution's average is .20, that is, if

$$\left(\sum_{i=1}^{n}\left(\frac{Y_i}{\mu}\right)^{1-e} f_i\right)^{\frac{1}{1-e}} = \left(\frac{y_{EDE}}{\mu}\right) = .20$$

then this means that the gains from redistribution to bring about equality would be equivalent to raising total income by 20 percent; or alternatively, that we could attain the same level of social welfare with only (1—.20 =) 80 percent of the present total income.

Atkinson's measure has many features that make it attractive to economists. These include the following:

1. Atkinson's measure is explicitly normative, as equality's distributional value is explicitly incorporated into the social welfare function via the value attached to ε.
2. Once we have determined ε's value we can assess whether proposals to reduce inequality or increase income are "worth it," by calculating their effect on the social welfare function (W). In particular, for most economists one outcome will be judged better than another if the value of its equally

21. Atkinson, *The Economics of Inequality*, p. 48.
22. Ibid, p. 49.

distributed equivalent income (y_{EDE}) is higher. This implication of Atkinson's measure has great practical appeal.

3. Atkinson's measure has the "virtue" of being neutral as between strikingly disparate distributional objectives. Thus, as noted previously, someone could accept his formal conditions for measuring inequality—that is, for measuring how bad situations are regarding "inequality"—who in fact attaches *no* weight to distributional considerations, and hence who assigns ε the value zero. By the same token, someone who is only concerned with the worst-off's level could accept Atkinson's measure, and assign ε the value infinity.
4. Atkinson's measure is neutral between those who think inequality is intrinsically bad and those who think it is only extrinsically bad. For example, strict utilitarians could assign a value to ε corresponding exactly to the extent to which there is diminishing marginal utility of income, in which case the value of I would reflect *only* the extent to which inequality was *inefficient*. Alternatively, ε's value could incorporate a distributional concern about inequality *itself* beyond the extent to which inequality is inefficient (for the standard utilitarian and economic reasons).
5. For all values of ε less than infinity, Atkinson's measure is compatible with the economic "requirement" of pareto optimality. That is, pareto improvements[23] will always increase the value of the social welfare function W, even if they worsen the situation's inequality.
6. For all values of ε between zero and infinity, Atkinson's measure reflects many of the egalitarian judgments noted previously in this work. For example, it captures the Pigou-Dalton condition in its most plausible form.[24] In addition, it implies that, other things equal, increases in the best-off or decreases in the worst-off will worsen inequality, and that for any transfer of any size or kind it will be best if the transfer goes from the "highest" possible person to the "lowest" possible person.[25]

Given the foregoing, it is not surprising that Atkinson's measure has been widely accepted. But it will also not be surprising that I have many worries about it. My main worry is that Atkinson's measure offers too many too much. While it provides common ground, or a common approach, for a wide variety of distributional concerns, I think it does this in a way that—at least in certain respects—eviscerates its plausibility as a measure of inequality. More cautiously, compatible as it is with utilitarian, Rawlsian, and antiegalitarian views, I believe Atkinson's measure obscures what is *distinctive* about the egalitarian's concern. Equality, I have urged, is an essentially *comparative* notion. Egalitarians believe it is bad (unfair or unjust) for some to be worse off than others through no fault of their own. It is this belief which a measure of inequality should reflect, and although Atkinson's measure is

23. Any increases in the income or welfare of some, which do not decrease the income or welfare of another.

24. Roughly, any (even) transfers from worse- to better-off will worsen inequality, while those from better- to worse-off will improve inequality (at least up to the point where the gaps between better- and worse-off are removed). See sections 3.4–3.6.

25. Again, see sections 3.4–3.6.

compatible with this belief,[26] it is not wedded to it. Thus, as suggested earlier, Atkinson's measure is also compatible with other nonegalitarian and even anti-egalitarian positions.[27]

I worry, then, about the third and fourth "advantages" noted previously. I also worry about the first and second. As the reader knows, I fully agree with Atkinson that one's measure of inequality should be normative, and not merely descriptive. That is why from the outset my concern has been with one situation's inequality being *worse* than another, rather than with one situation's having *more* inequality. But normative concerns can enter into our deliberations about inequality at two levels. First, there may be "purely internal" egalitarian reasons for caring about some inequalities more than others. Second, there may be "external" nonegalitarian reasons for caring about some inequalities more than others, because, for example, trade-offs are necessary between our egalitarian and other moral concerns. I agree with Atkinson that eventually we must address the latter concerns to arrive at a social welfare function reflecting our all-things-considered judgments. But I think we must probably get clear about the former concerns before we can get clear about the latter, not vice versa (see section 5.7).

Similarly, I worry about the simplifying assumption that we merely need to determine the trade-off between income and inequality to determine our social welfare function. Concern about income is but one part of our larger concern about utility, and our concerns about utility and equality are only two of many moral concerns we might have. What about freedom, perfectionism, proportional justice, maximin, and so on? One could care about each of these factors in addition to caring about utility and equality. Does this mean Atkinson's measure needs modification to reflect the relevant trade-offs between all such factors before we could compare situations regarding inequality? If so, Atkinson's measure would look very different than it does now. In addition, it would probably lose much of its appeal and point. Our measure of inequality would have turned into a rather odd way of revealing what factors matter to us and how much they matter relative to each other. It would be good to know such things, but I doubt they must be reflected in our measure of inequality.

I also have a worry about the fifth "advantage." I realize most economists think it conceptually true that if one outcome is pareto superior to another then it must be better all things considered. But for reasons noted in chapter 9 I am not convinced of this. In particular, it is arguable—though controversial—that egalitarians should maintain that in *some* cases pareto improvements worsen a situation all things considered. Consider, for example, the effect of raising one person a certain amount in an otherwise perfectly equal world. It is by no means clear—to me anyway—that the gains in utility and perfection accruing to that one person must *necessarily* outweigh the egalitarian complaints of those left behind. There might be ten billion left behind. Even if each of their complaints would be relatively small, together they might outweigh the gains

26. See subsequent discussion.

27. Many of these assertions have been previously articulated and/or defended, and many will receive further support in subsequent chapters.

of the one person benefited.[28] In sum, I think it is at least an open question whether the egalitarian objection to certain pareto improvements might be sufficiently strong as to condemn those improvements all things considered. Thus, I find it a shortcoming of Atkinson's approach that it appears to link inequality with social welfare in a way that rules this out.

Other objections to Atkinson's measure are possible. As noted already, Atkinson assumes the social welfare function (W) would be an additively separable and symmetric function of individual incomes. Sen raises doubts about the assumption of additive separability,[29] and I have doubts about the symmetry assumption. Suppose A and B both have the same incomes, but A is healthy and B severely handicapped. Why should we believe a given increase in income would improve A's and B's welfare the same amount? For that matter, why believe the individual welfare function U will not only be the same for everybody but be increasing and concave. This assumption assures that income transfers from the richer to poorer will improve inequality on Atkinson's measure. But under certain circumstances additional income might well increase A's welfare by *more* than it would increase B's, even if A were richer than B.[30] Finally, considering the distributive value of a $50 transfer from someone with $100,000 to someone with $10,000 we may judge that ε's value should be 3, whereas considering such a transfer from someone with $50,000 to someone with $30,000 we may judge that ε's value should be .5. Must one or both of these judgments be wrong? On Atkinson's measure we must ultimately determine a single value for ε, which would then be used in weighting each person's individual contribution to social welfare. Yet it may seem that different values for ε would be appropriate in different circumstances.

I think, then, that Atkinson's measure faces numerous problems. Nevertheless, it is widely accepted, and this isn't simply because it provides common ground for a variety of distributional concerns. Rather, Atkinson's measure seems plausible *as a measure of inequality* precisely because, as noted already, for all values of ε between zero and infinity, Atkinson's measure reflects many of the egalitarian's judgments. Let me next briefly suggest why this is so.

In essence, Atkinson's assumptions reflect the following views. First, that the best situation will be the one where income is equally distributed. Second, that deviations from the average worsen inequality. Third, that the total cost to social welfare arising from inequality will be an additive function of the "cost" in terms of individual welfare resulting from the inequality. That is, if we added up the gains in individual utility for each person below the average that would result if the total income in their world was equally distributed, we would know how much the inequality had adversely affected the situation. And finally, that greater deviations

28. At least on many of inequality's aspects. Whether or not one thinks this is so might well depend on the details of the case imagined. One example where I think a pareto improvement would be worse, all things considered, is suggested by Tim Scanlon in "Nozick on Rights, Liberty, and Property," *Philosophy and Public Affairs* 6 (1976): 3–25. The example I have in mind is cited in note 41 of chapter 9, and though it would have to be suitably modified to fit my current purposes, I think this could be done.

29. See Sen, *On Economic Inequality*, pp. 39–41.

30. For example, if A were healthy and rich, and B *severely* handicapped and poor, A might be a more efficient "utility machine" than B, so that A would get more utility from additional income than B.

from the average matter more than lesser deviations,[31] so that the worse off someone is the better (worse) it will be for income to be transferred to (from) him.

The preceding should sound familiar. If we think of each person worse off than the average as having an egalitarian complaint proportional to how much utility they would have gained had income been distributed equally—or alternatively, as having a complaint inversely proportional to the ratio between how much utility they have in the actual distribution to how much they would have if income were distributed equally—then we can see that basically Atkinson's measure reflects those intuitions underlying a relative to the average view of complaints with a weighted additive principle of equality.

Despite its shortcomings, then, it is no accident Atkinson's measure seems plausible even to those who genuinely value equality. This is because for values of ε between zero and infinity Atkinson's measure reflects the basic intuitions that lend plausibility to the variance (V), the coefficient of variation (C), and the standard deviation of the logarithm (H). Moreover, by asking us to explicitly incorporate our distributional concerns into the parameter ε, Atkinson's measure avoids the morally arbitrary squaring feature common to V, C, and H in attaching greater weight to larger deviations below the average. In addition, Atkinson's measure avoids the morally arbitrary square root feature common to C and H, and for values of ε less than infinity it avoids H's problem of becoming so insensitive to transfers at very high levels that it runs afoul of the Pigou-Dalton condition.[32]

In sum, despite the worries expressed, there is much to be said for Atkinson's measure. But as I have tried to indicate, Atkinson's measure is not plausible because of the way it inextricably links inequality to income, utility, and social welfare; rather, it is plausible despite the way it does this. Thus, I think Atkinson's measure does not pose a threat to this book's approach. To the contrary, I think this book helps illuminate both the shortcomings of Atkinson's measure and the source of its appeal.

In conclusion, I think Atkinson's measure, like the statistical measures, is best regarded not as a measure of the whole of inequality but as a measure of one important aspect of inequality. I also think that given the worries expressed here, it is probably best regarded as a useful but flawed method for capturing that aspect.

5.4 The Intersection Approach

In his book *On Economic Inequality*, A. K. Sen introduces an approach for measuring inequality, which I shall refer to as the *intersection approach*. In this section, I shall argue that the intersection approach has serious shortcomings in virtue of which it is not the most plausible method of capturing and reflecting a complex notion like inequality.

In fairness to Sen, let me emphasize at the outset that most of the criticisms of this section are not directed at the intersection approach as Sen himself understood

31. With the extent to which this is so being reflected in our value for ε.

32. Even in its most plausible form.

and advocated it. For the most part Sen's claims on behalf of the intersection approach are both modest and plausible. He merely suggests than an intersection approach may be a fairly useful measure of inequality, one that "opens up a new set of possibilities," that perhaps more adequately corresponds to our notion of inequality than the standard measures of inequality, and that, most important, "helps to sort out the relatively less controversial rankings from those that are more doubtful (p. 74)."[33]

Sen never suggests that an intersection approach will be fully adequate in the sense of being the most complete measure of inequality at which we might arrive. Still, there are two reasons for detailing the problems facing an intersection approach. First, doing so is a useful prophylactic against the temptations of those who would be less careful than Sen, and who might otherwise claim more on behalf of the intersection approach than Sen himself did. Second, and more important, as we shall see in the next section the shortcomings of the intersection approach help illustrate a more promising way of proceeding in order to accurately capture a complex notion like inequality.

Let us begin our discussion of the intersection approach with a bit of background. At the beginning of *On Economic Inequality*, Sen writes the following:

> We may not be able to decide whether one distribution x is more or less unequal than another, but we may be able to compare some other pairs perfectly well. The notion of inequality has many aspects, and a coincidence of them may permit a clear ranking, but when these different aspects conflict an incomplete ranking may emerge. There are reasons to believe that our idea of inequality as a ranking relation may indeed be inherently incomplete. If so, to find a measure of inequality that involves a complete ordering may produce artificial problems, because a measure can hardly be more precise than the concept it represents (pp. 5–6).

Later, after showing that each of the standard measures of inequality faces problems, and noting that each yields a complete ordering, Sen continues:

> Arbitrariness is bound to slip into the process of stretching a partial ranking into a complete ordering. It is arguable that each of these measures leads to some rather absurd results precisely because each of them aims at giving a complete-ordering representation to a concept that is essentially one of partial ranking (pp. 47–48).

So, Sen believes, as we do, that the notion of inequality is complex, multifaceted, and partially incomplete, and that none of the standard measures of inequality will alone suffice to capture that notion. It is in this context that he offers his proposal for how we might arrive at a useful measure of inequality. Sen's suggestion is that we take an intersection relation, Q, of a selected subset of the various measures of inequality that have been (or presumably might be) offered. He claims that "in eschewing exclusive reliance on any one measure and on the complete ordering generated by it, Q restrains the arbitrariness of such measures" (p. 74). On

33. Throughout this section page references are to Sen's *On Economic Inequality* unless noted otherwise.

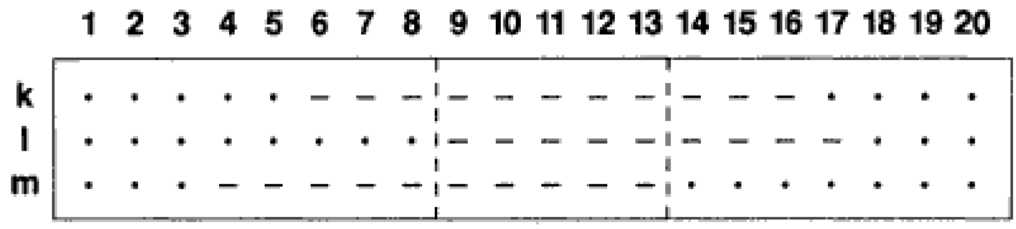

DIAGRAM 5.5

Sen's proposal, Q will yield the judgment that A's inequality is better than B's if and only if A is better than B according to each of the selected measures (and similarly for the relations "worse than," "equivalent to," "at least as good as," and "at least as bad as").

Intuitively, Sen's reasoning might be put as follows. A number of measures of inequality have been offered, each of which has some plausibility and yet each of which, by itself, is inadequate. If, therefore, we are interested in comparing the inequality of two situations, A and B, we would do well to look at more than one of the standard measures. In particular, we might do well to select carefully several of those measures and see what each says about A and B. If each of the measures is in agreement about how A and B compare, then we may confidently conclude that that is how A and B do compare. If, on the other hand, the measures yield opposing judgments, then the intersection approach would yield no judgment regarding A and B.[34]

My first criticism of the intersection approach may be made with the aid of diagram 5.5, where k, l, and m represent possible measures one might employ in arriving at an intersection relation. The marks next to them indicate how each measure corresponds to our intuitions about twenty cross-world comparisons. Each comparison is between two worlds, A_n and B_n. A dash (-) indicates we are certain the measure has yielded a correct judgment about how A_n compares to B_n. A dot (.) indicates this is not the case—either we are certain the measure has yielded an incorrect judgment, or we are not certain whether the judgment is correct or incorrect.

For instance, suppose that in judging between what we may refer to as the two

34. In such a case one might merely note that the intersection approach is silent on the issue of how A and B compare. Alternatively, one might conclude that due to the complex and incomplete nature of inequality, A and B are noncomparable regarding inequality. These two responses would correspond to weaker and stronger versions of the intersection approach. Sen's own view, expressed in correspondence, is the former—more plausible—version, that in such a case the intersection approach is just silent regarding how A and B compare.

Notice, to conclude that A and B are noncomparable regarding inequality would only be to conclude that one cannot say either that one is better than the other or that both are equivalent. That is, it would be to conclude that the notion of inequality does not permit us to make either of those judgments. It would not be to conclude that A and B are noncomparable along the dimension in question, as if one or the other was not the proper kind of object or situation about which judgments of inequality could be made. To the contrary, on the view in question one could make the judgment that neither was better than the other regarding inequality.

11-type worlds, A_{11} and B_{11}, we are certain the inequality is worse in A_{11} than in B_{11}, then the dashes in column 11 of diagram 5.5 imply that according to all three measures A_{11} is worse than B_{11}. Similarly, according to diagram 5.5, k and l both yield the correct judgment about 14-, 15-, and 16-type worlds, while m does not. For the 1-3- and 18-20-type worlds the diagram does not reveal whether the various measures agree with one another. Depending on the cases, the measures may all agree but we may be certain they are wrong, or they may disagree but we may be uncertain about which, if any, is correct.[35]

We are now in a position to see that an intersection relation employing measures k, l, and m will not accurately capture and reflect our notion of inequality. According to diagram 5.5, k, l, and m conflict in their judgments regarding the 4-8- and 14-17-type worlds. This means that an intersection relation employing those measures will yield no judgment about how A_n and B_n compare in those cases. Clearly, however, a fully adequate measure of inequality should yield a judgment in such cases. After all, by hypothesis, not only are we certain in each of those cases that A_n and B_n are comparable regarding inequality, but we are certain how they compare. Specifically, we are certain that l's judgments are correct for the 4-8-type worlds and that k's judgments are correct for the 14-17-type worlds.

The preceding considerations suggest a general and deep problem facing the intersection approach. The problem, whose roots will be explained in greater detail shortly, is that often an intersection approach will fail to yield an ordering when an ordering should be yielded, and this may be so even in cases where it is quite clear what the correct ordering is. Surely, however, a measure that accurately captures and expresses our notion of inequality ought to at least yield those orderings of which we are certain. More generally, just as Sen is correct that a fully adequate measure should not be *more* complete than the notion it expresses, yielding orderings where orderings shouldn't be yielded, so it shouldn't be *less* complete, failing to yield orderings where they should be yielded.

In defense of the intersection approach it might be suggested that one must distinguish between measures that are *strongly misdirecting* and ones that are *mildly misleading*, where a measure, m, will be strongly misdirecting if for at least one comparison between two situations we are certain that m yields the wrong judgment, and mildly misleading if for every instance where we are certain how two situations compare m yields the correct judgment. (The reason for regarding m as mildly misleading is that presumably it will yield judgments in some cases about which we are initially uncertain, and where, after employing the intersection relation Q, we will feel that no judgment can be made.) It might then be claimed that only measures that are mildly misleading may be used in arriving at Q.

Now I readily grant that if one may only employ mildly misleading measures in arriving at Q, the particular criticism just presented will collapse. Q will no longer fail to yield orderings in cases about which we are clear. However, I am most

35. Notice that we may be certain some of the measures are wrong and still be uncertain about whether a conflicting measure is correct. Such a situation could obtain, for example, if we were certain that A_n was not better than B_n, but were uncertain as to whether A_n was worse than, equivalent to, or perhaps noncomparable with B_n. (See the previous note for the sense of noncomparability in question.)

skeptical about the likelihood of our arriving at several intuitively acceptable non-*ad hoc* measures that are only mildly misleading and not strongly misdirecting. In fact, I believe that if we could find even one such measure we would gladly embrace it as our measure of inequality until someone greatly surprised us by revealing another measure that genuinely differed from the first and yet handled all the cases we were sure about equally well. I might add, here, that Sen himself seems to presume that Q will be derived from some of the standard measures of inequality, each of which, there is good reason to believe, is strongly misdirecting and not merely mildly misleading.[36] In short, if the intersection approach is to be regarded as a serious and realistic one it cannot be stipulated that only mildly misleading measures are to be employed in arriving at Q. Hence, there is good reason to think the intersection approach will have the previously noted shortcoming.

It might be claimed that the foregoing objection misses the whole *point* of the intersection approach, that we cannot be "certain" or "sure" of *any* of our egalitarian judgments in advance of, or unless they are yielded by, the intersection relation Q.[37] I believe this claim is deeply mistaken and begs the question against the most fundamental objection to the intersection approach. To see this, it may help to offer a more concrete illustration of why an intersection approach is not the best way of capturing a complex, multifaceted notion.

Consider the relatively trivial notion of basketball ability, shared by (among others) the players, coaches, and fans of basketball. Like inequality, the notion of basketball ability is a complex notion. A number of aspects or factors are relevant to our judgment of whether one person is a better basketball player than another.[38] Among these are the ability to score points, to get rebounds, to play defense, to handle the ball, and to lead the team. For the sake of discussion, let us suppose that the five suggested aspects constitute the notion of basketball ability, and that we are able to arrive at an accurate measure of each of these aspects. It is easy to see that the intersection relation Q derived from these measures would not adequately capture and reflect the notion of basketball ability.

Suppose there are two basketball players, D and F. D is the greatest scorer, rebounder, defender, and team leader in the game's history. F is well below average in all these respects. F is a better ball handler than D, but just barely (F diligently practiced ball handling since early childhood, but his natural talent is so minimal, and D's is so vast...). *Without question*, D would be a better basketball player than F. In fact, D's basketball ability is *vastly* superior to F's. Yet an intersection approach applied to the measures of the aspects of basketball ability would fail to

36. Many of the reasons for believing the standard measures of inequality are strongly misdirecting have already been noted, and were previously noted by Sen and others. Other important reasons for believing the standard measures are strongly misdirecting will be given in chapters 6 and 7.

37. This claim was suggested to me by both Amartya Sen and John Broome. I am not sure if they were actually endorsing it, or merely calling it to my attention as a possible response that might be given.

38. For the purposes of this discussion I am assuming a direct correlation between one person's being a better player than another and that person having greater basketball ability. This does not affect the substance of my remarks and greatly simplifies my presentation. I am aware that it is often said that certain players are better than their abilities would suggest and others are worse than their abilities permit.

yield the judgment that D is better than F. Clearly, then, such an approach would not suffice as a measure of basketball ability.

It is perhaps worth noting that one needn't resort to wild examples to illustrate the point in question. For instance, at the professional level we are frequently able to judge which of two players has more ability, yet it is rare for one professional to be at least as good as another in *all* of the respects relevant to determining our overall judgments. Similarly, among professionals almost every center is a better rebounder than almost every guard and almost every guard is a better ball handler than almost every center, yet in many, and perhaps most, cases we have little difficulty in judging whether a given center has more basketball ability than a given guard. Suffice it to say, such judgments cannot be accounted for on an intersection approach based on the aspects of basketball ability.

The problem with intersections is they don't allow trade-offs. They don't allow trade-offs that take account of the *number* of different measures supporting conflicting judgments, or of the relative *significance* of the different measures supporting conflicting judgments, or of the *degree* to which the different measures support conflicting judgments. So, for example, on an intersection approach, if the measures from which the intersection is derived yield opposing judgments regarding A and B, there is no room for contending that A might be better than B because it is better in four of five important aspects, or because the aspects in which A is better are more significant than those in which B is better, or even, as we have seen, because A is *much* better than B in terms of four aspects and only *slightly* worse in terms of a fifth.

It is no doubt true that sometimes when different facets of a multifaceted notion point in different directions a ranking cannot be expected to emerge. But it by no means follows that *whenever* different facets of a multifaceted notion point in different directions a meaningful ranking cannot emerge. Fortunately, this is as true for the notion of inequality as it is for the notion of basketball ability, for as we have seen inequality is an incredibly complex and multifaceted notion whose facets will diverge in many, and perhaps most, cases.[39]

39. In presenting the intersection approach Sen readily acknowledged that the orderings yielded by the intersection relation Q "might be rather severely incomplete and precisely how incomplete would depend on the extent to which the various . . . measures [from which Q is derived] conflict" (p. 72). The results of chapters 2 and 3 confirm Sen's fears on this score, assuming, as I believe Sen did, that Q was to be derived from the measures of inequality's different aspects.

There are ways of avoiding severe incompleteness in the orderings yielded by the intersection relation Q if one is "careful" in one's selection of measures from which Q is to be derived. For example, instead of including measures of each of the different aspects of a multifaceted notion—as we implicitly did in our basketball examples—one might "conveniently" omit measures of those more "troublesome" aspects that frequently conflict with other aspects of the notion in the judgments they yield. Or, instead of directly focusing on measures of a multifaceted notion's aspects, one might focus on several relatively plausible measures that have been (or presumably might be) offered as measures of the *whole* of the notion.

In such cases the orderings yielded by Q may not be "severely" incomplete, but Q will still face the central problems noted earlier. As long as one of the measures from which Q is derived is strongly misdirecting, there will be orderings Q fails to yield about which we are certain. Moreover, whatever the nature of the measures from which Q is derived, Q will not allow trade-offs of the sort we noted.

There are other, deep, problems with the suggestions in question. Let me mention one for each.

There are purposes for which an intersection approach is well suited. For instance, one could possibly determine the greatest all-around basketball players by appropriately employing an intersection approach. In addition, as we implicitly supposed throughout chapter 3, one might feel confident that any ordering yielded by a carefully selected intersection relation would be uncontroversial.[40] For the reasons noted, however, an intersection approach is not the most plausible way of attempting to capture fully a complex and multifaceted notion.

5.5 A Better Alternative

Let me next offer a suggestion for how best to proceed if one is interested in arriving at a measure that (as) accurately (as possible) captures and reflects a complex and multifaceted notion. To do this, it will be helpful to consider again the

---

Clearly, a measure that simply ignores one or more "troublesome" aspects of a multifaceted notion has no claim to being an accurate measure of the whole of that notion. At best, such a measure might reflect a truncated version of the notion in question, and might yield orderings that "generally" correspond with those of the original notion. Similarly, measures of the whole of a complex notion might be "relatively plausible" as long as the orderings they yielded corresponded "more or less" accurately with those that would be yielded by "most" of the different aspects of the notion "most" of the time. Correspondingly, there is no particular reason to believe that an intersection of such measures (or for that matter any other combination of such measures we might arrive at in a non-*ad hoc* way) would even reflect, let alone accurately capture, each of the complex notion's different aspects giving each its due weight in relation to the others.

40. There are various positions one might take regarding what we may conclude when employing a carefully constructed intersection relation Q. Let us distinguish four: (1) any judgment yielded by Q will be uncontroversial; (2) any judgment that Q fails to yield will be controversial; (3) any judgment yielded by Q will be less controversial than any judgment not yielded by Q; and (4) judgments yielded by Q will "tend" to be less controversial than judgments not yielded by Q.

As noted earlier, Sen's official line is simply to suggest that Q "helps to sort out the relatively less controversial rankings from those that are more doubtful" (p. 74). This cautious remark commits Sen to no more than position 4. However, unofficially, I think (based on conversation) that Sen is inclined to endorse, and even defend, both positions 2 and 3.

There are, I think, several perspectives from which positions 2 and 3 may initially seem plausible. Ultimately, however, my own view is that while 1 may be defensible, 2 and 3 are not. My reasons for rejecting 2 and 3 can be gleaned from our earlier example of the basketball players D and F. There is *no* question in my mind that D, who is the greatest scorer, rebounder, defender, and team leader in the game's history, is a better player than F, who is well below average in all those respects, even though F is a slightly better ball handler than D. To be sure, the extent to which D is better than F would be even greater if D were a better ball handler than F, but the fact that this is not so is not enough to render the judgment that D is a better player than F controversial. Indeed, I contend there would not be a single knowledgeable and unbiased player, coach, or fan of basketball who would be uncertain as to whether D was better than F. Hence, I reject 2.

Similarly, suppose there were a player, E, who was *slightly* better than F in terms of each of the aspects of basketball ability. I grant that the judgment that E would be a better player than F would be uncontroversial, but I reject the claim that it would be less controversial than the judgment that D would be a better player than F. Asked which was the better player among D, E, and F, I contend that knowledgeable and unbiased players, coaches, and fans would be unanimous in selecting D, and that, if anything, there would be less hesitation in choosing between D and F than in choosing between E and F. In addition, I think much more would be required to shake our confidence that D is better than F, than to shake our confidence that E is better than F. Hence, I also reject position 3.

notion of basketball ability, a notion that is more trivial, less widespread, and less controversial than the notion of inequality. The notion of basketball ability is analogous to that of inequality in certain respects, and disanalogous in others.[41] However, for my present purposes, I am mainly interested in the fact that both notions are complex and multifaceted.

Suppose we wanted to arrive at an accurate measure of basketball ability, that is, at a method or procedure for comparing any two people in terms of their basketball ability. How might we best go about doing this? As trivial and obvious as this might seem, I would suggest we should start by getting as clear as possible about what the various facets are constituting the multifaceted notion of basketball ability. Having done this, we should then combine these facets into a single measure in such a way as to accord each its proper weight. In this way, we might arrive at a measure that (as) accurately (as possible) captures and reflects the notion we want. (Note, as we will see in sections 5.6 and 5.7, giving each aspect its "proper" weight may or may not involve giving each aspect a constant weight.)

The idea here might be fleshed out as follows. If we were interested in arriving at a measure of basketball ability, we might begin by asking ourselves the following question: when we judge one player to be better than another, what aspects are relevant to the judgment we make? As suggested in the previous section, our answer to this question will presumably include such factors as the ability to score points, to get rebounds, to play defense, to handle the ball, and to lead the team. Once we were clear on what aspects are involved in our notion, we might then seek to arrive at a measure of each aspect. Each measure would tell us how skilled players were with respect to the aspect in question, and would give us a way of comparing players along that dimension. For instance, we might think we could accurately measure a player's scoring ability by how many points he averaged per game. If so, then someone who averaged eighteen points would have more scoring ability than someone who averaged twelve. Our ultimate aim, of course, would be to arrive at a measure of these measures accurately corresponding to the notion of basketball ability. Essentially, this would involve determining the relative importance of each of the aspects of basketball ability, and then combining the measures of the different aspects accordingly. For instance, if we determined that rebounding ability was twice as important as scoring ability, then in our final measure of the notion of basketball ability someone's rebounding ability, as determined by the appropriate measure for that aspect, would count twice as much as someone's scoring ability, as determined by the appropriate measure for *that* aspect.

The intuitive idea here is quite simple. If we are seriously interested in comparing two people in terms of their basketball ability, then we should compare them in terms of *each* aspect relevant to that ability, and count each aspect in accordance with how important it is to a person's ability to play basketball. If we do this, then we should arrive at the correct judgment given our multifaceted notion of basketball ability.

Now in the case of basketball ability, we *might* be able to come up with a precise and wholly satisfactory measure for each of the aspects composing that

41. In the following section I shall return to the significance, or lack thereof, of the disanalogies between basketball ability and inequality.

notion. Similarly, we *might* be able to come up with a precise and wholly satisfactory weighting of each of the various measures. (For instance, we might decide that each steal is worth so many rebounds, that each rebound is worth so many points, that each point is worth so many assists, and so forth.) If both these conditions held, then presumably we might arrive at a measure of basketball ability yielding a total ordering of all basketball players. Much more likely, however, we might be fairly confident about our methods for measuring certain aspects, such as rebounding[42] or scoring, but less confident about our methods for measuring other aspects—e.g. the ability to handle the ball or to lead the team. In addition, we may have no firm rule for counting each aspect exactly so much and no more in comparison to the others. Thus, for instance, while we may be confident that scoring twenty points is more important than getting three rebounds, we may be unclear whether scoring fifteen points is more important than getting eight rebounds.

The importance of such uncertainty is that, if we are truly faithful to the spirit of the procedure we have presented, our ultimate measure of the notion of basketball ability will undoubtedly be complex, messy, and incomplete. Still, as Sen rightly reminds us, "a measure can hardly be more precise than the concept it represents" (pp. 5–6). Neither our uncertainty in these matters nor a measure accurately mirroring these uncertainties will necessarily reflect a failure to understand or capture accurately the notion of basketball ability; rather, they might be the appropriate and inevitable[43] result of a complex, multifaceted, and incomplete notion.

The point to stress is whether our notion of basketball playing ability is complete or incomplete, if we want to capture that notion accurately we must accurately reflect each of the aspects composing that notion. It is highly unlikely that any method will just *happen* to do this. Rather, to arrive at the sort of method we want we must do some careful thinking. First, we must get clear on what the various aspects *are* constituting the notion of basketball ability. Then, we must arrive at a method accurately reflecting *each* of these aspects. I have suggested this would probably involve a two-step process. Step one would involve arriving at a measure for each of the aspects. This would enable us to compare each player in terms of each aspect. Step two would involve determining the relative importance of the various aspects comprising the notion of basketball ability and combining the values yielded by the measures arrived at in step one accordingly.

Naturally, in both steps our decisions will be guided, though not wholly determined, by our pretheoretical intuitions about who the better basketball players are.[44]

42. "Rebounding" is used here in place of "the ability to rebound." For brevity's sake, analogous substitutions are also sometimes used for basketball ability's other aspects.

43. So long as our thinking is guided by the notion itself.

44. More specifically, in both steps, though probably the second more than the first, we would go through the sort of process John Rawls describes in order to insure that our decisions reflect the "considered" judgments we might make in a state of "reflective equilibrium." (See *A Theory of Justice* [Harvard University Press, 1971], pp. 46–50. Also, see Norman Daniels's "Wide Reflective Equilibrium," *Journal of Philosophy* 5 [1979]: 256–82.) Thus, for instance, our initial judgments about how much to count each aspect will be influenced and altered by our pretheoretical notions about who the better players are, while our judgments about who the better players are will in turn be influenced and altered by our developed views about how much to count each aspect.

Central to the approach I have been advocating is that it will involve precisely the sort of trade-offs precluded by the intersection approach.[45] Thus, before rendering a final judgment as to how two alternatives compare, such an approach will take into account the number and relative significance of the aspects supporting possible judgments as well as the degree to which the different aspects support those judgments.

Incidentally, I believe comparisons between basketball players are often made in roughly the way discussed, as a typical assessment of two players might run something like this: "True, x scores more than y, but rebounding is as important as scoring, and y gets more rebounds than x." "Yes, but each blocked shot is worth about three points, and while x has about three blocks a game, y rarely blocks shots." "Still, one must not underestimate the way y sparks his team, x on the other hand . . ."

If the discussion is a fairly informed and dispassionate one—so the perceptions of the participants are not too distorted by partisan bias—eventually the two players will be compared in terms of each of the aspects of basketball ability. The relative importance of these aspects will be roughly and intuitively calculated, and an overall judgment will be made as to the relative abilities of the two players. Anyway, whether or not something like this actually *does* occur—I think it does—it seems such a process *should* occur both when our concern is with comparing basketball ability, and when our concern is with comparing inequality. More generally, I think such an approach is a reasonable way of attempting to capture and reflect our complex and multifaceted notions.

5.6 Objections and Replies

Some believe I have ignored certain obvious or characteristic disanalogies between basketball ability and inequality in ways rendering the preceding discussion suspect. Let me consider several such objections.

First, one might object that the analogy between basketball ability and inequality is strained because it is clear from the start that basketball ability is a complex of other abilities and that the question is one of weighting, whereas it is much less clear that this is so regarding inequality. I grant the disanalogy in question but fail to see its force. I have not employed the basketball analogy in (question-begging) support of my view that inequality is complex. I have offered arguments throughout this book in support of inequality's complexity, and asked the best way of proceeding in order to capture—that is, measure as accurately as possible—a complex notion. Once it is granted that a notion *is* complex, I think the best way of trying to measure that notion *will* involve a weighting of its aspects along the lines of the basketball example. So for my purposes the analogy is apt, as the relevant question is *whether* basketball ability and inequality are both complex, not the transparency or timing of our recognizing their complexity.

Second, one might object to my analogy because one thinks basketball ability is

45. See the previous section.

an additively separable[46] notion whereas inequality is not. That is, one might think that how good someone is regarding basketball ability is a simple (though weighted) *additive* function of how good that person is regarding each of the separate aspects of basketball ability, where our judgment about how good someone is regarding each aspect is independent of our judgment about how good that person is regarding the other aspects, but that how good a situation is regarding inequality is *not* a simple (though weighted) additive function of how good that situation is regarding inequality's aspects, where our judgment about how good a situation is regarding each aspect is independent of our judgment about how good that situation is regarding the other aspects.

This objection raises several interesting and important issues, but as with the preceding objection I do not see how it undermines my suggested approach for arriving at a measure of inequality. I confess that in my discussion I proceeded as if basketball ability were additively separable in the manner in question. But I only did this for reasons of clarity and simplicity and nothing about my suggested approach hinges on this.

Consider the following view. Though rebounding is an important aspect of basketball ability, and other things equal one would always prefer a better rebounder to a worse rebounder, how much someone's ability to rebound ultimately matters depends on his whole package of basketball skills. For example, great rebounding ability matters more in someone whose other characteristics make him well suited to play center, such as size, lack of speed, and a good inside shot, than in someone whose other characteristics make him well suited to play guard, such as lack of bulk, ballhandling, and a good outside shot. Similar remarks might be made about the other aspects of basketball ability. On this view, how good a basketball player someone is, all things considered, will depend on how good he is regarding the different aspects of basketball ability *and the complex non-additive relation between those aspects.*

In fact, I accept this view. Hence, I do not believe basketball ability is additively separable. But having said that, I would suggest that my basic account of how we might best go about arriving at an accurate measure of basketball ability will be unchanged. It is still the case that we must first get clear on what the various aspects *are* constituting the notion of basketball ability. Then, we must combine these aspects into a measure in such a way as to accord each its proper weight, where this would probably involve a two-step process. Step one would involve arriving at a measure for each aspect that would enable us to compare each player in terms of that aspect. Step two would involve determining the relative importance of the various aspects of basketball ability *in combination with each other* and then weighting and combining the values yielded by step one's measures accordingly.

I do not wish to minimize the difficulties one would confront in trying to construct an accurate measure of basketball ability if that notion is nonadditively separable, any more than I wanted to minimize the difficulties in constructing such a measure on the alternative assumption. My point is simply that my suggestion for

46. The notion of additive separability was introduced in section 5.3. See note 15, and the surrounding text.

how best to proceed if one wants to arrive at a measure of a complex notion does not depend on whether the notion with which one is concerned is additively separable. Correspondingly, the claimed disanalogy between basketball ability and inequality regarding additive separability—which may or may not obtain—does not affect my position.

Importantly, then, my claims about seeking a measure that will accord each of inequality's aspects its "proper" or "due" weight should *not* be construed as implying that each of inequality's aspects must be accorded a *constant* weight. Perhaps certain aspects will be more or less significant in combination with others—for example, if a situation is already very good in terms of aspect A, it may matter less that it (also) be good in terms of aspect B. Or perhaps promoting A will initially be more important than promoting B, but once a situation is "sufficiently" good regarding A further incremental gains will matter less and it will then be more important to promote B than A. In other words, whether additive or not, nothing in my claims precludes the trade-off function between aspects from being as complicated as necessary.

Third, one might object to the analogy between basketball ability and inequality by arguing that they are entirely different kinds of notions embedded within entirely different contexts. For instance, one might contend that basketball is a specific game with precise rules and well-defined aims in virtue of which one could expect basketball ability's aspects and their relative importance to be identifiable—say, in terms of their instrumental efficacy toward achieving the aims of basketball. The case is otherwise, one might contend, regarding morality in general and inequality in particular. There is no specific code or practice of morality or inequality with precise rules and well defined aims in virtue of which one could expect inequality's aspects and their relative importance to be identifiable. Correspondingly, although my suggested approach may seem perfectly plausible for arriving at a meaningful measure of basketball ability, there is no reason to expect my approach to meet with success in arriving at a meaningful measure of inequality.

This objection is significant, and some variation of it might be forwarded by two important, though very different, traditions in moral philosophy. Subjectivists deny the objectivity of value. Hence, they would deny the existence of any truths or facts about the world that could guide our reasoning about inequality's aspects in such a way as to result in an (objectively) meaningful measure of inequality. Aristotelians, on the other hand, are objectivists about ethics, but they deny that the subject matter of ethics permits precision. On an Aristotelian view the good is what the good man does. The good is recognized by the person of practical wisdom and is a matter of perception or judgment; it is decidedly *not* something that can be codified in an accurate system of rules or measures we might then consult in the course of moral deliberation.[47] (Presumably the subject matter of basketball might be far more precise and permit the sort of meaningful measure unavailable for ethics.)

This is not the place for me to respond to the deep worries of subjectivists or Aristotelians. Instead let me observe that their worries are not merely about my

47. See Aristotle's *Nicomachean Ethics*, for example, book 1, chapters 3 and 7, and book 2, chapters 2, 6, and 9.

particular approach, they are worries about *any* approach for trying to arrive at a meaningful measure of inequality.

My conclusion at the end of the preceding section was modest. All I claimed is that my suggested approach is "a reasonable way of attempting to capture and reflect our complex and multifaceted notions." This conclusion stands despite subjectivist and Aristotelian worries. After all, I did not claim my approach would necessarily *succeed*; rather, I offered it as "a suggestion for how best to proceed if one is interested in arriving at a measure which (as) accurately (as possible) captures and reflects a complex and multifaceted notion." Though here as elsewhere in philosophy there are no guarantees, I believe it is a suggestion worth considering.

If the subjectivists or Aristotelians are right, perhaps a meaningful measure of inequality is not to be had. Still, there are many respects in which it would be useful to have such a measure. Hence, unless we are already convinced they *are* right—and I take it many are not—there is reason to pursue my suggested approach.

In sum, although there are disanalogies between basketball ability and inequality, they do not impugn my earlier claims.

5.7 Methodology, Additive Separability, and the Hermeneutic Tradition

The preceding remarks about additive separability have some bearing on my methodology. In addition, I suspect they may have some bearing on the nature of the conflict between the analytic and hermeneutic traditions. Let me conclude this chapter by commenting on the former and speculating (no doubt foolishly!) about the latter.

As noted in the Introduction, some people are suspicious of my attempt to understand the notion of inequality independently of other ideals and the particular contexts with which it is normally associated. Although I suspect most who share this view are simply not egalitarians, at least not in my sense of the word, some might object to my methodology on general hermeneutic grounds. They will insist that trying to understand inequality, or any other ideal, out of the particular historical context in which it is embedded is to lose sight of the point of the philosophical enterprise, as any concept elucidated in such a manner will be unconnected with the real world and hence be lacking in significant grounding or relevance. "What" they might ask "does the inequality of my abstract sterile diagrams have to do with inequality as it is experienced in the world and as it influences the course of human events?" On this view my methodology commits the analytic philosopher's classic mistake of seeking to understand the whole by analyzing its parts, while failing to recognize that the parts themselves cannot be understood independently of the whole.

It is not my intention, here, to defend analytic philosophy or to compare the relative strengths and weaknesses of the analytic and hermeneutic traditions. But I do want to suggest that my methodology is compatible with certain important insights of both traditions. I believe there are many cases where one cannot *fully* understand the parts independently of the whole. On the other hand, as the hermeneutic tradition itself insists, neither can one fully understand the whole independently of the parts. This, of course, is the famous—or infamous—hermeneutic circle. And it raises the troubling question, "if to know the whole one must know the part, and to know the part one must know the whole, how is knowledge of anything possible?" How are we to enter (or exit?) the circle to establish a starting place or perspective (not to say foundation) from which to develop our understanding of both part and whole?

Not surprisingly, here my sympathies lie with the analytic approach. It is not as if the whole is composed of something *other* than its parts. Hence, if we want to understand the whole, it makes perfect sense to *begin* by trying to identify and understand its parts. And in doing this it does not seem unreasonable to *start* by carefully considering each part separately. But, as the hermeneutic tradition suggests, we must recognize that any understanding of the nature and significance of the different parts that we reach in this manner will only be tentative with respect to our ultimate goal, which is to understand the whole by understanding each part *and its nature and role in relation to the others.*

My view here parallels my earlier discussion regarding the best way of proceeding if one has a complex notion that is nonadditively separable. Suppose we wanted to know whether one situation was better than another all things considered, and we couldn't tell just by "looking" at them. (Alas, our judgment failed us and those we might [otherwise] have thought possessed practical wisdom were unable to help.) We might throw up our hands in despair, or perhaps decide that there couldn't "really" be any significant difference between them. Alternatively, we might proceed first by trying to get clear on what aspects *are* relevant to our all-things-considered judgments, and then by trying to combine these aspects into a measure in such a way as to accord each its due weight. As before, I think this might involve a two-step process. Step one would involve arriving at a measure for each aspect that would enable us to compare each outcome in terms of that aspect. Step two would involve determining the relative importance of the various aspects *in combination with each other* and then weighting and combining the values yielded by step one's measures accordingly—where, of course, this recognizes that according each aspect its "due weight" *may* require a complex, nonadditive, function involving variable weights.

Naturally, I favor the last approach, and see this book and its methodology as important for carrying out step one. As before, I grant that the described process may not succeed. Perhaps in seeking a method or measure—even a rough and incomplete one—for comparing situations all things considered, one is seeking yet another philosopher's stone. Still even if such an effort ultimately falls short of its aim, I believe the attempt to identify and clarify our different moral ideals reveals much about our moral concepts and the role they play in our moral thinking. I have found this to be so regarding Mill's contributions to our understanding of utility, Rawls's contributions to our understanding of justice, and my own modest contributions to our understanding of equality.

The foregoing suggests one way of interpreting a central insight of the hermeneutic tradition. Perhaps the insight that the part cannot be understood independently of the whole reflects the similar insight of the Gestalt tradition that the whole may be greater or less than the *sum* of its parts where the basic point is *not* that the whole is somehow composed of something *other* than its parts—it obviously isn't!—but rather that parts often combine in *nonadditive* ways in constituting the wholes of which they are parts. Thus, to note a nonphilosophical case, giving someone a more beautiful nose may or may not result in his having a more beautiful face, depending on the nature of the relation between a person's features and the beauty of his face. (Perhaps the more beautiful nose will be too big, or too small, for the size of the person's eyes or mouth.) Similarly, to note a moral case, increasing a situation's utility may or may not improve it—even a little—depending on the nature of the relation between utility and other ideals. So, for example, on a Kantian view of proportional justice the moral value of increased utility will depend on whether the beneficiary of the increased utility deserves her good fortune. (Recall Kant's dictum that "the sight of a being who is not graced by any touch of a pure and good will but who yet enjoys an uninterrupted prosperity can never delight a rational and impartial spectator. Thus a good will seems to constitute the indispensable condition of being even worthy of happiness.")[48]

Understood as an insight about additive separability, I think the hermeneutic claim that the part cannot be understood independently of the whole is interesting, important, and too often neglected. I have discussed additivity elsewhere[49] and shall not explore it here; but briefly, I think many positions in philosophy—especially analytic philosophy—implicitly assume additive separability where it is unclear such assumptions are defensible. Correspondingly, I think the hermeneutic tradition offers valuable criticism insofar as it calls such assumptions into question. On the other hand, I also think it is no accident additive assumptions are so pervasive. While such assumptions raise many worries, and are no doubt unwarranted in *some* cases, there are many cases where it is far from clear how we might think about morality or rationality in the absence of such assumptions.

In any event, for the reasons given here and in the preceding section I think accepting the nonadditive separability of our moral ideals does *not* undermine my methodological approach. To be sure, we may learn further insights about the nature and role of inequality when we later consider its relation to other ideals. Accordingly, we should regard the results of our present inquiry as preliminary. Still, the process of understanding must begin somewhere, and what better place to start than to try to identify and clarify as carefully as possible the different parts that compose the whole and whose relation to each other we ultimately seek to understand? This book begins the task for the ideal of equality.

48. Grounding for the Metaphysics of Morals, first section, first paragraph (James W. Ellington's translation).

49. Larry S. Temkin, "Additivity Problems," in *Encyclopedia of Ethics*, ed. by Lawrence Becker and Charlotte Becker (Garland Publishing, 1992) pp. 15–18. For a rich discussion of additivity, see Shelly Kagan's "The Additive Fallacy," *Ethics* 99 (1988): 5–31.

I conclude that my methodology is compatible with the hermeneutic insight that one cannot fully understand the part independently of the whole, at least on one plausible interpretation of that insight. More particularly, I conclude that my methodology is not objectionable as long as we do not naively assume that morally ranking outcomes will be a simple (perhaps weighted) additive function of separate values and ideals.