
Economic Inequality

10.1. Introduction

In this chapter, we begin to study a theme that recurs throughout the book: the analysis of the *distribution* of income or wealth among different groups in society. Economic growth that spreads its benefits equitably among the population is always welcome; growth that is distributed unequally needs to be evaluated not simply on the basis of aggregate or per-capita change, but on the grounds of equity.

There are two reasons to be interested in the inequality of income and wealth distribution. First, there are good ethical grounds to display an intrinsic aversion to inequality. There is no acceptable moral reason why individuals should be treated differently in terms of their access to lifetime economic resources.¹³⁰ It is, of course, possible to argue that people make choices — good and bad — over the course of their lifetimes for which only they are responsible: that they are poor because they made irresponsible decisions, and “had it coming to them.” In some cases this may indeed be true, but in most cases the unequal treatment of individuals begins from day one. Two children from two different families on an unequal footing simply because of parental wealth differences, and for this fact there is little ethical defense. To hold descendants responsible for the sins of their ancestors is excessive. At the same time, we run into a separate ethical dilemma. To counteract the unequal treatment of individuals from the first day of their lives, we must deprive parents of the right to bequeath their well-deserved (or perhaps ill-gotten) wealth to their children. There may be no way to resolve this dilemma in a satisfactory manner.

Nevertheless, we can work toward a society with tolerable levels of inequality in everyday life. This more pragmatic objective softens the dilemma in the preceding paragraph, by reasonably curbing the scope for drastically unequal levels of accumulation, though of course it cannot entirely eliminate the problem of an uneven playing field. But in any case, we cannot speak philosophically about “development” without a serious consideration of the problem of inequality.

¹³⁰I make this statement assuming that there is no fundamental difference, such as the presence of a disability or ailment, in the need for two people to have access to economic resources.

The second reason for studying inequality is that even if we are uninterested in inequality at an intrinsic level, there is still good reason to think about it for *functional* reasons. about it. Suppose you simply care about overall growth and not at all about inequality, but find that inequality in income and wealth somehow reduce the possibilities of that overall growth (or increase it; at this stage the direction of change is unimportant). Then you must care about inequality at what might be called an *instrumental level*; to you, inequality will be important not for its own sake, but because it has an impact on other economic features that you do care about.

In this book, we will pay attention to both the intrinsic and the instrumental features of inequality. That is, we view the reduction of inequality both as a development goal in itself, and as a strategic tool in that it impacts other development objectives: good health, the reduction of poverty, undernutrition and mortality, the right to gainful employment, and indeed even per-capita income. To address these twin goals — the intrinsic and the functional — we must first learn how to think about inequality at a conceptual level. This has to do with how we *measure* inequality, which is the subject of this chapter. In Chapter 11, we will examine, both at an empirical and theoretical level, how inequality interacts with other economic variables, such as per-capita income and the growth of that income.

10.2. What is Economic Inequality?

10.2.1. Economic Disparities in Context. At the level of philosophy, the notion of inequality can dissolve into an endless sequence of semantic issues. Ultimately, economic inequality is the fundamental disparity that permits one individual certain material choices, while denying another individual those very same choices. From this basic starting point begins a tree with many branches. João and José might both earn the same amount of money, but João may be physically handicapped while José isn't. John is richer than James, but John lives in a country that denies him many freedoms, such as the right to vote or travel freely. Shyamali earned more than Sahar did until they were both forty; thereafter Sahar did. These simple examples suggest the obvious: economic inequality is a slippery concept and is intimately linked to concepts such as lifetimes, personal capabilities and political freedoms.¹³¹

Nevertheless, this is no reason to throw up our hands and say that *no* meaningful comparisons are possible. Disparities in personal income and wealth, narrow though they may be in relation to the broader issues of freedom and capabilities, mean *something*. It is in this spirit that we study income and wealth inequalities: not because they stand for *all* differences, but because they represent an important component of those differences. But even within the economic context, there are caveats and qualifications to be kept in mind. Here are a couple of them.

10.2.2. Current and Lifetime Inequalities. Depending on the particular context, we may be interested in the distribution of current expenditure or income *flows*, the distribution of wealth (or asset *stocks*), or even the distribution of lifetime income. You can see right away that our preoccupation with these three possibilities leads us progressively from *short-term* to *long-term* considerations. Current income tells us about inequality at any one point of time, but such inequalities may be relatively

¹³¹On these and related matters, read the insightful discussions in Sen (1985).

harmless, both from an ethical point of view and from the point of view of their effects on the economic system, provided the inequality is temporary. As an example, imagine two societies. In the first, there are two levels of income: \$2,000 per month and \$3,000 per month. In the second society, there are also two levels of income, but they are more dispersed: \$1,000 per month and \$4,000 per month. Let us suppose that the first society is completely immobile: people enter their working life at one of the two levels of income but stay there forever. In the second society, people exchange jobs every month, switching between the low-paid job and the high-paid job. These societies are obviously unrealistic caricatures, but they make the point: the first society shows up as more equal if income is measured at any one point in time, yet in terms of average yearly income, everyone earns the same in the second society.

Thus our notions of cross-sectional inequality at any one point in time must be tempered by a consideration of *mobility*. Whether each job category is “sticky” or “fluid” has implications for the true distribution of income. Often, because of the lack of data, we are unable to make these observations as carefully as we would like, but that does not mean that we should be unaware of them.

10.2.3. Functional and Personal Distributions. It may also be of interest to know (and we will get into this later in the book) not only *how much* people earn, but *how* it is earned. This is the distinction between *functional* and *personal* income distribution. Functional distribution tells us about the returns to different factors of production, such as labor (of different skills), capital equipment of various kinds, land, and so on. But as you can easily see, this is just half the story. It is also important to describe how these different factors of production are owned by the individuals in society.

Figure 10.1 illustrates this process. Reading from left to right, the first set of arrows describes how income is generated from the production process. It is generated in varied forms. Production involves labor, for which wages are paid. It involves the use of land or capital equipment, for which rents are paid. It generates profits, which are paid out to owners and shareholders. Production also involves payments for various nonlabor inputs of production. These other inputs are in turn produced, so that in the ultimate analysis, all incomes that are generated can be classified under payments to labor of different skills, rents to non-produced inputs such as land, and profits to ownership. The distribution of income under these various categories is the *functional distribution of income*.




Figure 10.1. Functional and personal distribution of income.

The second set of arrows tells us how different categories of income are funneled to households. The direction and magnitude of these flows depend on who owns which factors of production (and how much of these factors). Households with only labor to offer (household 3 in the diagram, for instance) receive only wage income. In contrast, households that own shares in business, possess land to rent, and labor to supply (such as household 2) receive payments from all three sources. By combining the

functional distribution of income with the distribution of factor ownership, we arrive at the *personal* distribution of income—a description of income flows to individuals or households, not factors of production.

You might well ask: why should we care about this two-step process? Isn't knowledge of the personal distribution good enough for our analysis? The answer is that it isn't, and there are at least two reasons for this. First, the functional distribution tells us much about the relationship between inequality and other features of development, such as growth. Our understanding of how economic inequalities are created in a society necessitates that we understand both how factors are *paid* and how factors are *owned*. Second, the understanding of income sources may well influence how we judge the outcome. Money from charity or the welfare state may be viewed differently from the same amount received as income for work. Amartya Sen, in a closely related context, refers to this as the problem of “recognition” or self-esteem (see Sen 1975):

“Employment can be a factor in self-esteem and indeed in esteem by others. If a person is forced by unemployment to take a job that he thinks is not appropriate for him, or not commensurate with his training, he may continue to feel unfulfilled and indeed may not even regard himself as employed.”

Although there may not be much that we can do about this (so far as measurement goes), we should keep it in the back of our minds while we proceed to a final judgment about inequality.¹³²

10.2.4. The Road Ahead. The preceding discussion lays down a road map for our study of inequality. We look at economic inequalities from two angles. In this chapter, we put all sources of income into a black box and concentrate on the evaluation of income (or wealth or lifetime income) distributions. This part of the story is *normative*. All of us might like to see (other things being the same) an egalitarian society, but “egalitarian” is only a word. How exactly do we evaluate alternative income distributions? How do we rank or order these distributions? This part of the chapter discusses how we measure inequality, or equivalently, how we rank alternative distributions with respect to how much inequality they embody.

With measurement issues out of the way, we proceed in Chapter 11 to a study of the economics of income distributions: how inequality evolves in society, the effects that it has on other features of economic development, such as output, employment and growth rates, and how these other features feed back in turn on income and wealth distributions. This part of the story is *positive*. Whether or not we like the notion of egalitarianism per se, inequality affects other features of development.

10.3. Measuring Economic Inequality

10.3.1. Introduction. If there is a great deal of disparity in the incomes of people in a society, the signs of such economic inequality are often quite visible. We probably know a society is very unequal when we see it. If two people are supposed to share a cake and one person has all of it, that's unequal. If they split 50–50, that's equal. We can

¹³²Often, ingenious theories of measurement can go some way to resolve difficulties of this sort. For instance, it might matter for our measurement of literacy rate whether an literate person has access to *other* literate persons. On these matters, see Basu and Foster (1997).

even evaluate intermediate divisions (such as 30–70 and 40–60) with a fair amount of precision. All that goes away, however, once we have more than two individuals and we try to rank intermediate divisions of the cake. Is it obvious how to compare a 20–30–50 division among three people with a 22–22–56 division? In more complicated situations such as these, it might be useful to develop some yardsticks for comparison which rely on basic normative principles.

What are the properties or principles that a “desirable” inequality index should satisfy? It is difficult to have complete unanimity on the subject. If, to avoid controversy, we lay down only very weak criteria, then many inequality indices can be suggested, each consistent with the criteria, but probably giving very different results when used in actual inequality comparisons. If, on the other hand, we impose stricter criteria, then we sharply reduce the number of admissible indices, but the criteria loses wide approval. Sometimes there is debate even about the very basics.

So transparency is best. When you trust a measure of inequality, you are identifying your intuitive notions of inequality with that particular measure. If you are a policy maker or an advisor, such identification can be useful or dangerous, depending on how well you understand the underlying measurement criteria.

10.3.2. Four Criteria for Inequality Measurement. Suppose that society is composed of N individuals. An *income distribution* is a description of how much income is received by each individual i in that society.¹³³ We are interested in comparing the relative “inequality” of two income distributions. To this end, we will formalize some intuitive notions about inequality in the form of applicable criteria.

(1) *Anonymity Principle.* From a social perspective, it does not matter *who* is earning the income. A situation where Debraj earns x and Rajiv earns y should be viewed as identical (from the point of view of inequality) to one in which Debraj earns y and Rajiv earns x . Debraj may well be unhappy with the switch (if x happens to be larger than y), but it would be hard for him to argue that overall inequality in his society has deteriorated because of the change. Permutations of incomes among people should not matter for inequality judgments: this is the principle of *anonymity*.

(2) *Population Principle.* Cloning the entire population (with their incomes) should not alter inequality. More formally, if we compare an income distribution over N people and another population of $2N$ people with the same income pattern repeated, there should be no difference in inequality among the two income distributions.¹³⁴ The population principle thus asserts that all that matters are the *proportions* of the population that earn different levels of income, and not absolute numbers.

(3) *Relative Income Principle.* Just as population shares matter and the absolute values of the population itself do not, one can argue that only *relative* incomes should matter and the absolute levels of these incomes should not. If one income distribution is obtained from another by scaling *everybody's* income up or down by the same

¹³³In this section, we refer to “income” as the crucial variable whose inequality we wish to measure. You could replace this with wealth, lifetime income, and so on. Likewise, the recipient unit is called an individual. You could replace this by “household” or any other grouping that you might be interested in.

¹³⁴Warning: Cloning only one segment of the population while keeping the remainder unaltered may well inequality. Suppose that there are two incomes, \$100 and \$1,000. The population principle says that all income distributions are equally unequal provided the same percentage of people earn \$100. If the *proportion* of people earning the low income changes, then inequality will, in general, be affected.

percentage, then inequality should be no different across the two distributions. For instance, an income distribution over two people of (\$1,000, \$2,000) has the same inequality as (\$2,000, \$4,000). This is the relative income principle: absolute incomes have no meaning as far as *inequality measurement* is concerned.¹³⁵ Of course, absolute incomes continue to be centrally important in our overall assessment of development.

(4) *The Transfers Principle*. We are now in a position to state our final criterion for evaluating inequality. Formulated by Dalton (1920),¹³⁶ the criterion is fundamental to the construction of measures of inequality. An income transfer from one person to another will be called *regressive* if the money is going from a relatively poor person to a relatively rich person. The transfers principle states that if one income distribution can be converted into another by a sequence of regressive transfers, then the latter distribution must be deemed more unequal than the former.

A common misinterpretation of the transfers principle is that it actually requires those transfers to take place. That's not true. It applies as long as we (the researcher) can view two income distributions *as if* they are connected by such a chain of regressive transfers, whether or not those transfers have actually happened.

To summarize, an income distribution is described as $(y^1, y^2, \dots, y^N)$, where N is the total population. The anonymity principle tells us that we don't need to keep track of individual names. The population principle tells us that we can equivalently describe the income distribution by $(y_1, \dots, y_m; n_1, \dots, n_m)$, where we are given m income *classes* $\{y_j\}$, along with the population *shares* $\{n_j\}$ peopling each class. That is, $n_1 + n_2 + \dots + n_m = 1$.¹³⁷ (In actual practice, we don't know anyone's incomes perfectly, so an income class y_j typically stands for a *range* of incomes centered at y_j ; for example, "\$100 per month or less," "\$300–400," and so on.)

An *inequality measure* I maps income distributions to a number, which we call the "inequality" of that distribution. A higher value of the measure signifies the presence of greater inequality. Thus an inequality measure can be interpreted as a function of the form $I(y^1, \dots, y^N)$, or $I(y_1, \dots, y_m; n_1, \dots, n_m)$, depending on which way of writing an income distribution is more convenient. The relative income principle tells us that

$$I(y^1, \dots, y^N) = I(\lambda y^1, \dots, \lambda y^N)$$

for all $\lambda > 0$. Finally, the transfers principle requires that

$$I(y^1, \dots, y^N) < I(y^1, \dots, y^{i-1}, y^i - \delta, y^{i+1}, \dots, y^{j-1}, y^j + \delta, y^{j+1}, \dots, y^N),$$

whenever $y^i \leq y^j$ and $\delta > 0$. As we shall see, some of these "first principles" or axioms can indeed be challenged. But in my opinion, they have enough going for them that it is worth seeing exactly where they take us.

¹³⁵Is it as easy to buy the relative income principle as the population principle? Not really. What we are after, in some sense, is the inequality of "happiness" or utility, however that may be measured. As matters stand, our presumption that inequality can be quantified at all forces us to make the assertion that the utilities of different individuals can be compared. (On the analytical framework of interpersonal comparability that is required to make systematic egalitarian judgments, see, for example, Sen 1970 and Roberts 1980.) However, the relative income principle needs more than that. It asserts that utilities are proportional to incomes. This is a strong assumption. We make it nevertheless because Chapter 13 will partially make amends by studying the effects of absolute income shortfalls below some poverty line.

¹³⁶See also Pigou (1912). The criterion is often called the Pigou-Dalton principle.

¹³⁷Superscripts like k denote the income y^k of a person k , and subscripts like j stand for income class y_j .

10.3.3. Maximal Inequality. Here's a little application of the concepts we've developed so far. Think of societies with different populations: $N = 2, 3, 4, \dots$. Let's try to conceive of *maximal* inequality within and across societies. The Relative Income Principle tells us that we may as well think of 100 units of income to divide among the N people in each society, while the Anonymity Principle tells us two societies each with the same number of people can be treated as identical from the viewpoint of inequality measurement.¹³⁸ Now consider, say, a two-person society, and suppose that one of the two has all the income. I think we would easily agree that inequality is maximal across all other possible two-person allocations (such as a 70-30 division). That agreement is in line with the Dalton Principle, because a regressive transfer can take us from any allocation of the form $(x, 100 - x)$ to either 100-0 or 0-100.

But here's a trickier question. Compare a two-person society with a *four*-person society. Suppose that in each society, just one person has all the income. Which allocation is more unequal now: 100-0 or 100-0-0-0? One might argue — and there is a viewpoint among scholars that does argue — that the two allocations are *equally unequal*.¹³⁹ After all, just one person has all the income in either society. But it is worth noting that our axioms prohibit this conclusion. To see this, note by the Population Principle that 100-0 is just as unequal, no more and no less, than 100-100-0-0. However, by the Transfers Principle, 200-0-0-0 is *strictly* more unequal than 100-0-0-0. The former allocation in turn has the same inequality as 100-0-0-0, by the Relative Income Principle. So our axioms insist that 100-0-0-0 is *more unequal* than 100-0. It is in this sense that the axioms help us discriminate between different allocations. With allocations that do not involve just one person getting the entire cake, the comparisons are more subtle still, and we need additional tools to make headway.

10.3.4. The Lorenz Curve. There is a useful pictorial way to see what the four criteria of the previous section give us. The *Lorenz curve* nicely summarizes the distribution of income in any society. Because it is very often used in economic research and discussion, it is worthwhile to invest a little time in order to understand it.

Figure 10.2 shows a typical Lorenz curve. On the horizontal axis, we depict *cumulative* percentages of the population arranged in increasing order of income. Thus points on that axis refer to the poorest 1/4 of the population, the poorest half of the population, and so on. On the vertical axis, we measure the percentage of national income accruing to any particular fraction of the population thus arranged. The first marked point, for example, corresponds to a value of 1/3 on the population axis and 1/10 on the income axis. The interpretation of this is that the poorest 1/3 of the population earns only 1/10 of overall income. The second point corresponds to 2/3

¹³⁸Clearly, either Principle can be objected to. At 100 units of income everybody might starve no matter how that income is divided, so who cares? We set this objection aside by noting that this is a question of *poverty*, not inequality. Nevertheless, the argument could be continued by saying that we should be interested in the inequality of *utility*, not money, and utility may well not be linear in money. That takes us on to slippery philosophical ground, as it necessitates the use of cardinal utility comparisons across people. Of course, inequality measurement is not free of those concerns, using — as it implicitly does — a linear utility function $u(y) = y$. In similar fashion, the Anonymity Principle could be criticized on the grounds that two individuals may have different needs, owing to a disability or handicap or other differential circumstances. We set this one aside as well by noting that we could redefine “exchange rates” across differently-abled individuals by stating that each dollar accruing to person *A* is the same as, say, \$1.10 to person *B*. That said, notice how even the most basic principles can — and should — be questioned.

¹³⁹See, for instance, the mini-debate across the authors in Sethi et al. (2023).

on the population axis and 1/4 on the income axis. This point, therefore contains the information that the “poorest” 2/3 possess 1/4 of the national income.

Notice that the Lorenz curve must begin and end on the grey diagonal: after all, the poorest 0% earn 0% of total income by definition and the poorest 100% must earn 100%. Indeed, if everybody had the same income, the curve would then coincide *everywhere* with that diagonal or the 45° line of the box. The poorest $x\%$ (however selected, and for any x) would then have exactly $x\%$ of national income. Because the 45° line expresses the relationship $Y = X$, it is our Lorenz curve in this case, and is marked “perfect equality” in the diagram. With increasing inequality, the Lorenz curve starts to fall below the diagonal in a convex loop bowed out to the right of the diagram; it cannot curve the other way. After all, the slope of the Lorenz curve at any point is simply the contribution of the person at that point to the cumulative share of national income. Because we have ordered individuals from poorest to richest, this “marginal contribution” cannot ever fall, and so the Lorenz curve must be convex.

The “overall gap” between the diagonal and the Lorenz curve is indicative of the amount of inequality present in the society. The greater the extent of inequality, the further the Lorenz curve will be from the 45° line. Hence, even without writing down any formula for the measurement of inequality, we can obtain an intuitive idea of how much inequality there is by simply inspecting the Lorenz curve.

For instance, Figure 10.3 shows the Lorenz curves of two different income distributions, marked $L(1)$ and $L(2)$. Because the second curve $L(2)$ lies entirely *below* the first one, it is natural to expect a good index to indicate greater inequality in the second case. Let’s try to understand why this is so. If we choose a poorest $x\%$ of the population (it does not matter what x you have in mind), then $L(1)$ always has this poorest $x\%$ earning at least as much as they do under $L(2)$. Thus regardless of which precise value of x you pick, the curve $L(1)$ is always “biased” toward the poorest $x\%$ of the population, relative to $L(2)$. It stands to reason that $L(1)$ should be judged more equal than $L(2)$.

This eyeballing can be elevated to the status of a formal rule: the *Lorenz criterion*. It says that if the Lorenz curve of a distribution lies everywhere below the Lorenz curve of some other distribution, the former should be judged to be more unequal




Figure 10.2. The Lorenz curve.




Figure 10.3. The Lorenz criterion.

than the latter. Just as we required an inequality measure to be consistent with the criteria of the previous section, we might require it to be consistent with this particular criterion. Thus an inequality measure I is *Lorenz-consistent* if, for every pair of income distributions $(y^1, \dots, y^N)$ and $(z^1, \dots, z^k)$ (over n and k individuals respectively),

$$I(y^1, \dots, y^N) \geq I(z^1, \dots, z^k)$$

whenever the Lorenz curve derived from $(y^1, \dots, y^N)$ lies everywhere below its counterpart derived from $(z^1, \dots, z^k)$.

This is all very well, but now confusion starts to set in. We just spent an entire section discussing four reasonable criteria for inequality comparisons and now we have introduced a fifth! Are these all independent restrictions that we have to observe? Fortunately, there is a neat connection between the four criteria of the previous section and the Lorenz criterion that we just introduced: *an inequality measure is consistent with the Lorenz criterion if and only if it is simultaneously consistent with the anonymity, population, relative income, and transfer principles.*

This observation is very useful.¹⁴⁰ It captures our four ethical criteria in one clean picture that gives us exactly their joint content. It is so central to our conceptual understanding of inequality that it is worth taking a minute to see why it is true.

First, observe that the Lorenz curve automatically incorporates the principles of anonymity, population, and relative income, because the curve drops all information on income or population *magnitudes* and retains only information about income and population *shares*. To see that the Lorenz curve also embodies the transfers principle, consider this example: suppose there are just three persons with incomes (10, 20, 30). How might we compare this to a six-person society with incomes (14, 14, 44, 44, 62, 62)? Anonymity just tells us that individual identities don't matter. The population principle tells us that the latter society has the same inequality as a three-person society with incomes (14, 44, 62). The relative income principle tells us that we can divide all incomes by 2 to obtain yet another three-person society with incomes (7, 22, 31). We now complete the comparison by returning to the initial society. We can hypothetically transform it into the final society by transferring 2 units of income from the poorest person to the middle person, and another unit from the poorest person to the richest person. These are regressive transfers, and the transfers principle says that inequality goes up as a result. In summary, (14, 14, 44, 44, 62, 62) is more unequal than (10, 20, 30).




Figure 10.4. The transfers principle and the Lorenz criterion

Figure 10.4 tells us what happens to the Lorenz curve. The thicker curve marks the original Lorenz curve and the thinner curve shows us the Lorenz curve after the transfer. Under the old curve, the first third of the society (namely, the poorest person) has 1/6

¹⁴⁰For a useful discussion of the history of this result, see the survey by Foster (1985).

of the income, and the cumulative second third (namely, the first two combined) have $1/2$ of the income. What about the new curve? Because resources have been transferred away, *both* these cumulative shares have fallen. The new Lorenz curve therefore dips below and to the right of the old Lorenz curve. This state of affairs persists until the 100th percentile comes along, at which point the cumulative shares return to 100%. The new Lorenz curve lies below the old, and so the Lorenz criterion also tells us that inequality has gone up, thus agreeing with the principles we introduced above.¹⁴¹

Happily, the converse comparison is true as well: if two Lorenz curves are comparable according to the Lorenz criterion, as in the case of $L(1)$ and $L(2)$ in Figure 10.3, then it *must* be possible to construct a set of regressive transfers leading from $L(1)$ to $L(2)$. We leave the details to an exercise at the end of this chapter.

At this point, it looks like we are all set. We have a set of axioms at our disposal. We also have the Lorenz curve that can be displayed easily on a diagram. And — best of all — the axiomatic approach and the diagrammatic approach are one and the same! Whenever a pair of distributions can be compared according to the axiomatic criteria, the Lorenz curve criterion gives the same answer, and vice versa. So it appears that we are ready to close the book on inequality measurement. Alas, there is a catch. All of this presumes that the pair of income distributions are in fact comparable under the axioms (or the Lorenz criterion). That's not always the case: two Lorenz curves can *cross*.




Figure 10.5. Ambiguous comparisons: A Lorenz crossing.

Figure 10.5 illustrates. There are two income distributions, represented by the Lorenz curves $L(1)$ and $L(2)$. Observe that neither Lorenz curve is uniformly above the other, so that the Lorenz criterion does not apply to this case. By the equivalence result discussed previously, it follows that the transfers principles cannot apply either, but what does it mean for that principle “not to apply”? It means that we *cannot* get from one distribution to the other by a sequence of regressive transfers.

The following example illustrates this point. Try and compare two income distributions within a four-person society, given by $(75, 125, 200, 600)$ and $(25, 175, 400, 400)$. Observe that we can “travel” from the first distribution to the second in two steps. First transfer 50 from the first person to the second: this is a regressive transfer. Next transfer 200 from the fourth person to the third: this is a progressive transfer. We have arrived at the second distribution. Of course, these transfers are just a construction, and we can try other routes. For instance, transfer 50 from the first person to the third: this is regressive. Transfer now 150 from the fourth to the third: this is progressive. Finally, transfer 50 from the fourth to the second: this is also progressive. Again, we

¹⁴¹It should be noted that the new curve does not have to lie below the old curve *everywhere*, though that's the case in our example. For instance, if a regressive transfer only occurs between the 40th and 80th percentiles of the population, the new Lorenz curve will coincide with the old all the way up to the 40th percentiles (nothing has happened so far) and also above the 80th percentile (no cumulative shares have been altered there either), only falling below the old Lorenz curve between the 40th and the 80th percentiles.

arrive at the second distribution, this time in three steps. And indeed, there are many ways to “travel” from the first to the second distribution, but the point is that they *all* necessarily involve at least one regressive *and* at least one progressive transfer. (Try it.) In other words, the four principles of the previous section are just not enough to permit a comparison.

Therefore, in the sort of cases typified by this example, we *have* to somehow weigh in our minds the “cost” of the regressive transfer(s) against the “benefit” of the progressive transfer(s). These trade-offs are generally impossible to quantify in a way so that everybody will agree.¹⁴² In our specific example, the poorest 25% of the population earn 7.5% of the income in the first distribution and only 2.5% of the income in the second. The opposite comparison holds when we get to the poorest 75% of the population, who enjoy only 40% of the total income under the first distribution, but 60% of the income under the second distribution.

Now go back and look at Figure 10.5 once again. You can see that $L(1)$ and $L(2)$ are precisely the Lorenz curves for the two distributions in this example.

Databases for Inequality Measurement

Unlike per-capita income, which needs no more than a national aggregate divided by a headcount, inequality asks for the whole *distribution*. That means household surveys: expensive, infrequent, and awkward at both ends of the income range. Rich households are the least likely to answer and the most likely to understate, and the poorest are the hardest to find and track. To add insult to injury, some countries measure income while others measure consumption, some report gross figures and others report income net of taxes and transfers, and some adjust for household size while others don't. The resulting numbers are serviceable for broad comparisons and for trends, but let's agree not to trust them down to the decimal point. With that warning in place, here are the main sources.

1. The World Income Inequality Database (WIID), from UNU-WIDER, <https://www.wider.unu.edu/database/world-income-inequality-database-wiid>. This is the closest thing we have to a census of inequality estimates: around 200 countries and historical entities, close to 4,000 unique country-year observations, running through 2022 in the version that's current as I write. (The database advertises over 26,000 “data points,” but that is a count of rows rather than of facts: the same country and year recurs once for every source and every income concept, so the smaller number is really the fairer measure of coverage.) It reports decile and quintile shares along with Gini coefficients, Theil and Atkinson indices, and Palma ratios, and it deliberately does *not* force everything into a single mold: each observation is tagged with the income concept it came from. The WIID Companion datasets do the harmonizing — a curated series of per-capita net income at the percentile level, one file for countries and one for the world as a whole, so that a Lorenz curve can be drawn straight from the data. A cautionary tale, incidentally, about the whole enterprise: the 2026 update was pulled back not long after release when errors turned up in it, and the site reverted to the 2025 version.

2. The Standardized World Income Inequality Database (SWIID), maintained by Frederick Solt, <https://fsolt.org/swiid/>. Where the WIID gives you everything

¹⁴²For instance, Shorrocks and Foster (1987) consider “transfer sensitivity,” a principle that compares progressive transfers at the lower end of the income distribution with regressive transfers at the upper end, arguing that if both involve the same sized transfer, then inequality should come down under the composite transfer. Such additional criteria are generally all open to debate.

and leaves you to sort it out, the SWIID makes comparability the entire point. It benchmarks other sources against the Luxembourg Income Study and reports Gini coefficients for both market and disposable income — the gap between the two being a serviceable measure of how much redistribution a government actually accomplishes — for close to 200 countries from 1960 onwards. Coverage that complete is often bought with imputation, and to its credit the SWIID is candid about the implicit price: every number arrives with a confidence interval attached. Use the numbers with those intervals in mind. The method is described in Solt (2020).

3. The Luxembourg Income Study (LIS), <https://www.lisdatacenter.org>. LIS does not hand you inequality measures. Instead, it hands you harmonized household micro-data — some fifty-odd countries and over a thousand datasets spanning five decades — and you compute whatever measure you please. A parallel wealth database (LWS) covers around 25 countries. This is the gold standard for comparability, which is exactly why the SWIID uses it as its benchmark, but the LIS coverage leans towards the richer countries. High quality, but understandably restrictive.

4. The Poverty and Inequality Platform (PIP), from the World Bank, <https://pip.worldbank.org>. This replaced PovcalNet, retired in 2022. Built on the Bank's archive of household surveys, it reports Ginis, decile shares and mean log deviations alongside the poverty counts. Two things are worth knowing. Since June 2025 everything is expressed in 2021 PPP, with international poverty lines of \$3.00, \$4.20 and \$8.30 per person per day, so older numbers (\$1.90 a day at 2011 PPP, say) are not comparable to current ones. And the Bank itself warns that the machinery was built to replicate poverty counts, not to trace out entire distributions: the estimated densities near the very top and the very bottom should not be leaned on.

5. The World Inequality Database (WID), <https://wid.world>. This grows out of the work of Atkinson, Piketty, Saez and Zucman on the reconstruction of top income shares from historical tax records (Piketty 2003, Piketty and Saez 2003), and is now maintained by the World Inequality Lab. Its distinguishing feature is a refusal to rely on surveys alone: tax records, national accounts and surveys are spliced together so that the pieces of the distribution add up to national income — “distributional national accounts.” It reports the income and wealth shares of the top 1%, the top 10%, the middle 40% and the bottom 50%, before and after taxes and transfers, with series for a few countries reaching back into the nineteenth century. Precisely because it corrects for the top incomes that surveys miss, WID inequality *runs higher* than survey-based inequality; the two are not to be spliced or casually compared. The readable summary of it all is the World Inequality Report, most recently Chancel et al. (2026).

A warning: the WID's historical reach is largely an illusion outside the rich countries. Where a country has no usable source data, it is handed its region's series, with those same numbers repeated country by country: Ethiopia, Kenya and Nigeria carry *identical* top-decile shares in every year before 1980, as do Bangladesh, Pakistan, the Philippines and Vietnam — and Bangladesh is furnished with an income distribution for 1820, a century and a half before the country existed. Long historical series for poor countries are not to be trusted. WID does flag them, in that every observation carries a data-quality score. Those flags should be applied before anything else is done with the data.

Nor does the problem stop in 1980. The series is *reported* annually, but the underlying sources are updated only in the years a country actually fielded a survey or published tax tabulations, and the years in between are filled in by interpolation. Inference using intertemporal kinks in the reported decile shares suggests that Nigeria's apparently annual series are based on exactly five survey rounds, and Brazil's carries none at all before 1999,

while France and the United States genuinely are updated every year. Movements between updates are an artifact of the interpolation process.

There are other compendia worth a look. The OECD maintains income and wealth distribution databases for its member countries. Branko Milanovic's *All the Ginis* (Milanovic 2026) assembles Ginis from the major primary sources into a single long panel. The *Chartbook of Economic Inequality* (Atkinson et al. 2017) puts together century-long series for 25 countries. And for the United States, the Census Bureau's Current Population Survey and the Internal Revenue Service's Statistics of Income are the two workhorses, with usefully complementary blind spots: the survey top-codes high incomes, while the tax data simply omit those too poor to file a return.

As a final observation on comparability: Ceriani and Verme (2026) collect 122,351 Gini estimates covering 222 countries and territories over the years 1867–2024, and find that Ginis built on income exceed Ginis built on consumption by 4.7 points on average. That is a large number — often larger than the cross-country or over-time difference one is trying to detect in the first place. So before setting any two inequality numbers side by side, ask what they are measures of.

Despite these ambiguities, Lorenz curves provide a clean, visual image of the overall distribution of income in a country. Figure 10.6 provides several examples of Lorenz curves for different countries, arranged so that the curves bow further away from the diagonal as we read across and down.¹⁴³ By looking at these figures, you can get a sense of income inequalities in different parts of the world, and with a little mental superimposition you can compare inequalities across any two countries.

10.3.5. Complete Measures of Inequality. As we've seen, our four criteria (and the Lorenz curve) do well in comparing income distributions, but they are necessarily silent on some comparisons. At the same time, researchers, policy makers and governments are often interested in summarizing inequality by a single number; which implies a complete ranking over all pairs of distributions. As we will see, that completeness does not come free of charge: it means that in some situations, inequality measures will disagree with one another. We now survey some commonly used inequality measures, but will keep this caveat in mind.¹⁴⁴

Our income distributions of the form $(y_1, \dots, y_m; n_1, \dots, n_m)$, where the y_i 's are income classes and the n_i 's are population shares. The mean of such a distribution is given by the average of incomes weighted by population shares:

$$\mu \equiv \sum_{j=1}^m n_j y_j.$$

We sometimes invoke the equivalent description $(y^1, y^2, \dots, y^N)$, where the y 's are now a full listing of everyone's income. This will sometimes be useful in examining the properties of the measures.

(1) *Kuznets Ratios.* Simon Kuznets introduced these ratios in his pioneering study of income distributions in developed and developing countries. These ratios refer to

¹⁴³Countries whose surveys measure consumption rather than income — India, Indonesia, Egypt, Nigeria and Thailand among them — are deliberately left out, because consumption is distributed more equally than income and the two cannot be set side by side: see the box on databases earlier in this chapter.

¹⁴⁴See Sen (1997) for a discussion of these and other measures, and for a comprehensive overall treatment of the subject of economic inequality.




Figure 10.6. Lorenz curves for nine countries, ordered by increasing inequality. The figures along each curve are cumulative percentage shares of income at each decile. *Source:* World Bank Poverty and Inequality Platform (PIP), most recent survey available for each country as of 2026, spanning 2021–2024. All nine estimates are income-based rather than consumption-based, and the income concept is household *disposable* income — that is, income measured after taxes and transfers.

the share of income owned by the poorest 20 or 40% of the population, or by the richest 10%, or sometimes to the *ratio* of the shares of income of the richest $x\%$ to the poorest $z\%$, where x and z stand for numbers such as 10, 20, or 40. The Kuznets ratios are essentially “pieces” of the Lorenz curve and serve as a useful shorthand in situations where detailed income distribution data are missing.

(2) *The Mean Absolute Deviation.* This is the first of several measures that take advantage of the entire income distribution. The idea is simple: take all income distances from the mean, add them up, and divide by the mean to express average deviation as a fraction of total income. That yields the *mean absolute deviation*:

$$M = \frac{1}{\mu} \sum_{i=1}^m n_i |y_i - \mu| \tag{10.1}$$

where the notation $|\cdot|$ stands for absolute differences (neglecting negative signs). But M has one serious drawback: it is often insensitive to the transfers principle. We can

equivalently write M as

$$M = \frac{1}{N\mu} \sum_{i=j}^N |y^j - \mu|,$$

where $(y^1, \dots, y^N)$ is the full listing of incomes. Let $y^j \leq y^k$, and consider a regressive income transfer from person j to k . If these incomes are on opposite sides of the mean, then both $|y^j - \mu|$ and $|y^k - \mu|$ go up (and all other income distances are unchanged), so M must rise. But the transfers principle is meant to apply to *all* regressive transfers. Suppose that both y_j and y_k are above the mean. Now, if the transfer is small enough so that both incomes stay above the mean after the transfer, there will be no change in the *sum* of the absolute differences from mean income. The mean absolute deviation therefore fails the transfers principle. That isn't good news for the measure.

(3) *The Coefficient of Variation.* One way to avoid the insensitivity of the mean absolute deviation is by giving more weight to larger deviations from the mean. A familiar statistical measure that does just this is the standard deviation, which squares all deviations from the mean. Because the square of a number rises more than proportionately to the number itself, this is effectively the same as attaching greater weight to larger deviations from the mean. The coefficient of variation (C) is just the standard deviation divided by the mean, so that only relative incomes matter. Thus

$$C = \frac{1}{\mu} \sqrt{\sum_{i=1}^m n_i (y_i - \mu)^2}. \quad (10.2)$$

The measure C , it turns out, has satisfactory properties. It satisfies all four principles and so it is Lorenz-consistent. In particular, it always satisfies the transfers principle. In terms of the full listing of incomes $(y^1, \dots, y^N)$, equation (26.7) is the same as

$$C = \frac{1}{\mu} \sqrt{\frac{1}{N} \sum_{j=1}^N (y^j - \mu)^2}.$$

Consider a transfer from j to k , where $y_j \leq y_k$. If y^j and y^k are on opposite sides of the mean, it is immediate that C rises. If both y^j and y^k are on the same side of the mean before and after the transfer, say, both greater than μ , we have a transfer from a smaller number (i.e., $y^j - \mu$) to a larger one (i.e., $y^k - \mu$), which increases the *square* of the larger number by more than it decreases the square of the smaller number. The net effect is that C invariably registers an increase when such a regressive transfer is made. So C is in agreement with all our axioms, but it is not the only one that is.

(4) *The Gini Coefficient.* This measure starts from a base that is fundamentally different from the others. Instead of taking deviations from the mean, it takes the absolute difference between *all* pairs of incomes and adds them all up. It is as if inequality is the sum of all conceivable pairwise comparisons of "two-person inequalities." In symbols, the Gini coefficient G is given by

$$G = \frac{1}{2\mu} \sum_{i=1}^m \sum_{j=1}^m n_i n_j |y_i - y_j|. \quad (10.3)$$

The double summation sign signifies that we sum all *pairs* of income differences (weighted by the share of such pairs, $n_i n_k$). Because each $|y_j - y_k|$ is counted twice (again as $|y_k - y_j|$), the whole expression is divided by 2 as well as by mean income.

The Gini coefficient has pleasing properties. It satisfies all four principles and is therefore Lorenz-consistent, just like the coefficient of variation. Again, we need only check the transfers principle, the others being obvious. Figure 10.7 depicts the listing of incomes from lowest to highest. Take two incomes, say y^j and y^k , with $y^j \leq y^k$, and transfer some amount δ from y^j to y^k , small enough so that $y_j - \delta$ remains at least as rich as anyone with strictly lower income than y_j , and $y_k + \delta$ is no richer as anyone with strictly higher income than y_k .¹⁴⁵ To figure out how the Gini changes, all we need to do is consider pairs in which j or k figure. Consider incomes no larger than $y^j - \delta$. The difference between these incomes and that of j has narrowed by δ . This narrowing is exactly counterbalanced by the fact that y^k has gone up by the same amount, so the distance between k 's income and those incomes no larger than $y^j - \delta$ has gone up by an equal amount. An exactly parallel argument holds for pairs that involve incomes no poorer than $y^k + \delta$. That leaves us with incomes between $y^j - \delta$ and $y^k + \delta$. Every pairwise distance between these incomes and those of j and k has gone up. So has the distance between the incomes of j and k . All told, then, the Gini must increase. This argument shows why the Gini coefficient is Lorenz-consistent.




Figure 10.7. The Lorenz consistency of the Gini coefficient.

There is another interesting property of the Gini coefficient that ties it closely to the Lorenz curve. Recall that the more “bowed out” the Lorenz curve, the higher is our intuitive perception of inequality. It turns out that the Gini coefficient is precisely the ratio of the *area* between the Lorenz curve and the 45° line of perfect equality, to the *area* of the triangle below the 45° line. Why this is so is left to the problem set.

We have thus surveyed four indexes. The first is a crude but nevertheless useful indicator of inequality when detailed data are unavailable. The second should not be used. Finally, both the coefficient of variation (C) and the Gini coefficient (G) appear to be perfectly satisfactory indexes, going by our four principles (or what is equivalent, Lorenz consistency),¹⁴⁶ but this gives rise to a puzzle. If both C and G are satisfactory in this sense, why use both measures? Why not just one?

This brings us back full circle to Lorenz crossings. We have just seen that both C and G are Lorenz-consistent. This means that when Lorenz curves *can* be compared, both C and G give us exactly the same ranking, because they both agree with the Lorenz criterion. The problem arises when two Lorenz curves cross. In that case, it is possible for the Gini coefficient and the coefficient of variation to give contradictory rankings. This is nothing but a reflection of the fact that our intuitive sense of inequality

¹⁴⁵The argument for larger δ follows by breaking up the transfer into smaller pieces satisfying the condition above, and applying the same logic.

¹⁴⁶Of course, other measures are used as well. There is the use of *log variance* as an inequality measure, which is just the standard deviation of the logarithm of incomes. Although it is easy to compute and use — and is used — the log variance disagrees with the transfers principle in some cases, and I would not recommend it. Another measure, introduced to inequality evaluation by Henri Theil and known as the *Theil index*, is derived from entropy theory. Although it looks bizarre at first, it turns out to be the *only* measure that satisfies the four principles and a convenient decomposability principle that permits us to separate overall inequality into between-group and within-group components (Foster 1983). This makes the Theil index uniquely useful in situations where we want to decompose inequality into various categories, for example, inequality within and across ethnic, religious, caste, occupational, or geographical lines.

is essentially incomplete. In such situations, we should probably not rely entirely on one particular measure of inequality, but rely on a whole set of measures. It may be a good idea to simply study the two Lorenz curves as well.

As a hypothetical example, consider two societies, each consisting of only three persons. Let the distribution of income in the two societies be (3, 12, 12) and (4, 9, 14), respectively. You can easily check that for the first of our hypothetical societies, the coefficient of variation is 0.47, whereas it is 0.45 for the second. Using *C* as an index, therefore, we reach the conclusion that the first society is more unequal than the second. However, if we calculate the Gini coefficient, the values come out to be 0.22 and 0.25, respectively. On the basis of the latter measure, therefore, inequality seems to be higher in the second society compared to the first.¹⁴⁷

To be sure, this isn't just a hypothetical possibility. Such contradictory movements of our inequality indices occur in real life as well. (Fields, 1980, Chapter xx) contains a fascinating discussion of these and related matters.

Clearly we have a dilemma here: the result of our comparison is sensitive to the choice of the index, but we have no clear intuitive reason to prefer one over the other. There are two ways out of this dilemma. The first, as we said before, is to examine our *notion* of inequality more closely and to come up with stricter criteria after such introspection. The result will likely be subjective and controversial. The second escape is to realize that human thought and ideas abound with *incomplete* orderings: everyone agrees that Shakespeare is a greater writer than the Saturday columnist in the local newspaper; however, you and I might disagree whether he is greater than Tagore or Tolstoy, and even I may not be very sure *myself*. Relative inequality, like relative literary strength, may be perfectly discernible some of the time and difficult to judge in other cases. We can learn to live with that. If a society manages to significantly increase economic fairness and humane distribution among its members, then this fact will be captured in every reasonable inequality index, and we will not have to quibble over technicalities! It pays, however, to be aware of the difficulties of measurement.

10.4. Beyond Economic Inequality

As the struggle proceeds, the whole society breaks up more and more into two hostile camps, two great, directly antagonistic classes: bourgeoisie and proletariat. The classes *polarize*, so that they become internally more homogeneous and more and more sharply distinguished from one another in wealth and power. (Deutsch 1971, p. 44)

The theory of economic inequality that we've discussed so far is based on a long and venerable tradition in welfare economics. The transfers principle lies at the heart of this theory, arguing that every regressive transfer must serve to increase inequality, no matter what specific measure we use. Yet there is something above the transfers principle that fails to capture the global picture in Morton Deutsch's observation, quoted above. Deutsch talks about the formation of two great income classes, each homogeneous, yet sharply distinct from each other. If we looked at the formation of each of the classes *separately* instead, we would see "mini-equalizations" of income

¹⁴⁷Warning: There is no connection between a value of, say, 0.25 achieved by the Gini coefficient compared to the same number achieved by *C*. That's like comparing apples and oranges. All this example is doing is contrasting different *movements* of these indexes as the distribution of income changes.

in each class. The transfers principle would celebrate these changes. Yet the overall situation may well be increasing in its “polarization” across classes. It is in this sense that the transfers principle is “local” in nature: it passes a welfare judgment on a transfer independently of other changes in the system. This may be good welfare economics and a good basis for studying inequality, but may not capture some of the important distributional changes that can create social conflict.

10.4.1. The Transfers Principle Revisited. Consider an example. Suppose that the population is uniformly distributed over ten values of income, spaced apart equally at the values (1000, 2000, . . . , 10000). Now we imitate Deutsch: we collapse this distribution into a two-spike configuration, with half the population at 3000, and the other half at 8000. See Figure 10.8. Which distribution exhibits greater “polarization”? The answer should be quite clear. Two groups are now perfectly well formed in the second situation, while in the former a sense of group identity is more fuzzy. Under the second distribution, population is either “rich” or “poor,” with no “middle class” bridging the gap between the two, and one may be inclined to perceive this situation as more antagonistic than the initial one. There is a greater sense of Us and Them.




Figure 10.8. Inequality and polarization.

As another interpretation, think of the above not as incomes but as an opinion index over a given political issue, running from “left” to “right.” Many would agree that political conflict is more likely under a two-spike distribution — with perhaps not completely extreme political opinions, but sharply defined and involving population groups of significant size — rather than under the uniform dispersion of opinions.

But the point is this. If you do admit the possibility that the second distribution is more polarized (or more conflictual or tense), you are forced to depart from the domain of inequality measurement. For under *any* Lorenz-consistent inequality measure, inequality has come *down* in moving from the first distribution to the second.

10.4.2. Polarization. This is not so much an indictment of inequality measurement as a statement that there is more to distributions than just inequality. Just as the mean of an income distribution is central to economic growth, but deliberately kept aside in the study of inequality (recall the relative income principle), there are other aspects of a distribution that inequality may not fully capture. One of them is polarization.

That’s not to say that polarization is somehow the opposite of inequality. Sometimes the two don’t go together, as in the example above, but sometimes they might. For instance, start with the same distribution and collapse the entire distribution into a single population spike at 5000. Now inequality has come down (as it did before), but so, presumably, has polarization. Polarization isn’t the negation of inequality. It’s just a different concept.¹⁴⁸

¹⁴⁸For a fuller development, see Esteban and Ray (1994, 2011a), Duclos, Esteban, and Ray (2004) and Foster and Wolfson (2010).

Just as in the case of inequality, it is possible to write down criteria to evaluate polarization. Anonymity, the population principle, and the relative income principle do not need to be discarded. The same considerations apply to these criteria as they do for inequality measurement. The main departure comes by discarding the transfers principle and replacing it by criteria that respond to clustering or bunching of the income distribution into groups. Esteban and Ray (1994) develop such criteria by referring to three essential features of a polarized situation: society exhibits groupings which (1) are each significantly sized, (2) feature a high degree of homogeneity *within* each group, and (3) exhibit a high degree of heterogeneity *across* groups. They propose a class of measures, of which the central one is

$$P = \frac{1}{\mu} \sum_{i=1}^m \sum_{j=1}^m n_i^2 n_j |y_i - y_j|. \quad (10.4)$$

This looks very much like the Gini coefficient, but there is one important difference. The measure P adds up terms of the form $n_i^2 n_j |y_i - y_j|$ instead of the terms $n_i n_j |y_i - y_j|$ that appear in the Gini. To understand this, note that inequality, as expressed by the Gini, is a measure of the “overall alienation” in the system. Each person in each income class i feels an alienation $|y_i - y_j| n_j$ from income class j : they have different incomes and of course we need to weight by the share of the population having those incomes, which is n_j . Adding these alienations across everyone (and normalizing by the mean income), we get the Gini.

Polarization also adds up alienations, but factors in a sense of “group identification” as well. Person i ’s alienation with respect to income class j is still $|y_i - y_j| n_j$, but this antagonism is additionally bolstered by a sense of identification with her own group, which is proportional to the share of her own compatriots n_i . The “socially effective alienation” she feels is therefore given by $|y_i - y_j| n_i n_j$. Adding these alienations across everyone (and normalizing by the mean income), we get the polarization index P .¹⁴⁹

10.4.3. Social Groupings. We’ve applied the idea of polarization to income or wealth, but in principle it could be over any other social feature. We hinted at political opinion above, but it could be over any groupings, based on geography, kin, ethnicity, religion ... and income, of course, but let’s go beyond income for now.

Suppose, then, that there are m “social groups,” with population shares $(n_1, \dots, n_m)$, just as before. For each pair of groups i and j we can think of an inter-group “distance,” call it $d(ij)$. Then a natural transplant of (10.4) yields the measure

$$P = \sum_{i=1}^m \sum_{j=1}^m n_i^2 n_j d(ij). \quad (10.5)$$

In this case of income, $d(ij)$ was just the (mean-normalized) income distance $|y_i - y_j|/\mu$, but we can replace it by other notions of distance, depending on what we know about the history and origins of the groups in question. For instance, in Chapter 26 we will use linguistic differences across ethnic groups, but sometimes, matters can be far simpler. Given the groupings at hand, individuals may be interested only in the dichotomous perception Us/Them. Perhaps the simplest instance of this is what might

¹⁴⁹In particular, n_i appears again as a consequence of the addition, which accounts for the term n_i^2 . You should note that this is not an asymmetric treatment of group j , because their term n_j will be correspondingly squared when we add in *their* “effective alienation.”

be called a “pure contest” situation, in which the “distance” across two groups is the same (say 1) irrespective of group. It is then easy to see that

$$P = \sum_{i=1}^m \sum_{j=1}^m n_i^2 n_j = \sum_{i=1}^m n_i^2 (1 - n_i), \quad (10.6)$$

a polarization measure first introduced by Esteban and Ray (1999) and applied to studies of social conflict by Montalvo and Reynal-Querol (20xx).

The parallel with the Gini isn’t over yet. Notice that the Gini could just as easily be applied to social groupings, and would correspondingly take the form

$$\sum_{i=1}^m \sum_{j=1}^m n_i n_j d(ij).$$

In the special case of a pure contest, this formula takes on a particularly well-known form, known as the *fractionalization index*. First introduced in the Soviet *Atlas Narodov Mira* which classified and catalogued ethnolinguistic groups around the world, this index is given by the formula

$$F = \sum_{i=1}^m n_i (1 - n_i) \quad (10.7)$$

and going by the parallel formula for polarization in (10.6), is easily seen to just be the Gini coefficient adapted to social groupings with 0-1 distances. The fractionalization index has been widely used in economics and political science, and we will encounter it again — and polarization as well — when we study social conflict later in the book.

10.5. Summary

In this chapter, we studied the measurement of *inequality* in the distribution of wealth or incomes. We argued that there are two reasons to be interested in inequality: the *intrinsic*, in which we value equality for its own sake and therefore regard inequality reduction as an objective in itself, and the *functional*, in which we study inequality to understand its impact on *other* features of the development process.

As a prelude to the study of measurement, we recognized that there were several conceptual issues. For instance, inequality in incomes may be compatible with overall equality simply because a society might display a high degree of *mobility*: movement of people from one income class to another. We also paid attention to the *functional distribution* of income as opposed to the *personal distribution* of income: *how* income is earned may have just as much social value as *how much* is earned.

With these caveats in mind, we then introduced four criteria for inequality measurement: (1) the *anonymity principle* (names do not matter), (2) the *population principle* (population size does not matter as long as the *composition* of different income classes stay the same in percentage terms), (3) the *relative income principle* (only relative incomes matter for the measurement of inequality, and not the absolute amounts involved), and (4) the *transfers principle* (if a transfer of income is made from a relatively poor to a relatively rich person, then inequality, however measured, registers an increase). It turns out that these four principles create a ranking of income distribution identical to that implied by the *Lorenz curve*, which displays how cumulative shares of income are earned by cumulatively increasing fractions of the population, arranged from poorest to richest.

However, the ranking is not complete. Sometimes two Lorenz curves cross. In such situations the four principles are not enough to make an unequivocal judgment about inequality. We argued that in this sense, our notions of inequality are fundamentally incomplete, but that forcing an additional degree of completeness by introducing more axioms may not necessarily be a good idea.

Complete measures of inequality do exist. These are measures that assign a degree of inequality (a number) to *every* conceivable income distribution, so they generate complete rankings. We studied examples of such measures that are popularly used in the literature: the *Kuznets ratios*, the *mean absolute deviation*, the *coefficient of variation*, and the *Gini coefficient*. Of these measures, the last two are of special interest in that they agree fully with our four principles (and so agree with the Lorenz ranking). That is, whenever the Lorenz ranking states that inequality has gone up, these two measures do not disagree. However, it is possible for these measures (and others) to disagree when Lorenz curves *do* cross.

Thus the theory of inequality measurement serves a double role. It tells us the ethical principles that are widely accepted and that we can use to rank different distributions of income or wealth, but it also warns us that such principles are incomplete, so we should not treat the behavior of any one complete measure at face value. We may not have direct information regarding the underlying Lorenz curves, but it is a good idea to look at the behavior of more than one measure before making a provisional judgment about the direction of change in inequality (if any such judgment can be made at all).

Finally, we go beyond inequality to *polarization*. Over and above the consideration of income differences, polarization emphasizes the bunching of population into distinct groups, thus reinforcing a sense of Us and Them. We showed that this notion *cannot*, in general, be compatible with the transfers principle. That does not dethrone the transfers principle; rather, it shows that polarization is a different measure of distributional characteristics that is conceptually separate from inequality. We end the chapter by showing how both polarization and inequality can be extended to other groupings not based on income.