We begin the chapter by defining *inequality* and *poverty,* terms that are commonly used in informal conversation but need to be measured more precisely to provide a meaningful understanding of how much progress has already been made, how much remains to be achieved, and how to set incentives for government officials to focus on the most pressing needs. You will see that the most important measures of poverty and inequality used by development economists satisfy properties that most observers would agree are of fundamental importance. After a discussion of why attention to inequality as well as poverty is important, we then use the appropriate measures of poverty and inequality to evaluate the welfare significance of alternative patterns (or "typologies") of growth. After reviewing the evidence on the extent of poverty and inequality in the developing world, we conclude with an overview of the key issues in poverty policy. Some important principles of effective poverty policies are considered, together with some initial examples of programs that have worked well in practice. We conclude the chapter with a comparative case study of Ghana and Côte d'Ivoire, which illustrate issues of the quality of growth and the difficulties of achieving it.

5.1 Measuring Inequality and Poverty

In this section, we define the dimensions of the income distribution and poverty problems and identify some similar elements that characterize the problem in many developing nations. But first we should be clear about what we are measuring when we speak about the distribution of income and absolute poverty.

Measuring Inequality

Economists usually distinguish between two principal measures of income distribution for both analytical and quantitative purposes: the personal or size distribution of income and the functional or distributive factor share distribution of income.

Personal distribution of income (size distribution of income) The distribution of income according to size class of persons—for example, the share of total income accruing to the poorest specific percentage or the richest specific percentage of a population—without regard to the sources of that income.

Quintile A 20% proportion of any numerical quantity. A population divided into quintiles would be divided into five groups of equal size.

Decile A 10% portion of any numerical quantity; a population divided into deciles would be divided into ten equal numerical groups.

Size Distributions The **personal** or **size distribution of income** is the measure most commonly used by economists. It simply deals with individual persons or households and the total incomes they receive. The way in which that income was received is not considered. What matters is how much each earns irrespective of whether the income was derived solely from employment or came also from other sources such as interest, profits, rents, gifts, or inheritance. Moreover, the locational (urban or rural) and occupational sources of the income (e.g., agriculture, manufacturing, commerce, services) are ignored. If Ms. X and Mr. Y both receive the same personal income, they are classified together irrespective of the fact that Ms. X may work 15 hours a day as a doctor while Mr. Y doesn't work at all but simply collects interest on his inheritance.

Economists and statisticians therefore like to arrange all individuals by ascending personal incomes and then divide the total population into distinct groups, or sizes. A common method is to divide the population into successive **quintiles** (fifths) or **deciles** (tenths) according to ascending income levels and

TABLE 5.1 Typical Size Distribution of Personal Income in a Developing Country by Income Shares—Quintiles and Deciles

Individuals	Personal Income (money units)	Share of Total Income (%)	
		Quintiles	Deciles
1	0.8		
2	1.0		1.8
3	1.4		
4	1.8	5	3.2
5	1.9		
6	2.0		3.9
7	2.4		
8	2.7	9	5.1
9	2.8		
10	3.0		5.8
11	3.4		
12	3.8	13	7.2
13	4.2		
14	4.8		9.0
15	5.9		
16	7.1	22	13.0
17	10.5		
18	12.0		22.5
19	13.5		
20	15.0	51	28.5
Total (national income)	100.0	100	100.0

then determine what proportion of the total national income is received by each income group. For example, Table 5.1 shows a hypothetical but fairly typical distribution of income for a developing country. In this table, 20 individuals, representing the entire population of the country are arranged in order of ascending annual personal income, ranging from the individual with the lowest income (0.8 units) to the one with the highest (15.0 units). The total or national income of all individuals amounts to 100 units and is the sum of all entries in column 2. In column 3, the population is grouped into quintiles of four individuals each. The first quintile represents the bottom 20% of the population on the income scale. This group receives only 5% (i.e., a total of 5 money units) of the total national income. The second quintile (individuals 5 through 8) receives 9% of the total income. Alternatively, the bottom 40% of the population (quintiles 1 plus 2) is receiving only 14% of the income, while the top 20% (the fifth quintile) of the population receives 51% of the total income.

A common measure of **income inequality** that can be derived from column 3 is the ratio of the incomes received by the top 20% and bottom 40% of the population. This ratio, sometimes called a Kuznets ratio after Nobel laureate Simon Kuznets, has often been used as a measure of the degree of inequality between high- and low-income groups in a country. In our example, this inequality ratio is equal to 51 divided by 14, or approximately 3.64.

Income inequality The disproportionate distribution of total national income among households.

To provide a more detailed breakdown of the size distribution of income, decile (10%) shares are listed in column 4. We see, for example, that the bottom

FIGURE 5.1 The Lorenz Curve

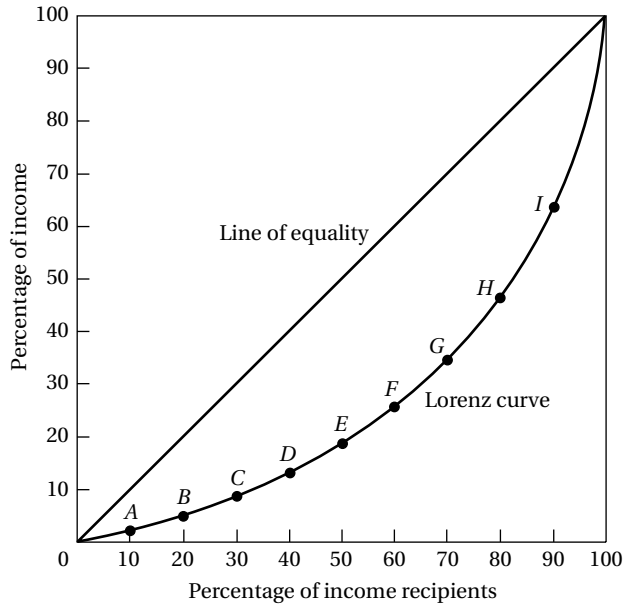

10% of the population (the two poorest individuals) receives only 1.8% of the total income, while the top 10% (the two richest individuals) receives 28.5%. Finally, if we wanted to know what the top 5% receives, we would divide the total population into 20 equal groups of individuals (in our example, this would simply be each of the 20 individuals) and calculate the percentage of total income received by the top group. In Table 5.1, we see that the top 5% of the population (the twentieth individual) receives 15% of the income, a higher share than the combined shares of the lowest 40%.

Lorenz curve A graph depicting the variance of the size distribution of income from perfect equality.

Lorenz Curves Another common way to analyze personal income statistics is to construct what is known as a **Lorenz curve**.[1] Figure 5.1 shows how it is done. The numbers of income recipients are plotted on the horizontal axis, not in absolute terms but in *cumulative percentages.* For example, at point 20, we have the lowest (poorest) 20% of the population; at point 60, we have the bottom 60%; and at the end of the axis, all 100% of the population has been accounted for. The vertical axis shows the share of total income received by each percentage of population. It is also cumulative up to 100%, meaning that both axes are the same length. The entire figure is enclosed in a square, and a diagonal line is drawn from the lower left corner (the origin) of the square to the upper right corner. At every point on that diagonal, the percentage of income received is *exactly equal* to the percentage of income recipients—for example, the point halfway along the length of the diagonal represents 50% of the income being distributed to exactly 50% of the population. At the three-quarters point on the diagonal, 75% of the income would be distributed to 75% of the

population. In other words, the diagonal line in Figure 5.1 is representative of "perfect equality" in size distribution of income. Each percentage group of income recipients is receiving that same percentage of the total income; for example, the bottom 40% receives 40% of the income, while the top 5% receives only 5% of the total income.[2]

The Lorenz curve shows the *actual* quantitative relationship between the percentage of income recipients and the percentage of the total income they did in fact receive during, say, a given year. In Figure 5.1, we have plotted this Lorenz curve using the decile data contained in Table 5.1. In other words, we have divided both the horizontal and vertical axes into ten equal segments corresponding to each of the ten decile groups. Point *A* shows that the bottom 10% of the population receives only 1.8% of the total income, point *B* shows that the bottom 20% is receiving 5% of the total income, and so on for each of the other eight cumulative decile groups. Note that at the halfway point, 50% of the population is in fact receiving only 19.8% of the total income.

The more the Lorenz line curves away from the diagonal (line of perfect equality), the greater the degree of inequality represented. The extreme case of perfect inequality (i.e., a situation in which one person receives all of the national income while everybody else receives nothing) would be represented by the congruence of the Lorenz curve with the bottom horizontal and right-hand vertical axes. Because no country exhibits either perfect equality or perfect inequality in its distribution of income, the Lorenz curves for different countries will lie somewhere to the right of the diagonal in Figure 5.1. The greater the degree of inequality, the greater the bend and the closer to the bottom horizontal axis the Lorenz curve will be. Two representative distributions are shown in Figure 5.2, one for a relatively equal distribution (Figure 5.2a) and the other for a more unequal distribution (Figure 5.2b). (Can you explain why the Lorenz curve could not lie above or to the left of the diagonal at any point?)

FIGURE 5.2 The Greater the Curvature of the Lorenz Line, the Greater the Relative Degree of Inequality

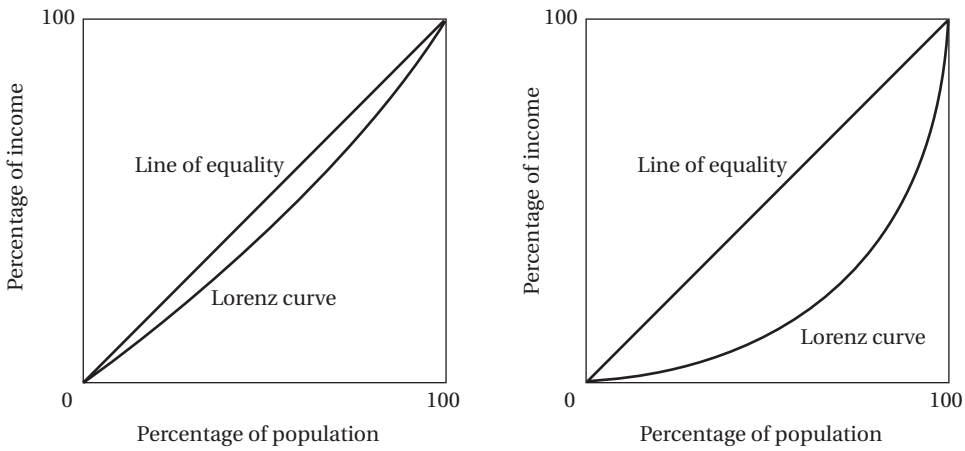

(a) A relatively equal distribution

(b) A relatively unequal distribution

FIGURE 5.3 Estimating the Gini Coefficient

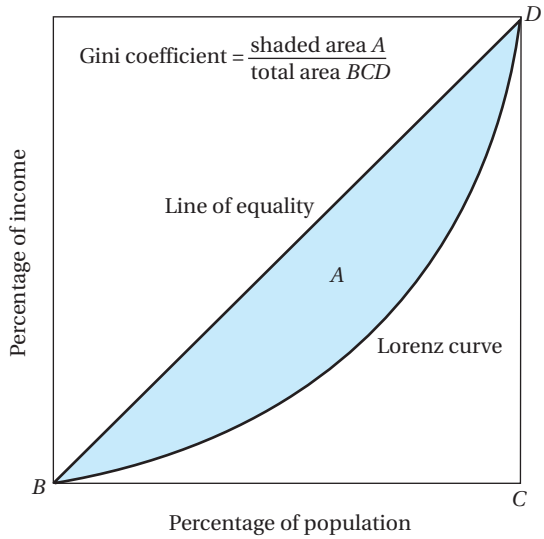

Gini Coefficients and Aggregate Measures of Inequality A final and very convenient shorthand summary measure of the relative degree of income inequality in a country can be obtained by calculating the ratio of the area between the diagonal and the Lorenz curve divided by the total area of the half-square in which the curve lies. In Figure 5.3, this is the ratio of the shaded area *A* to the total area of the triangle *BCD*. This ratio is known as the *Gini concentration ratio* or **Gini coefficient**, named after the Italian statistician who first formulated it in 1912.

Gini coefficient An aggregate numerical measure of income inequality ranging from 0 (perfect equality) to 1 (perfect inequality). It is measured graphically by dividing the area between the perfect equality line and the Lorenz curve by the total area lying to the right of the equality line in a Lorenz diagram. The higher the value of the coefficient, the higher the inequality of income distribution; the lower it is, the more equal the distribution of income.

Gini coefficients are aggregate inequality measures and can vary anywhere from 0 (perfect equality) to 1 (perfect inequality). In fact, as you will soon discover, the Gini coefficient for countries with highly unequal income distributions typically lies between 0.50 and 0.70, while for countries with relatively equal distributions, it is on the order of 0.20 to 0.35. The coefficient for our hypothetical distribution of Table 5.1 and Figure 5.1 is approximately 0.44—a relatively unequal distribution.

Four possible Lorenz curves such as might be found in international data are drawn in Figure 5.4. In the "Lorenz criterion" of income distribution, whenever one Lorenz curve lies above another Lorenz curve, the economy corresponding to the upper Lorenz curve is more equal than that of the lower curve. Thus economy *A* may unambiguously be said to be more equal than economy *D*. Whenever two Lorenz curves cross, such as curves *B* and *C*, the Lorenz criterion states that we "need more information" or additional assumptions before we can determine which of the underlying economies is more equal. For example, we might argue on the grounds of the priority of addressing problems of poverty that curve *B* represents a more equal economy, since the poorest are richer, even though the richest are also richer (and hence the middle class is "squeezed"). But others might start with the assumption

FIGURE 5.4 Four Possible Lorenz Curves

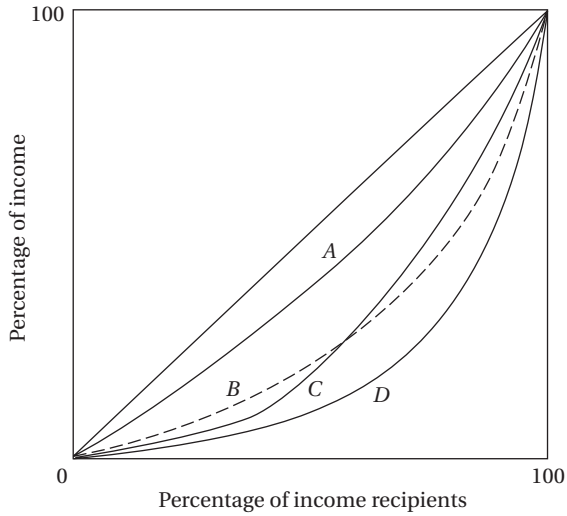

that an economy with a stronger middle class is inherently more equal, and those observer might select economy *C*.

One could also use an aggregate measure such as the Gini coefficient to decide the matter. As it turns out, the Gini coefficient is among a class of measures that satisfy four highly desirable properties: the anonymity, scale independence, population independence, and transfer principles.[3] The *anonymity principle* simply means that our measure of inequality should not depend on who has the higher income; for example, it should not depend on whether we believe the rich or the poor to be good or bad people. The *scale independence principle* means that our measure of inequality should not depend on the size of the economy or the way we measure its income; for example, our inequality measure should not depend on whether we measure income in dollars or in cents or in rupees or rupiahs or for that matter on whether the economy is rich on average or poor on average—because if we are interested in inequality, we want a measure of the dispersion of income, not its magnitude (note that magnitudes are very important in poverty measures). The *population independence principle* is somewhat similar; it states that the measure of inequality should not be based on the number of income recipients. For example, the economy of China should be considered no more or less equal than the economy of Vietnam simply because China has a larger population than Vietnam. Finally, we have the *transfer principle* (sometimes called the *Pigou-Dalton principle* after its creators); it states that, holding all other incomes constant, if we transfer some income from a richer person to a poorer person (but not so much that the poorer person is now richer than the originally rich person), the resulting new income distribution is more equal. If we like these four criteria, we can measure the Gini coefficient in each case and rank the one with the larger Gini as more unequal. However, this is not always a perfect solution. For example, the

Gini coefficient can, in theory, be identical for two Lorenz curves that cross; can you see why by looking at curves *B* and *C* in Figure 5.4? And sometimes different inequality measures that satisfy our four properties can give different answers as to which of two economies are more unequal.[4]

Note that a measure of dispersion common in statistics, the coefficient of variation (CV), which is simply the sample standard deviation divided by the sample mean, is another measure of inequality that also satisfies the four criteria. Although the CV is more commonly used in statistics, the Gini coefficient is often used in studies of income and wealth distribution due to its convenient Lorenz curve interpretation. Note, finally, that we can also use Lorenz curves to study inequality in the distribution of land, in education and health, and in other assets.

Functional Distributions The second common measure of income distribution used by economists, the **functional** or **factor share distribution of income**, attempts to explain the share of total national income that each of the **factors of production** (land, labor, and capital) receives. Instead of looking at individuals as separate entities, the theory of functional income distribution inquires into the percentage that labor receives as a whole and compares this with the percentages of total income distributed in the form of rent, interest, and profit (i.e., the returns to land and financial and physical capital). Although specific individuals may receive income from all these sources, that is not a matter of concern for the functional approach.

Functional distribution of income (factor share distribution of income) The distribution of income to factors of production without regard to the ownership of the factors.

Factors of production Resources or inputs required to produce a good or a service, such as land, labor, and capital.

A sizable body of theoretical literature has been built up around the concept of functional income distribution. It attempts to explain the income of a factor of production by the contribution that this factor makes to production. Supply and demand curves are assumed to determine the unit prices of each productive factor. When these unit prices are multiplied by quantities employed on the assumption of efficient (minimum-cost) factor utilization, we get a measure of the total payment to each factor. For example, the supply of and demand for labor are assumed to determine its market wage. When this wage is then multiplied by the total level of employment, we get a measure of total wage payments, also sometimes called the *total wage bill.*

Figure 5.5 provides a simple diagrammatic illustration of the traditional theory of functional income distribution. We assume that there are only two factors of production: capital, which is a fixed (given) factor, and labor, which is the only variable factor. Under competitive market assumptions, the demand for labor will be determined by labor's marginal product (i.e., additional workers will be hired up to the point where the value of their marginal product equals their real wage). But in accordance with the principle of diminishing marginal products, this demand for labor will be a declining function of the numbers employed. Such a negatively sloped labor demand curve is shown by line D_L in Figure 5.5. With a traditional neoclassical upward-sloping labor supply curve S_L, the equilibrium wage will be equal to W_E and the equilibrium level of employment will be L_E. Total national output (which equals total national income) will be represented by the area $0REL_E$.[5] This national income will be distributed in two shares: $0W_EEL_E$ going to workers in the form of wages and W_ERE remaining as capitalist profits (the return to owners of capital). Hence in a competitive market economy with constant-returns-to-scale

FIGURE 5.5 Functional Income Distribution in a Market Economy: An Illustration

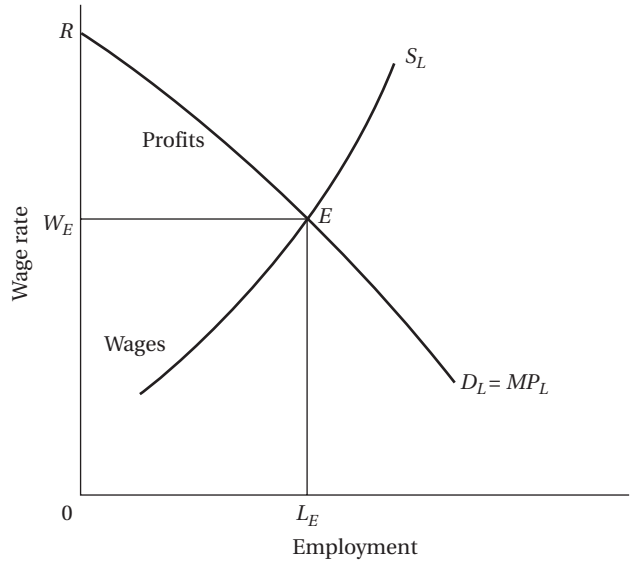

production functions (a doubling of all inputs doubles output), factor prices are determined by factor supply and demand curves, and factor shares always combine to exhaust the total national product. Income is distributed by function—laborers are paid wages, owners of land receive rents, and capitalists obtain profits. It is a neat and logical theory in that each and every factor gets paid only in accordance with what it contributes to national output, no more and no less. In fact, as you may recall from Chapter 3, this model of income distribution is at the core of the Lewis theory of modern-sector growth based on the reinvestment of rising capitalist profits.

Unfortunately, the relevance of the functional theory is greatly diminished by its failure to take into account the important role and influence of nonmarket forces such as power in determining these factor prices—for example, the role of collective bargaining between employers and trade unions in the setting of modern-sector wage rates and the power of monopolists and wealthy landowners to manipulate prices on capital, land, and output to their own personal advantage. Appendix 5.1 examines the economic implications of factor price distortions, and we return to consider their implications for policy at the end of this chapter.

Measuring Absolute Poverty

Now let's switch our attention from relative income shares of various percentile groups within a given population to the fundamentally important question of the extent and magnitude of **absolute poverty** in developing countries.

Absolute poverty The situation of being unable or only barely able to meet the subsistence essentials of food, clothing, and shelter.

In Chapter 2, we defined the extent of absolute poverty as the number of people who are unable to command sufficient resources to satisfy basic needs. They are counted as the total number living below a specified minimum level of real income—an international poverty line. That line knows no national boundaries, is independent of the level of national per capita income, and takes into account differing price levels by measuring poverty as anyone living on less than $1.25 a day or $2 per day in PPP dollars. Absolute poverty can and does exist, therefore, as readily in New York City as it does in Kolkata, Cairo, Lagos, or Bogotá, although its magnitude is likely to be much lower in terms of percentages of the total population.

Headcount index The proportion of a country's population living below the poverty line.

Absolute poverty is sometimes measured by the number, or "headcount," H, of those whose incomes fall below the absolute poverty line, Y_p. When the headcount is taken as a fraction of the total population, N, we define the **headcount index**, H/N. The poverty line is set at a level that remains constant in real terms so that we can chart our progress on an absolute level over time. The idea is to set this level at a standard below which we would consider a person to live in "absolute human misery," such that the person's health is in jeopardy.

Of course, to define a minimum health standard that is invariant across historical epochs is an impossibility, in part because technology changes over time. For example, today we have 15-cent oral rehydration therapy packets that can save the life of a child in Malawi. Not long ago, the death of a child after a diarrheal disease would be taken as a sad but inevitable part of life, whereas today we regard such a death as a catastrophic moral failure of the international community. We simply come as close as we can to establishing a reasonable minimum standard that might hold over a few decades so that we can estimate more carefully how much progress we have made on a (more) absolute rather than a (highly) relative scale.

Certainly one would not accept the international poverty level of $1.25 a day in an unquestioning way when planning local poverty work. One practical strategy for determining a local absolute poverty line is to start by defining an adequate basket of food, based on nutritional requirements from medical studies of required calories, protein, and micronutrients. Then, using local household survey data, one can identify a typical basket of food purchased by households that just barely meet these nutritional requirements. One then adds other expenditures of this household, such as clothing, shelter, and medical care, to determine the local absolute poverty line. Depending on how these calculations are done, the resulting poverty line may come to more than $1.25 per day at PPP.

In many respects, however, simply counting the number of people below an agreed-on poverty line can have its limitations. For example, if the poverty line is set at U.S. $450 per person, it makes a big difference whether most of the absolute poor earn $400 or $300 per year. Both are accorded the same weight when calculating the proportion of the population that lies below the poverty line; clearly, however, the poverty problem is much more serious in the latter instance. Economists therefore attempt to calculate a **total poverty gap (TPG)** that measures the total amount of income necessary to raise everyone who is below the poverty line up to that line. Figure 5.6 illustrates how we could measure the total poverty gap as the shaded area between poverty line, PV, and the annual income profile of the population.

Total poverty gap (TPG) The sum of the difference between the poverty line and actual income levels of all people living below that line.

FIGURE 5.6 Measuring the Total Poverty Gap

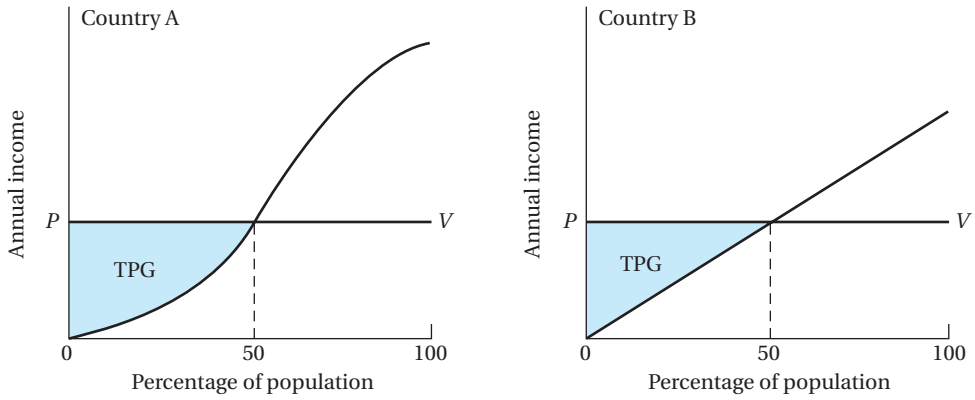

(a) A relatively large poverty gap

(b) A relatively small poverty gap

Even though in both country A and country B, 50% of the population falls below the same poverty line, the TPG in country A is greater than in country B. Therefore, it will take more of an effort to eliminate absolute poverty in country A.

The TPG—the extent to which the incomes of the poor lie below the poverty line—is found by adding up the amounts by which each poor person's income, Y_i, falls below the absolute poverty line, Y_p, as follows:

$$\text{TPG} = \sum_{i=1}^{H} (Y_p - Y_i) \tag{5.1}$$

We can think of the TPG in a simplified way (i.e., no administrative costs or general equilibrium effects are accounted for) as the amount of money per day it would take to bring every poor person in an economy up to our defined minimum income standards. On a per capita basis, the *average poverty gap* (APG) is found by dividing the TPG by the total population:

$$\text{APG} = \frac{\text{TPG}}{N} \tag{5.2}$$

Often we are interested in the size of the poverty gap in relation to the poverty line, so we would use as our income shortfall measure the *normalized poverty gap* (NPG): NPG = APG/Y_p; this measure lies between 0 and 1 and so can be useful when we want a unitless measure of the gap for easier comparisons.

Another important poverty gap measure is the *average income shortfall* (AIS), which is the total poverty gap divided by the headcount of the poor: AIS = TPG/H. The AIS tells us the average amount by which the income of a poor person falls below the poverty line. This measure can also be divided by the poverty line to yield a fractional measure, the *normalized income shortfall* (NIS): NIS = AIS/Y_p.

The Foster-Greer-Thorbecke Index We are also often interested in the degree of income inequality among the poor, such as the Gini coefficient among those who are poor, G_p, or alternatively, the coefficient of variation (CV) of incomes among the poor, CV_p. One reason that the Gini or CV among the poor can be important is that the impact on poverty of economic shocks can differ greatly, depending on the level and distribution of resources among the poor. For example, if the price of rice rises, as it did in 1998 in Indonesia, low-income rice producers, who sell a little of their rice on local markets and whose incomes are slightly below the absolute poverty line, may find that this price rise increases their incomes to bring them out of absolute poverty. On the other hand, for those with too little land to be able to sell any of the rice they grow and who are net buyers of rice on markets, this price increase can greatly worsen their poverty. Thus the most desirable measures of poverty would also be sensitive to the distribution of income among the poor.

As was the case with inequality measures, there are criteria for a desirable poverty measure that are widely accepted by development economists: the anonymity, population independence, monotonicity, and distributional sensitivity principles. The first two principles are very similar to the properties we examined for inequality indexes: Our measure of the extent of poverty should not depend on who is poor or on whether the country has a large or small population. The monotonicity principle means that if you add income to someone below the poverty line, all other incomes held constant, poverty can be *no greater* than it was.[6] The distributional sensitivity principle states that, other things being equal, if you transfer income from a poor person to a richer person, the resulting economy should be deemed strictly poorer. The headcount ratio measure satisfies anonymity, population independence, and monotonicity, but it fails on distributional sensitivity. The simple headcount fails even to satisfy the population independence principle.

Foster-Greer-Thorbecke (FGT) index A class of measures of the level of absolute poverty.

A well-known poverty index that in certain forms satisfies all four criteria is the **Foster-Greer-Thorbecke (FGT) index**, often called the P_α class of poverty measures.[7] The P_α index is given by

$$P_\alpha = \frac{1}{N}\sum_{i=1}^{H}\left(\frac{Y_p - Y_i}{Y_p}\right)^\alpha \tag{5.3}$$

where Y_i is the income of the ith poor person, Y_p is the poverty line, and N is the population. Depending on the value of α, the P_α index takes on different forms. If $\alpha = 0$, the numerator is equal to H, and we get the headcount ratio, H/N. If $\alpha = 1$, we get the normalized poverty gap.

If $\alpha = 2$, the impact on measured poverty of a gain in income by a poor person increases in proportion to the distance of the person from the poverty line. For example, raising the income of a person from a household living at half the per capita poverty line by, say, one penny per day would have five times the impact on poverty reduction as would raising by the same amount the income of a person living at 90% of the poverty line.

If $\alpha = 2$, the resulting measure, P_2, can be rewritten as[8]

$$P_2 = \left(\frac{H}{N}\right)[\text{NIS}^2 + (1 - \text{NIS})^2\,(\text{CV}_P)^2] \tag{5.4}$$

As Equation 5.4 shows, P_2 contains the CV_p measure, and it satisfies all four of the poverty axioms.[9] Clearly, P_2 increases whenever H/N, NIS, or CV_p increases. Note from the formula that there is a greater emphasis on the distribution of income among the poor (CV_p) when the normalized income shortfall is small and a smaller emphasis when the NIS is large.

P_2 has become a standard of income poverty measure used by the World Bank and other agencies, and it is used in empirical work on income poverty because of its sensitivity to the depth and severity of poverty. For the same reason, the P_2 measure has now become part of the Mexican constitution (chap. 5, art. 34). Mexico uses the P_2 poverty measure to allocate funds for education, health, and welfare programs for the poor (in particular in the Progresa Opportunidades program, described at the end of Chapter 8), in accordance with the regional intensity of poverty.[10]

Values of P_0 and P_2 for selected developing countries are found in Table 5.6 later in this chapter.

The Newly Introduced Multidimensional Poverty Index Poverty cannot be adequately measured with income, as Amartya Sen's capability framework, examined in Chapter 1, makes apparent. Income is imperfectly measured, but even more important, the advantages provided by a given amount of income greatly differ, depending on circumstances. To capture this idea the United Nations Development Program used its *Human Poverty Index*[11] from 1997 to 2009. In 2010, the UNDP replaced the HPI with its new **Multidimensional Poverty Index (MPI)**; by building up the index from the household level, the MPI takes into account that there are negative interaction effects when people have multiple deprivations—worse poverty than can be seen by simply adding up separate deprivations for the whole country, taking averages, and only then combining them. The first step in measuring poverty is to know which people are poor. In the multidimensional poverty approach, a poor person is identified through what is called the "dual cutoff method": first, the cutoff levels within each of the dimensions (analogous to falling below a poverty line such as $1.25 per day if income poverty were being addressed), and second, the cutoff of the number of dimensions in which a person must be deprived (below the line) to be deemed multidimensionally poor.

Multidimensional Poverty Index (MPI) A poverty measure that identifies the poor using dual cutoffs for levels and numbers of deprivations, and then multiplies the percentage of people living in poverty times the percent of weighted indicators for which poor households are deprived on average.

In applied studies, we need proxy measures, called *indicators,* for each of the selected dimensions. The index's creators report that they selected the three dimensions (health, education, and standard of living) and each of their corresponding indicators because they reflect problems often mentioned by the poor, they have been long considered important by the development community particularly as reflected in the Millennium Development Goals (see Chapter 1), and they are well established philosophically as human rights or basic needs; naturally, reliable data also had to be available for enough countries.

With respect to health, two indicators—whether any child has died in the family and whether any adult or child in the family is malnourished—are weighted equally (so each counts one-sixth toward the maximum possible deprivation in the MPI). Regarding education also, two indicators—whether not even one household member has completed five years of schooling and whether any school-age child is out of school for grades one through eight—are given equal weight (so again, each counts one-sixth toward the MPI). Finally, in

terms of standard of living, equal weight is placed on six deprivations (each counting one-eighteenth toward the maximum possible): lack of electricity, insufficiently safe drinking water, inadequate sanitation, inadequate flooring, unimproved cooking fuel, and lack of more than one of five assets—telephone, radio, television, bicycle, and motorbike or similar vehicle.

Calculating deprivation in this way, individuals in a family are then identified as "multidimensionally poor" when deprived by a "weighted sum" of 0.3 or more (3 out of 10 points as calculated in practice). For concreteness, a person would get a value of 33% and thus be considered poor by having a child in the family who is malnourished while at the same time the most educated person in the family received only three years of schooling. Or a multidimensionally poor person might live in a household that has experienced a child death and is also deprived in at least three of the six living standards indicators, which also sums to $1/6 + 1/18 + 1/18 + 1/18 = 1/3$, or 33%. Or they could live in a household that is deprived in the other three living standard indicators and in which there is a school-age child not attending school. But if there were no health or education deprivations, a person would have to be deprived in all six standard-of-living indicators to be deemed poor. Thus the MPI approach identifies the very poor by measuring a range of important household deprivations directly, rather than only indirectly through income, then building the index from household measures up to the aggregate measure. Rather than using already aggregated statistics in an index, the approach takes into account the *multiplied or interactive harm* done when multiple deprivations are experienced by the *same individual or family*. In essence, the approach assumes that an individual's lack of capability in one area can to a degree be made up for by other capabilities—but only to a degree. (Put differently, capabilities are treated as substitutes up to a point but then as complements.) This greatly augments measures used previously.

Finally, the actual MPI for the country (or region or group) is computed; a convenient way to express the resulting value is the product of the headcount ratio H_M (the percentage of people living in multidimensional poverty) and the average intensity of deprivation A (the percentage of weighted indicators for which poor households are deprived on average). The adjusted headcount ratio $H_M A$ is readily calculated, and it also satisfies some desirable properties. One of these is *dimensional monotonicity,* meaning that when a person deemed poor becomes deprived in another indicator, he or she is deemed even poorer; this would not be the case if using the simple headcount ratio. Moreover, $H_M A$ is a special case of a broader class of multidimensional poverty measures developed by Sabira Alkire and James Foster; for example, with better data another of these measures may be used to show the severity of multidimensional poverty, analogous to P_2.[12]

The UNDP reports the MPI for 104 developing countries, based on the currently available data; some examples are given in Table 5.2. More countries may be added as data availability improves, but already the countries encompass nearly 78% of the world's population. The index can range from 0 to 1. Slovenia and Slovakia receive an MPI of 0, the lowest possible value, indicating the least poverty, while the world's most impoverished country for which data were available to compute a ranking, Niger, ranks 104th, with an MPI value of 0.642. Based on the thresholds just described, the report found there

TABLE 5.2 MPI Rankings and Poverty Headcounts for Selected Countries

			Multidimensional Poverty			Income Poverty			
						$1.25 a Day (proportion of poor)		$2 a Day (proportion of poor)	
Country	Year	MPI Value	MPI Rank	H_M (proportion of poor)	A (average intensity of deprivations)	Value	Rank	Value	Rank
Kazakhstan	2006	0.002	7	0.006	0.369	0.031	23	0.172	29
Thailand	2005	0.006	16	0.016	0.385	0.020	1	0.115	20
Ecuador	2003	0.009	24	0.022	0.416	0.047	26	0.128	23
Mexico	2006	0.015	29	0.040	0.389	0.020	1	0.048	16
Brazil	2003	0.039	39	0.085	0.460	0.052	29	0.127	21
Colombia	2005	0.041	40	0.092	0.441	0.160	42	0.279	35
Dominican Republic	2000	0.048	42	0.111	0.433	0.050	28	0.151	27
China	2003	0.056	44	0.125	0.449	0.159	41	0.363	41
Viet Nam	2002	0.075	50	0.143	0.525	0.215	50	0.484	51
Indonesia	2007	0.095	53	0.208	0.459	0.075	31	0.490	52
Ghana	2008	0.140	57	0.301	0.464	0.300	57	0.536	56
Zimbabwe	2006	0.174	60	0.385	0.452				
Bolivia	2003	0.175	61	0.363	0.483	0.196	46	0.303	38
Nicaragua	2001	0.211	64	0.407	0.519	0.158	40	0.318	40
Lao	2006	0.267	68	0.472	0.565	0.440	46	0.768	73
Pakistan	2007	0.275	70	0.510	0.540	0.226	53	0.603	59
Yemen	2006	0.283	71	0.525	0.539	0.175	43	0.466	49
Bangladesh	2007	0.291	73	0.578	0.504	0.496	71	0.813	80
India	2005	0.296	74	0.554	0.535	0.416	64	0.756	70
Kenya	2003	0.302	76	0.604	0.500	0.197	47	0.399	43
Haiti	2006	0.306	77	0.573	0.533	0.549	76	0.721	67
Côte d'Ivoire	2005	0.320	78	0.522	0.614	0.233	55	0.468	50
Nepal	2006	0.350	82	0.647	0.540	0.551	77	0.776	76
Tanzania	2008	0.367	84	0.653	0.563	0.885	93	0.966	93
DR Congo	2007	0.393	88	0.732	0.537	0.592	79	0.795	77
Madagascar	2004	0.413	91	0.705	0.585	0.678	86	0.896	87
Angola	2001	0.452	93	0.774	0.584	0.543	89	0.900	88
Ethiopia	2005	0.582	103	0.900	0.647	0.390	62	0.775	75
Niger	2006	0.642	104	0.927	0.693	0.659	85	0.856	85

were nearly 1.7 billion people living in what was termed "acute" poverty—several hundred million more than the estimated number living on income of less than $1.25 per day. At the broadest level, the results are not out of line with what one might expect; sub-Saharan Africa has the highest *proportion* of people living in poverty, and South Asia has the largest *number* of people living in poverty (almost twice as many as in Africa).

Only Niger had an MPI higher than 0.6. Seven other countries had an MPI higher than 0.5, all in sub-Saharan Africa: Ethiopia, Mali, Burkina Faso, Burundi, Somalia, Central African Republic, and Guinea. A further eight countries had MPIs between 0.4 and 0.5; all of these were African nations as well. The severity of poverty in Africa is also brought home by some of the findings. In Guinea, Mali, and Niger, more than 50% are poor and live in a household in which at least one child has died. In Mozambique, Guinea, Burundi, Mali, Ethiopia, Burkina Faso, and Niger, more than 50% live in a poor household where no one has completed five years of education.

The poorest non-African countries were Nepal (with an MPI of 0.350), Haiti (0.306), India (0.296), Bangladesh (0.291), Yemen (0.283), and Pakistan (0.275). But these six countries have a combined population of about 1.6 billion people. And 39% in India and 37% in Bangladesh live in a poor household where at least one child or woman is undernourished.

The results showed that knowing income poverty is not enough if our concern is with multidimensional poverty. For example, multidimensionally, Bangladesh is substantially less poor and Pakistan substantially poorer than would be predicted by these countries' income poverty (this finding may be related to some of the comparisons in the case study in Chapter 2). In Africa, Ethiopia is far more multidimensionally poor and Tanzania much less so than predicted by income poverty. Most Latin American countries studied rank worse on multidimensional poverty than on income poverty, but Colombia's income and MPI poverty ranks are about the same.

Different regions in the same country can have very different MPIs. Overall, the MPI for Kenya (whose regional HDI differences were examined in Chapter 2) is similar to that of India. But in Kenya, the MPI for Nairobi is close to that of Brazil. Central Kenya's MPI is similar to that of Bolivia. And northeastern Kenya has a worse MPI even than Niger. There were also great inequalities across ethnic groups in Kenya, with 29% of the Embu considered multidimensionally poor, compared with a staggering 96% of the Turkana and Masai peoples. Great inequalities were also found in India, in which indigenous ("tribal") peoples and low-ranked ("scheduled") castes were far poorer than people from high-ranking castes. In the Delhi and Kerala regions, just 14 to 16% are MPI poor, but in Jharkhand and Bihar, 77 to 81% are MPI poor. Finally, changes in the MPI over time were examined for three countries: Ghana saw its MPI halved from 0.29 to 0.14; Bangladesh saw its MPI reduced by a more modest 22%; and in Ethiopia, the MPI fell by 16% in the periods studied.

As with all indexes, the MPI has some limitations. As mentioned, data are from the household rather than the individual level (such as whether *any* child of school age is out of school or whether *any* family member is undernourished). It does not fully distinguish between past and present conditions (because its measure is whether a child has *ever* died). It does not distinguish differences

within households (such as who may make use of the bicycle or whether the undernourished individuals are females). Proxies are imperfect; for example, nourishment does not capture micronutrient deficiencies. Sometimes a person has to be labeled nondeprived if data are missing, so the numbers may understate poverty somewhat. Education considers only inputs such as enrolling or attending for five years, not outputs such as being able to read. And the choice of basic assets is questionable; for example, even where a radio and a simple bicycle are present, a woman may have just one dress and the children may sleep on a rough concrete floor.

No doubt there will be future improvements in the MPI. The MPI provides a new and fundamentally important way to measure poverty, to help us understand how poverty levels differ across and within countries and also how the dimensions (or composition) of poverty can differ greatly in different settings. Ultimately, this should assist with better design and targeting of programs and policies and help us evaluate their performance more quickly and effectively.

For now, because of the way living standards and human development surveys are conducted, most of the usable data is at the household level, making it difficult to "drill down" to the individual level. Household data are far better than what used to be available; in fact, the availability of household data has already had a substantial impact on improving the study of development economics. It is a great improvement to be able to focus on what is happening at the family rather than the national level. Well-designed income poverty measures such as P_2 will always be used for many purposes; but the MPI is likely to help usher in an era in which multidimensional poverty is examined in most assessments.

The ACWI A final approach to accounting for the distribution of income in assessing the quality of growth is to value increases in income for all individuals but to assign a higher weight to income gains by lower-income individuals than to gains by higher-income individuals. Perhaps the best-known example is the Ahluwalia-Chenery Welfare Index (ACWI), which is explained in Appendix 5.2.

5.2 Poverty, Inequality, and Social Welfare

What's So Bad about Extreme Inequality?

Throughout this chapter, we are assuming that social welfare depends positively on the level of income per capita but negatively on poverty and negatively on the level of inequality, as these terms have just been defined. The problem of absolute poverty is obvious. No civilized people can feel satisfied with a state of affairs in which their fellow humans exist in conditions of such absolute human misery, which is probably why every major religion has emphasized the importance of working to alleviate poverty and is at least one of the reasons why international development assistance has the nearly universal support of every democratic nation. But it may reasonably be asked, if our top priority is the alleviation of absolute poverty, why should *relative inequality*