

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 illustrates, issues of the quality of growth and the difficulties of achieving it.

5.1 Measuring Inequality

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.

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 they received that income is not considered. What matters is how much each earns irrespective of whether the income is derived solely from employment or comes 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 then determine what proportion of the total national income is received

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

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.

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 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

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

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 Curves

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

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

FIGURE 5.1 The Lorenz Curve




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 relatively 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



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 is, 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 that

FIGURE 5.3 Estimating the Gini Coefficient





FIGURE 5.4 Four Possible Lorenz Curves



an economy with a stronger middle class is inherently more equal, and those observers 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


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.

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.

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 $OREL_E$.⁵ This national income will be distributed in two shares: OW_EEL_E going to workers in the form of wages and $W_ER E$ 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.

The Ahluwalia-Chenery Welfare Index (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 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.

Income Poverty

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.

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 (also referred to as the “headcount ratio”). 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.

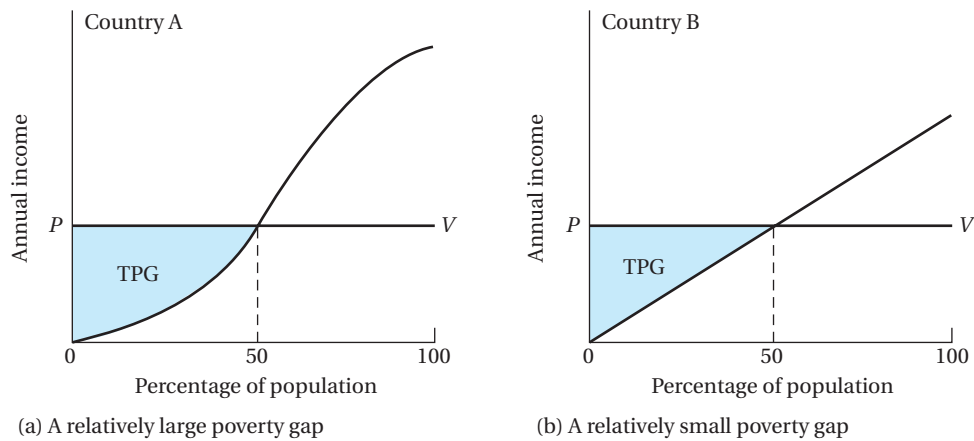
Headcount index The proportion of a country's population living below the poverty line.

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.

However, simply counting the number of people below an agreed-on poverty line has serious 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

FIGURE 5.6 Measuring the Total Poverty Gap



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 can measure the total poverty gap as the shaded area between poverty line, PV , and the annual income profile of the population.

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:

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

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:

$$\text{APG} = \frac{\text{TPG}}{N} \quad (5.2)$$

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

Total poverty gap (TPG)

The sum of the difference between the poverty line and actual income levels of all people living below that line.

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 is 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.

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

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

where Y_i is the income of the i th 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 . Unfortunately, this measure is the same whether those in poverty earn 90 cents per day or 50 cents per day, so it cannot reveal the depth of poverty.

If $\alpha = 1$, we get the normalized (per capita) poverty gap. An alternative formula that can be derived for P_1 is given by $P_1 = (H/N) \cdot (NIS)$, that is, the headcount ratio (H/N) times the normalized income shortfall (NIS). So, P_1 has

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

the properties that poverty goes up whenever either the fraction of people in poverty goes up or the fractional income deficits (poverty depth) go up (or both)—in general, this makes it a better measure than P_0 .

If $\alpha = 2$, we account for poverty severity, in that the impact on measured poverty of a gain in income by a poor person increases in relation to the square of 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; this differing magnitude results from squaring the poverty gaps, so the P_2 measure captures the *severity* of poverty.

As a numerical example of the calculation of P_2 , consider an 8-person economy with a poverty line of 1, and a hypothetical income distribution of: (0.6, 0.6, 0.8, 0.8, 2, 2, 6, 6). The headcount is 4, because two people have incomes of 0.6 and two people have incomes of 0.8; but the others have incomes above the poverty line. Using these numbers, we can find the P_2 level of poverty from equation 5.3:

$$P_2 = (1/8)[0.4^2 + 0.4^2 + 0.2^2 + 0.2^2] = (1/8) [0.16 + 0.16 + 0.04 + 0.04] = 0.4/8 = 0.05$$

Note that P_2 can be expressed in an alternative form to add further intuition. If $\alpha = 2$, the resulting measure, P_2 , can be rewritten as⁸

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

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 lesser emphasis when the NIS is large.

The **P_2 poverty measure**, also known as the *squared poverty gap index*, 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. Mexico uses the P_2 poverty measure to allocate funds for education, health, and welfare programs for the poor (in particular in the Progresa/Oportunidades Program, described at the end of Chapter 8), in accordance with the regional intensity of poverty.¹⁰

Another reason to prefer P_2 (or at least P_1) over P_0 is that standard headcount measures also have the perverse property of creating incentives for officials to focus efforts on the poor who are closest to the poverty line—because that is the easiest and cheapest way for them to demonstrate progress. We encountered a version of this problem in Chapter 1—a critique of the Millennium Development Goals focus on reducing the fraction of those living below the poverty line.

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

Person-Equivalent Headcounts Although P_1 and P_2 are more informative measures, which provide better incentives to poverty programs than P_0 , many agencies (including U.S. Agency for International Development—USAID)

continue to report progress primarily if not exclusively in terms of P_0 headcount measures—apparently responding to public and legislative expectations to discuss poverty in terms of numbers of people. Given a political need to feature “headline” headcount measures, a partial improvement is to convert changes in the poverty gap into its headcount-equivalent (based on the initial average income shortfall). If aid agencies featured a supplementary headcount-equivalent, they could report in terms of numbers of people while accounting for changes in poverty depth. Estimates using this approach show progress against poverty in many countries is significantly greater than revealed using conventional headcount measures alone.¹¹

Multidimensional Poverty Measurement Poverty cannot be adequately measured with income alone, as Amartya Sen’s capability framework, examined in Chapter 1, makes apparent. To fill this gap, Sabina Alkire and James Foster have extended the FGT index to multiple dimensions.¹²

As always, 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. Using calculations analogous to the single-dimensional P_α index, the multidimensional M_α index is constructed. The most basic measure is the fraction of the population in multidimensional poverty—the multidimensional headcount ratio H_M .

The most common measure in practice is M_0 , the *adjusted* headcount ratio, which uses ordinal data and is similar conceptually to the poverty gap P_1 (which again can be expressed as the headcount ratio times the normalized income shortfall). M_0 may be represented by the product of the multidimensional headcount ratio times the average fraction of dimensions in which the poor are deprived (or “average intensity of poverty” A , that is, $M_0 = H_M * A$). (In contrast to the simple multidimensional headcount ratio, the adjusted multidimensional headcount ratio satisfies the desirable property (called “dimensional monotonicity”) that if the average fraction of deprivations increases, so does M_0).

In applied studies, proxy measures, called *indicators*, are used for each of the selected dimensions. Details of the way this measure has been constructed and applied in the UNDP Multidimensional Poverty Index and findings across countries are reported in Section 5.4, when we apply the poverty measures to examine the extent of poverty in different countries and regions. Another wisely used application is the Women’s Empowerment in Agriculture Index, referred to in Chapter 9.

5.3 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* be a concern? We have seen that inequality among the poor is a critical factor in understanding the severity of poverty and the impact of market and policy changes on the poor, but why should we be concerned with inequality among those *above* the poverty line?

There are three major answers to this question. First, extreme income inequality leads to economic inefficiency. This is partly because at any given average income, the higher the inequality is, the smaller the fraction of the population that qualifies for a loan or other credit. Indeed, one definition of *relative poverty* is the lack of collateral. When low-income individuals (whether they are absolutely poor or not) cannot borrow money, they generally cannot adequately educate their children or start and expand a business. Moreover, with high inequality, the overall rate of savings in the economy tends to be lower, because the highest rate of marginal savings is usually found among the middle classes. Although the rich may save a larger dollar amount, they typically save a smaller fraction of their incomes, and they almost always save a smaller fraction of their marginal incomes. Landlords, business leaders, politicians, and other rich elites are known to spend much of their incomes on imported luxury goods, gold, jewelry, expensive houses, and foreign travel or to seek safe havens abroad for their savings in what is known as *capital flight*. Such savings and investments do not add to the nation's productive resources; in fact, they represent substantial drains on these resources. In short, the rich do not generally save and invest significantly larger proportions of their incomes (in the real economic sense of productive domestic saving and investment) than the middle class or even the poor.¹³ Furthermore, inequality may lead to an inefficient allocation of assets. As you will see in Chapter 8, high inequality leads to an overemphasis on higher education at the expense of quality universal primary education, which not only may be inefficient but is also likely to beget still more inequality in incomes. Moreover, as you will see in Chapter 9, high inequality of land ownership—characterized by the presence of huge *latifundios* (plantations) alongside tiny *minifundios* that are incapable of supporting even a single family—also leads to inefficiency because the most efficient scales for farming are family and medium-size farms. The result of these factors can be a lower average income and a lower rate of economic growth when inequality is high.¹⁴

The second reason to be concerned with inequality above the poverty line is that extreme income disparities undermine social stability and solidarity. Also, high inequality strengthens the political power of the rich and hence their economic bargaining power. Usually this power will be used to encourage outcomes favorable to themselves. High inequality facilitates *rent seeking*, including actions such as excessive lobbying, large political donations, bribery, and cronyism. When resources are allocated to such rent-seeking behaviors, they are diverted from productive purposes that could lead to faster growth. Even worse, high inequality makes poor institutions very difficult to improve,

because the few with money and power are likely to view themselves as worse off from socially efficient reform, and so they have the motive and the means to resist it (see Chapter 2). Of course, high inequality may also lead the poor to support populist policies that can be self-defeating. Countries with extreme inequality, such as El Salvador and Iran, have undergone upheavals or extended civil strife that have cost countless lives and set back development progress by decades. High inequality is also associated with pathologies such as higher violent crime rates. In sum, with high inequality, the focus of politics often tends to be on supporting or resisting the redistribution of the existing economic pie rather than on policies to increase its size (Chapter 11 examines these concerns in more detail).¹⁵

Finally, extreme inequality is generally viewed as unfair. The philosopher John Rawls proposed a thought experiment to help clarify why this is so.¹⁶ Suppose that before you were born into this world, you had a chance to select the overall level of inequality among the earth's people but not your own identity. That is, you might be born as Bill Gates, but you might be born as the most wretchedly poor person in rural Ethiopia with equal probability. Rawls calls this uncertainty the "veil of ignorance." The question is, facing this kind of risk, would you vote for an income distribution that was more equal or less equal than the one you see around you? If the degree of equality had no effect on the level of income or rate of growth, most people would vote for nearly perfect equality. Of course, if everyone had the same income no matter what, there would be little incentive to work hard, gain skills, or innovate. As a result, most people vote for *some* inequality of income outcomes, to the extent that these correspond to incentives for hard work or innovation. But even so, most vote for *less* inequality than is seen in the world (or in virtually any country) today. This is because much of the inequality we observe in the world is based on luck or extraneous factors, such as the inborn ability to kick a football or the identity of one's great-grandparents.

For all these reasons, for this part of the analysis we will write welfare, W , as

$$W = W(Y, I, P) \quad (5.5)$$

where Y is income per capita and enters our welfare function positively, I is inequality and enters negatively, and P is absolute poverty and also enters negatively. These three components have distinct significance, and we need to consider all three elements to achieve an overall assessment of welfare in developing countries. (A similar framework can be applied to health and education.)

Dualistic Development and Shifting Lorenz Curves: Some Stylized Typologies

As introduced by Gary Fields, Lorenz curves may be used to analyze three limiting cases of dualistic development:¹⁷

1. The *modern-sector enlargement* growth typology, in which the two-sector economy develops by enlarging the size of its modern sector while maintaining constant wages in both sectors. This is the case depicted by the

Lewis model in Chapter 3. It corresponds roughly to the historical growth pattern of Western developed nations and, to some extent, the pattern in East Asian economies such as China, South Korea, and Taiwan.

2. The *modern-sector enrichment* growth typology, in which the economy grows but such growth is limited to a fixed number of people in the modern sector, with both the numbers of workers and their wages held constant in the traditional sector. This roughly describes the experience of many Latin American and African economies.
3. The *traditional-sector enrichment* growth typology, in which all of the benefits of growth are divided among traditional-sector workers, with little or no growth occurring in the modern sector. This process roughly describes the experiences of countries whose policies focused on achieving substantial reductions in absolute poverty even at very low incomes and with relatively low growth rates, such as Sri Lanka, and the state of Kerala in southwestern India.

Using these three special cases and Lorenz curves, Fields demonstrated the validity of the following propositions (reversing the order just presented):

1. In the *traditional-sector enrichment* typology, growth results in higher income, a *more equal* relative distribution of income, and less poverty. Traditional-sector enrichment growth causes the Lorenz curve to shift uniformly upward and closer toward the line of equality, as depicted in Figure 5.7.
2. In the *modern-sector enrichment* growth typology, growth results in higher incomes, a *less equal* relative distribution of income, and no change in poverty.

FIGURE 5.7 Improved Income Distribution under the Traditional-Sector Enrichment Growth Typology





FIGURE 5.8 Worsened Income Distribution under the Modern-Sector Enrichment Growth Typology




Modern-sector enrichment growth causes the Lorenz curve to shift downward and farther from the line of equality, as shown in Figure 5.8.

3. Finally, in the case of Lewis-type, *modern-sector enlargement* growth, absolute incomes rise and absolute poverty is reduced, but the Lorenz curves will always cross, indicating that we cannot make any unambiguous statement about changes in relative inequality: It may improve or worsen. Fields shows that if, in fact, this style of growth experience is predominant, inequality is likely first to worsen in the early stages of development and then to improve. The crossing of the Lorenz curves is demonstrated in Figure 5.9.

The explanation for the crossing in Figure 5.9 is as follows: The poor who remain in the traditional sector have their incomes unchanged, but these incomes are now a smaller fraction of the larger total, so the new Lorenz curve, L_2 , lies below the old Lorenz curve, L_1 , at the lower end of the income distribution scale. Each modern-sector worker receives the same absolute income as before, but now the share received by the richest income group is smaller, so the new Lorenz curve lies *above* the old one at the higher end of the income distribution scale. Therefore, somewhere in the middle of the distribution, the old and new Lorenz curves must cross.¹⁸

These three typologies offer different predictions about what will happen to inequality in the course of economic growth. With modern-sector enrichment, inequality rises steadily, while under traditional-sector enrichment, inequality falls steadily. Under modern-sector enlargement, inequality first rises and then falls;¹⁹ if this admittedly highly stylized process of development were occurring, we would not be concerned about the temporary rise in inequality, because in addition to being temporary, it would be reflecting a

FIGURE 5.9 Crossing Lorenz Curves in the Modern-Sector Enlargement Growth Typology

process in which citizens are, one by one, achieving incomes above the absolute poverty line.²⁰

These observations tell us that we have to qualify our conclusion that a rise in inequality is inherently bad. In some cases, inequality may increase on a temporary basis due to causes that will eventually make everyone better off and ultimately lower inequality. However, with modern-sector enrichment growth, the increase in inequality is not later reversed, and the poor do not escape their poverty.²¹ So we need to be careful about drawing conclusions from short-run changes in economic statistics before we know more about the underlying changes in the real economy that have given rise to these statistics. The process of modern-sector enlargement growth suggests a possible mechanism that can give rise to Kuznets's "inverted-U" hypothesis, so we turn to this question next.

Kuznets's Inverted-U Hypothesis

Simon Kuznets suggested that in the early stages of economic growth, the distribution of income will tend to worsen; only at later stages will it improve.²² This observation came to be characterized by the "inverted-U" **Kuznets curve** because a longitudinal (time-series) plot of changes in the distribution of income—as measured, for example, by the Gini coefficient—seemed, when per capita GNI expanded, to trace out an inverted U-shaped curve in some of the cases Kuznets studied, as illustrated in Figure 5.10.

Explanations as to why inequality might worsen during the early stages of economic growth before eventually improving are numerous. They almost always relate to the nature of structural change. Early growth may, in accordance with the Lewis model, be concentrated in the modern industrial sector, where employment is limited but wages and productivity are high.

Kuznets curve A graph reflecting the relationship between a country's income per capita and its inequality of income distribution.