

Draft – Please do not cite

August 2008

Development Economics: Theory, Empirical Research and Policy Analysis

Julie Schaffner

Chapter 5

Poverty, Inequality and Vulnerability

Most development analysts would agree that the objective in development is broader than mere growth in average income, in at least two ways. First, what matters in development is not just what happens *on average*, but what happens to every person, paying special attention to the poor. Second, what matters is not just what happens to people's current income, but also what happens to dimensions of current living standards that vary independently of income, and what happens to peoples' hopes and fears regarding the future. Thus it is important to complement measures of economic growth with additional measures of development success at the aggregate level, which illuminate these other dimensions of aggregate change.

This chapter considers a broader array of aggregate development outcomes by examining the definition and measurement of three ills that we might wish to reduce through successful development: poverty, inequality and vulnerability. After defining each of these concepts, and discussing the strengths and weaknesses of commonly employed measures, it begins to examine the complex inter-relationships between economic growth, poverty, inequality and vulnerability. It argues that while economic growth creates resources that are useful for reducing the ills of poverty, inequality and vulnerability, it does not guarantee that such reduction will take place. The chapter also raises the currently debated possibility that policy efforts to reduce poverty, inequality and vulnerability might even speed economic growth. The chapter thus suggests the importance of avoiding the extremes of both exclusive attention to economic growth and exclusive attention to short-term poverty reduction when designing and evaluating development policies. Finally, it argues that concerns with poverty, inequality and vulnerability, and the desire to identify policies that will reduce poverty and vulnerability while also promoting growth, provide additional reasons (beyond those already presented in Chapter 3) for taking a more careful, "micro" look at the many decisions that households in developing countries must make and of what happens when they interact with each other in markets and non-market institutions. Parts III and IV of the text equip the reader with analytical tools to guide such study, while Part V applies the tools in the analysis of a wide range of development policies and programs.

Poverty

Economic growth is valuable for raising well-being, but it does not necessarily make right the wrong we sense when we learn of millions of children falling asleep hungry every night and dying of easily preventable childhood illnesses, or when we smell the raw sewage running in ditches between the shacks of the developing world's many shanty towns. The instinctive desire to do something for human beings living in

deprivation is given shape and substance in policy debates through the identification of “poverty reduction” as a component of the development objective. In this section we discuss the definition and measurement of poverty. Later in the chapter we examine the empirical relationship between growth and poverty reduction.

Defining poverty

Defining poverty in a general way is easy: being poor means not having enough. Using the language of Chapter 2, we might also say that **poverty** means living at a level of well-being that falls below some minimally acceptable level. Construction of more specific definitions and measures of poverty is more complicated, because this requires wrestling with three questions, to which people with differing values and beliefs might have different answers. The first two pertain to the evaluation of poverty at the individual or household level:

- What dimensions of living standards should be considered when evaluating whether someone’s well-being is above or below a minimally acceptable level? and
- What constitutes the minimally acceptable level of these living standards?

The third question has to do with the aggregation of information on poverty for many people into a single summary measure for a group or region:

- What formula should we use for constructing an aggregate poverty statistic as a function of the poverty levels of many individuals?

The following sub-sections address each of these questions in turn.

Selecting dimensions of living standards for use in constructing poverty measures

The first step in defining poverty is to identify the specific elements of living standards that are relevant for evaluating whether well-being falls short of a minimally acceptable level. The questions that must be addressed here are the same as those addressed in Chapter 2, when we considered the measurement of well-being more generally. There we saw that commonly used measures of **real per capita income** within households offer useful summaries of peoples’ current ability to purchase goods and services available in markets, but can fail to capture differences in well-being across households that arise out of differences in their access to goods or services that are rationed or cannot be purchased at all, or out of differences in their future prospects. Measures of income in short reference periods also serve as poor measures of households’ typical ability to consume goods and services when income fluctuates widely and households engage in borrowing, saving and other activities with the aim of achieving rates of consumption expenditure that are more constant over time than are their income receipts. Measuring only the average of income per person for the entire household also fails to shed light on the distribution of well-being across males and females, young and old, within households. Measures of **real per capita consumption expenditure** within households offer

somewhat better summaries (than do measures of per capita income) of peoples' typical ability to purchase goods and services when they "smooth" consumption expenditure over time relative to income fluctuations, but still fail to capture many inter- and intra-household differences in well-being.

We may often wish to supplement these summary measures of well-being with multiple **direct measures** of the enjoyment by household members of many features of life circumstances that they care about – including what goods and services they consume, how long they work, and what housing and environmental conditions they enjoy, paying special attention to items that cannot be purchased in well-functioning markets and items that can be measured at the individual level. Unfortunately, while we may wish such measures to capture differences across households in their ability to enjoy a wide variety of life qualities, they may instead reflect differences in taste. For example, one household may consume less meat than another, not because it faces inferior opportunities for consuming meat, but because its members are vegetarians. We may also wish to supplement measures of current income or consumption expenditure with measures related to **household assets**, which tend to render a household's members better prepared for the future.

Choosing poverty lines

The next step in defining poverty is to ask: What is a minimally acceptable standard of living? That is, how do we define the **poverty lines** to which we will compare our measures of well-being when drawing a conclusion about whether a household is poor (and how poor it is)? When describing this step in the definition of poverty, it is common for textbooks to draw a sharp distinction between **absolute** and **relative** notions of poverty. Absolute poverty lines are defined with respect to standards that are of unchanging importance to human experience. For example, the starkest definition of absolute poverty compares an individual's nutrient intake to the minimum physiologically necessary for sustaining human life. At the other extreme, pure notions of relative poverty find living standards unacceptable if they fall too far below the mean or median levels within a region or country. Such poverty lines are higher in wealthier regions and rise over time during periods of economic growth, reflecting beliefs that the severity of deprivation is related not just to the conditions in which people find themselves, but also to how their living conditions compare to those of others in the same society. In practice, most practical definitions of poverty are **hybrids**, recognizing the importance both of unchanging physiological requirements and of components that are shaped by prevailing economic and social conditions, as we will see.

Defining poverty lines for summary measures based on household income or consumption expenditure. Whether one's concept of poverty is absolute, relative or a hybrid has important implications for the determination of traditional income- or consumption expenditure-based poverty lines. For the identification of absolute poverty based on physiological survival requirements, researchers seek to identify the level of income or consumption expenditure required to purchase a package of food that just provides adequate nutrition for sustaining life and activity. This turns out to be a complex undertaking, because the minimal package of *nutrients* can be obtained from

many different combinations of *foods*, each combination with a different total cost. Given data on food prices in any region, researchers can identify the bundle of foods that provides minimal nutrient levels at lowest cost there, and can set the poverty line equal to the cost of that bundle. Unfortunately, even though the poor tend to seek out relatively low-cost ways of satisfying their nutritional needs, their consumption choices are also naturally influenced by tastes and cultural perceptions of good and bad foods. They may thus choose bundles of food that differ from the lowest cost bundle, and may thus require more income to achieve minimal nutrition. Thus poverty measures involving poverty lines tied to minimum cost approaches for achieving minimum nutrient requirements may tend to understate nutrient poverty.

Most official income- or expenditure-based poverty lines build in not only enough money to purchase minimal nutrient consumption, but also money for clothing, shelter and other needs. These components of poverty line construction are often determined in much more arbitrary ways, and often (implicitly or explicitly) reflect the idea that poverty is to some extent relative. For example, after determining the expenditure necessary to buy a minimum standard of nutrition, national statistical organizations may add in to their poverty line calculation some amount of expenditure that reflects the cost of “participating in every day life of society”. This tends to introduce relative notions of poverty into the poverty line calculation, because in some regions one needs only clothes and shoes to participate fully in society, while in other (more affluent) regions, one may also need a minimum acceptable standard of housing, as well as access to electricity and telephone services. For an interesting discussion of how politics enters into the determination of official poverty lines, see Deaton (2006).

Even though sometimes based (in principle) on absolute notions of poverty, in practice countries’ own poverty lines (in terms of real per capita household income) tend to rise with country income level, from around \$300 in South Asia to around \$4000 per person in the U.S. (where the poverty line is more often stated in per family terms, e.g. \$17,000 for a family of four). As a result, poverty rates based on countries’ own poverty lines are not strictly comparable. It has become common in recent years to make cross-country poverty comparisons based on poverty lines of \$1 and \$2 per day, which are similar to official poverty lines chosen in the world’s low and middle income countries, respectively, as is the case in Table 1.2.¹

Poverty lines for use with direct measures of life circumstances. Given the potential weaknesses of poverty measures based solely on income or consumption expenditure, it is often useful to construct complementary measures of whether households’ circumstances fall short of minimally acceptable levels along a variety of specific dimensions, especially dimensions along which income plays only a small role in determining what the household enjoys. Often reasonably natural poverty lines suggest themselves. For example, it would be natural to compare measures of calorie consumption and micro-nutrient consumption to well accepted dietary standards for

¹ See Chapter 1, footnote 3, for a discussion of methods for converting poverty lines in local currency into poverty lines expressed in U.S. dollar equivalents. For a more detailed discussion of technical difficulties with poverty comparisons across countries and over time, see Deaton (2001).

adequate nutrition, and to compare measures of children's weight and height to internationally accepted levels (by gender and age) used to identify children suffering from short- and long-term malnutrition (see the third column of Table 1.2).

For many other living standard measures of interest there is little choice in how to draw poverty lines, because the measures are dichotomous, meaning that they can take only two values. School-aged children either are or are not in school; households either do or do not have access to improved water and sanitation; and individuals either did or did not become ill in the last 7 days, either did or did not become the victim of violent crime in the last 12 months, and either do or do not feel that they have some say in village investment decisions. For dichotomous measures the only choice is to consider those without the desired characteristic as poor. The statistics reported in the last three columns of Table 1.2 can be thought of as indicating the shares of the relevant populations that are not poor along such dimensions.

Distinguishing the “very poor” from the “poor”. Increasingly policymakers and policy analysts are drawing not just one poverty line, but two or more, with the aim of identifying populations at different levels of poverty. For example, the \$2 per day poverty lines identifies those who would be considered poor by middle income country standards, while the \$1 per day poverty line identifies people in deeper levels of poverty, who would be considered poor even by the standards in very low income countries.

Distinguishing those in “persistent” poverty from those in “transitory” poverty. Another distinction whose significance has become widely appreciated in the last decade is the distinction among those who are poor today (according to standard income-based measures) between those who are likely to remain in poverty only a short time and those who are likely to remain poor for a long time. The motivation for this distinction is the observation that the predictable and unpredictable fluctuations described in Chapter 2 may cause households' current living standards to vary greatly from year to year and month to month. The incomes of households with many assets will tend to vary around a high average, while the incomes of households with few assets will tend to vary around a low average. Among those who are income-poor in any one year, some have moderate asset holdings but are experiencing an unusually bad year. They are likely to exit poverty soon and may be called the “temporarily” or “transitorily” poor. Others among the current income poor, however, have few assets and are likely to remain poor for a long time. They may be called the “permanently”, “chronically” or “persistently” poor.

Identifying who among today's poor whose poverty is **persistent** and **transitory** is a difficult task, and no consensus has been achieved regarding the best approaches for making this distinction in practice. One suggestion is to complement standard income- or consumption expenditure-based measures of current poverty with “asset-based” measures that distinguish between those whose asset holdings are and are not great enough to provide them with income above the income-based poverty line in “average” years or in some minimum percentage of years. Drawing such a line is complicated because peoples' “assets” may take many forms, and an approach must be devised for turning data on a wide variety of assets into a single asset index. Carter and Barrett (2006)

propose a way of doing this, and suggest drawing an asset-based poverty line at the level of this index that should provide households with average income at least as great as the income-based poverty line.² The science of constructing asset-based poverty measures is still in its infancy, however, and no consensus has yet emerged.

Choosing formulas for summary measures of poverty at the aggregate level

Once we have chosen a measure to use in assessing each individual's well-being, and once we have constructed the relevant poverty line, we can use survey or census methods to gather data describing the entire distribution of that well-being measure across individuals in a country or region of interest, and can identify how far above or below the poverty line each individual lies. For example, if we choose to use per capita income within the household as our measure of individuals' well-being, we can collect data that describe the distribution of this income measure across all members of the population, and can compare this information to an income-based poverty line, to construct measures of poverty for every person in the population. With this mass of detailed information, we would be able to answer such questions as:

- How many people are poor?
- What fraction of the total population is poor?
- What is the average depth of poverty? (That is, how far below the poverty line is income of the average poor person?)
- How deep is the deepest poverty?
- How much variation is there in the depth of poverty among the poor?
- What is the total amount by which the incomes of the poor fall short of the poverty line?

Having information to answer this array of interesting questions is valuable, but the mass of information containing such diverse information can be overwhelming and difficult to talk about in policy conversations. Thus we will often wish to employ **aggregate** statistics that offer meaningful and succinct summaries of important features of this amassed information. In what follows we first offer a graphical depiction of the full range of information we wish to summarize with an aggregate poverty statistic. We then define many commonly used aggregate poverty statistics and discuss the values or priorities built in to each.

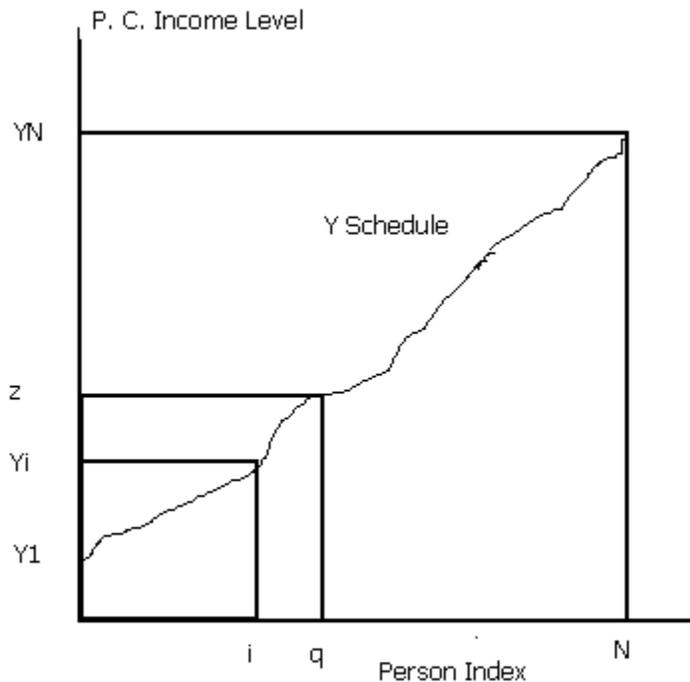
Graphing the information we wish to summarize. Figure 4.1 illustrates a useful way of presenting the information we wish to summarize. For clarity in exposition, let's assume we have decided to use the real per capita income within a person's household as a

² The asset-based poverty line described here is what Carter and Barrett (2006) call a "static asset poverty line". They also introduce a notion called the "dynamic asset poverty line" which attempts to distinguish between the poor whose very low level of assets render them caught in "poverty traps" (a concept introduced in Appendix 3A and later in this chapter), and those whose assets are great enough that they can profitably save and invest, and thus have hope of eventually raising themselves out of poverty through investment.

measure of that person’s well-being. (We could use the same approach for summarizing information on any other continuous measure of well-being.)

Let’s imagine ordering the N individuals in the population from 1 to N , with individual 1 being the poorest and individual N being the richest. We will refer to each individual’s position in this rank ordering as his or her “person index.” Thus the poorest person has person index 1, and the i th poorest person has person index i . We will use Y_i to denote person i ’s income level. The horizontal axis in Figure 1 indicates the range of person indices. As we move to the right along the horizontal axis, we are moving up the income distribution to people with higher and higher incomes. Each point along the horizontal axis references a particular person in the population. The upward-sloping Y schedule associates each individual’s person index (i) with his or her per capita income (Y_i). That is, we can find person i ’s income by moving along the horizontal axis to the point marked i , and then observing the height of the Y schedule directly above that point.

Figure 4.1



The height and shape of the Y schedule tell us everything there is to know about incomes in this population. The average height of the Y schedule indicates average income, while the slope of the Y schedule indicates how rapidly incomes rise as we move from the poorest to the richest individuals. If the Y schedule is quite flat, incomes do not vary much around the mean. Where the Y schedule is steeper, income inequality is greater.

If we decide that z is the minimally acceptable level of per capita household income, then the line at height z in the diagram is a literal “poverty line.” We know that person i is

poor if her income Y_i is less than z . The poverty line hits the Y schedule at person index q , indicating that q of the N people in this population are poor. The extent to which a poor person's income, Y_i , is below the poverty line, $(z - Y_i)$, is a measure of the depth of that person's poverty, in units of currency. We call $z - Y_i$ the **income gap** for individual i . Because income may be measured in a currency unit with which we are not very familiar, we may find it more convenient to express the gap between an individual's income and the poverty line in proportional terms, as in $(z - Y_i)/z$. This **proportional income gap** is expressed as a fraction of the poverty line. For example, if this number equals 0.30, it indicates that the person's income falls 30 percent below the poverty line.

Defining frequently-used poverty statistics. Each of the most commonly used summary measures of poverty (summarized in Box 4.1) captures some, but not all, of this information. Because the formulas for most of these measures require summing or averaging across a large number of people, it is convenient to express them using summation notation. For a brief refresher on the definition and uses of summation notation, see problem 1 at the end of the chapter.

Box 4.1 Summary of Poverty Measure Formulas

In preparation for defining common poverty measures and studying the relationships between them, let's let the individuals whose poverty is to be summarized be ordered by income level (or by level of some other measure of living standards) and numbered from 1 to N , with individual 1 being the poorest and individual N being the richest. Individuals 1 through q have incomes below the poverty line, z , and are thus considered poor, while individuals $q+1$ through N are the non-poor. Let Y_i be the income level of individual i . We can now define the following poverty measures:

$$\text{Headcount Ratio} = H = q/N$$

$$\text{Total Income Gap} = TYG = \sum_{i=1}^q (z - Y_i)$$

$$\text{Average Proportional Income Gap Among the Poor} = APYG = \frac{1}{q} \sum_{i=1}^q \left(\frac{z - Y_i}{z} \right)$$

$$\text{Poverty Gap Index} = PG = \frac{1}{N} \sum_{i=1}^q \left(\frac{z - Y_i}{z} \right)$$

$$\text{Squared Proportional Income Gap Index} = P_2 = \frac{1}{N} \sum_{i=1}^q \left(\frac{z - Y_i}{z} \right)^2$$

Added insight into these measures can be gained by observing ways in which the formulas may be re-expressed. For example, the Total Income Gap may be re-expressed as

$$TYG = (z - Y_1) + (z - Y_2) + \dots + (z - Y_N) = zq - \sum_{i=1}^q Y_i .$$

This shows that the *TYG* is just the difference between the total amount of income the poor would have if all of their incomes were just brought up to the poverty line, and the total amount of income they actually have.

The Poverty Gap Index may be re-expressed as

$$PG = \frac{q}{N} \left(\frac{1}{q} \sum_{i=1}^q \left(\frac{z - Y_i}{z} \right) \right) = H * APYG$$

This shows that the Poverty Gap can be calculated as the product of the Headcount Ratio (*H*) and the Average Proportional Income Gap Among the Poor (*APYG*). It thus takes into account both the average severity of poverty among the poor and the prevalence of poverty in the population.

H, *PG* and *P₂* are three members of what is called the Foster-Greer-Thorbecke or FGT class of poverty indices, all of which rank poverty on scales from 0 to 1. (Sometimes they are multiplied by 100 and presented on a 0 to 100 scale.) The class is defined by the general formula

$$P_\alpha = \frac{1}{N} \sum_{i=1}^q \left(\frac{z - Y_i}{z} \right)^\alpha$$

and individual members of the class are defined by the value given to α . The parameter α is the power to which the proportional income gaps of the poor are raised in calculating the overall measure of poverty among the poor. When α equals 0, the contribution to the aggregate measure is 1 for each poor individual. This means that we sum up *q* 1's and then divide by *N*. That is, it equals the Headcount Ratio. For this reason, *H* is sometimes called *P₀*. Each poor person contributes the same amount to this aggregate measure, no matter how poor he or she is. A one dollar increase for a poor person has no impact on this measure, as long as it does not move the poor individual out of poverty.

When α equals 1, the formula yields the Poverty Gap Index (PG), which is sometimes called *P₁*. The contribution to this aggregate measure for each poor individual is equal to his proportional income gap. Any one dollar increase in a poor person's income reduces this measure, but the effect is the same whether the dollar goes to the poorest person or to the least poor person among the poor.

When α equals 2, the formula yields the *P₂* measure defined above. The contribution to this aggregate measure for each poor individual is the square of his proportional income gap. The effect of squaring is to magnify the contributions of the poorest, whose incomes are furthest away from the poverty line. Now any one dollar increase in a poor person's income reduces this poverty measure, but the reduction is greater when the dollar is given to a very poor person than when it is given to a less poor person. This measure is thus sensitive to the distribution of poverty among the poor, and thus to the prevalence of those in very deep poverty among the poor, as well as to the average level of poverty and the incidence of poverty in the population. Setting α to a value greater than 2 would place even more emphasis on what is happening to the poorest among

the poor.

The **Headcount Ratio** is the fraction (sometimes expressed as a percentage) of households in the region with incomes below the poverty line. It measures the prevalence of poverty in the population of concern, but is not sensitive to variation in the severity of poverty. That is, regardless of the level of income enjoyed by the average poor person or the poorest of the poor, the headcount measure is the same so long as the total number of poor and the total size of the population are the same.

The **Total Income Gap** is the total amount of money that would be required to bring every income up to the poverty line (if it could be directed perfectly and without cost to exactly the people who need it). Graphically, it is equal to the area between the z line and the income line. It tells us the order of magnitude of the cost of eliminating poverty, but does not tell us about the number of people involved, the prevalence of poverty, the average severity of poverty, or the diversity of poverty levels among the poor. The total income gap could be large either because a very large number of people are slightly poor or a smaller number of people are extremely poor.

A measure that captures the average severity of poverty among the poor is the **Average Proportional Income Gap** (APYG), which is just the simple average over all the poor of their proportional income gaps. It can also be calculated by taking the Total Income Gap and dividing by $z \cdot q$. An APYG of .35 means that on average the incomes of the poor fall short of the poverty line by 35 percent. While this tells us about the typical depth of poverty, it is insensitive to the number of the poor, the prevalence of poverty in the population, and the diversity of poverty levels among the poor.

The **Poverty Gap Index** (PG) averages proportional income gaps across everyone in the population, treating the non-poor as having income gaps of zero (because they need nothing to bring their incomes up to at least the poverty line). A PG of .05 indicates that the average person in the population at large has a relative income gap of 5 percent. The Poverty Gap Index can be calculated as the product of the Headcount Ratio and the Average Proportional Income Gap. It is sensitive to both the prevalence of poverty and its average severity. It is still insensitive, however, to differences in the levels of income among the poor, and thus to differences in the fractions of the poor who are in very deep poverty.

The **Squared Proportional Income Gap Index** also averages an individual-level poverty measure across all members of the population. Instead of letting the individual's contribution to the aggregate measure be his proportional income gap, however, it lets the individual's contribution be the square of his proportional income gap. The effect of this squaring is to magnify the role played by the depth of an individual's poverty. While an individual whose income is 10 percent below the poverty line adds .01 to the sum, an individual whose income is 50 percent below the poverty line adds .25 to the sum. The ratio of the poverty gaps is 5 to 1, but the ratio of their contributions to this aggregate measure is 25 to 1. The measure thus places greater weight on what is happening to the

incomes that are furthest below the poverty line, and is thus sensitive to changes in the diversity of poverty levels among the poor, as well as to the prevalence and average severity of poverty.

Selecting poverty measures. Policymakers and policy analysts must sometimes select just one of the many possible poverty statistics for use in identifying poverty reduction priorities or evaluating poverty reduction success. For example, in the interest of replacing discretion and patronage by objective rules, the Mexican government tied its allocation of poverty reduction funding across regions to regional values of the Squared Proportional Income Gap Index. In the interest of promoting accountability and comparability, a donor organization might require all the recipients of its funds to quantify their impacts using specific poverty measures.

As should be clear by now, the selection of a single aggregate poverty measure for use in such situations is a value-laden activity that should be undertaken with care. Different poverty statistics can produce different conclusions regarding poverty reduction priorities across regions. This is illustrated in Table 4.1, which presents Headcount Ratio, Poverty Gap and P2 Index values for selected provinces in Indonesia in 1990. The provinces are listed in decreasing order of the headcount ratio. E. Nusa Tenggara exhibited the greatest poverty by all three measures. But some relative rankings differ depending on the aggregate poverty measure employed. Irian Jaya is less poor than W. Sumatra by the Headcount Ratio, but more poor by the other two measures. Evidently, while the share of the population that is poor is lower in Irian Jaya, the average depth of poverty is higher there. N. Sulawesi is less poor than E. Java by both the headcount and poverty gap measures, but contains more poverty according to the P2 Index, indicating that while the share of the population that is poor *and* the average level of poverty among the poor are lower in N. Sulawesi, the variation in income around the mean for the poor, and thus the share of the poor living in especially deep poverty, must be higher there.

Table 4.1
Poverty Measures for Selected Provinces in Indonesia, 1990
(All measures multiplied by 100 to render them in percentage terms.)

Province	Headcount	Poverty	
		Gap	P2
E. Nusa Tenggara	45.62	11.84	4.12
E. Java	21.8	3.84	0.95
N. Sulawesi	18.79	3.79	1.09
W. Sumatra	13.35	1.71	0.30
Irian Jaya	12.61	2.41	0.63
Indonesia	19.63	3.46	0.87

Source: Bidani and Ravallion (1993).

Different poverty measures can also produce quite different conclusions as to which of several competing programs is “the best.” A program that makes small improvements in

the incomes of many “slightly poor” people whose incomes fall only just below the poverty line may be judged the most successful in reducing poverty as measured by the headcount ratio (provided the individuals’ incomes are boosted enough to raise them above the poverty line). In contrast, a program that makes a dramatic improvement in lives of the some of the severely poor, but raises them up only to milder levels of poverty without raising their incomes above the poverty line, would be rated as completely ineffective by the headcount ratio measure, while it might be judged “the best” when poverty is measured by the Squared Proportional Income Gap measure.

Inequality

We may wish to improve the lot of the worst off in society for two quite different reasons. We may believe it simply isn’t right for human beings to live at less than some minimally acceptable standard of living. In this case, our policy objective is to reduce poverty. We may instead, or in addition, believe it is troubling or wrong for some people to have so much less than others in their society, and similarly wrong for some people to have so much more than others in their society. In this case our policy objective is to reduce inequality.

Often, in practice, the two objectives will provide similar guidance for our activities; and indeed the distinction between poverty and inequality are often blurred in popular discussions. Poverty reduction efforts to raise the well-being of the poor also reduce inequality if the incomes of the non-poor remain unaffected, and even more so if the financing for the poverty reduction efforts is derived by taxing the non-poor.

Despite their overlaps, however, the objectives of reducing absolute poverty and reducing inequality are really quite different, and might lead to very different conclusions about the best routes to development success. For example, if all incomes in an economy are rising, but the incomes of the richest people are rising faster than the incomes of the poorest, then by the criterion of absolute poverty reduction it is a period of success (because the incomes of the poor are rising), while by the criterion of inequality reduction it is a failure. Similarly, if the incomes of the richest fall, while the incomes of the middle class rise and the incomes of the poor stay the same, the period would be one of success in reducing (at least some measures of) inequality, though one of failure in reducing poverty. (Try picturing each of these scenarios using a graph like Figure 4.1. What would the change in the Y schedule look like?)

Defining and measuring inequality

In general terms, **inequality** means some people having less than others. That is, inequality has to do with the differences within a group in levels of living standards. Constructing a specific measure of inequality requires answering two questions:

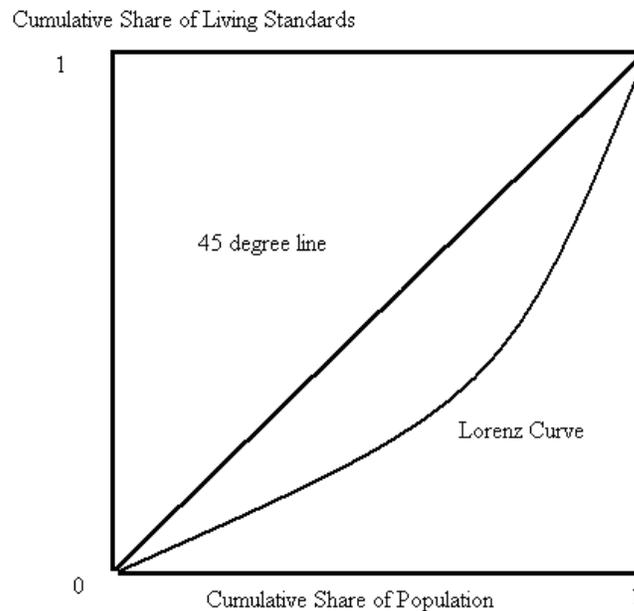
- What dimensions of living standards should be considered when comparing living standards across members of the group?

- What formula should be used for summarizing the relevant features of the distribution of living standards in the group?

Answering the first question here involves much the same choices as were relevant in the choice of living standards for use in the construction of poverty statistics. For the purposes of exposition, we will assume here that we are content to measure living standards using a single measure: per capita household income, and will concentrate on the second question.³

Graphing the information we wish to summarize. The information we must summarize when constructing measures of inequality is the same information on income levels for every individual in the population that we graphed in Figure 4.1. That figure emphasizes the absolute levels of living standards and their relationship to the poverty line. Here we will prefer to graph the same information in a way that emphasizes the *relative* shares of individuals in the distribution of income. A diagram often used for this purpose is a **Lorenz curve** diagram, such as that presented in Figure 4.2.

Figure 4.2
Lorenz Curve



³ We focus here on the **size distribution** of income. That is, we are concerned with the shares of income held by groups defined by their level of income. This may be contrasted with the **functional distribution** of income, which is concerned with the shares of income held by groups defined by their main sources of income: laborers, farmers, owners of capital and natural resources. We assume that policymakers' ultimate interest (if they are concerned with inequality at all) is with the size distribution of income, but we will later make great use of functional distinctions in analyzing policy impacts on distribution. Where we know, for example, that the poor earn their income primarily as low-skilled laborers, then understanding how policy impacts the relative remuneration of unskilled labor is useful for understanding its ultimate impacts on the size distribution of income.

To construct a Lorenz curve, we order everyone in the population from lowest income to highest income. We then ask: What percent of total income is enjoyed by the bottom one percent of the population? What percent of total income is enjoyed by the bottom two percent of the population? And so on all the way up to: What percent of total income is enjoyed by 100 percent of the population? (The answer to this last question is, of course, always 100 percent.) We plot the answers to this series of questions directly in the Lorenz curve. The horizontal axis in the Lorenz curve plots the percent of population (ranged from poorest to richest), while the vertical axis plots the percent of total income. Both percentages range from zero to 100, thus the Lorenz curve will always fall within a square box.

What would the Lorenz curve look like if the distribution of income were perfectly equal? Each successive one percent of the population should enjoy exactly one percent of total income. Thus the first 10 percent of the population should enjoy 10 percent of total income, and the first 50 percent of the population should enjoy 50 percent of total income. As we move to the right by one percentage point of population, we should move up by exactly one percentage point of income, and we obtain a straight line of slope 1, starting at the origin and ending at the northeast corner of the Lorenz curve square. It is useful to draw such a reference line into our Lorenz curve diagram, as is done in Figure 4.2. This perfect equality reference line is often called the **45 degree line**, because it rises up at an angle of 45 degrees relative to the horizontal axis.

What would perfect inequality look like? If one person enjoyed all the income in the economy, then the percent of income enjoyed by the first 20, 40 or 90 percent of the population would be zero. The Lorenz curve would be flat along the horizontal axis, until it shoots up to 100 percent at the very final point, indicating that 100 percent of the population -- which includes the one rich person -- does indeed enjoy 100 percent of the income. Thus perfect inequality would be depicted by a backwards L-shaped "curve" following the bottom and right side of the Lorenz curve box.

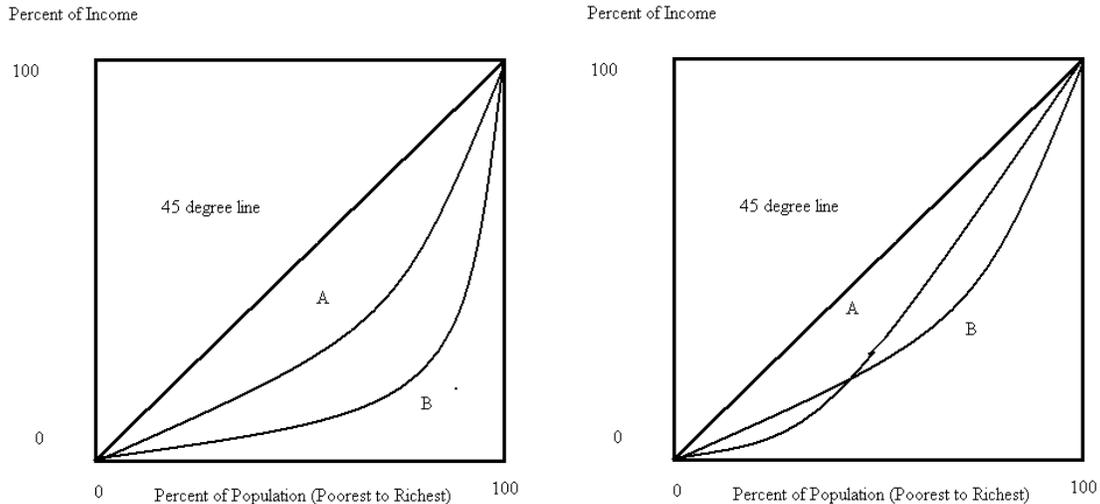
For distributions between the extremes, we expect Lorenz curves to look something like that shown in Figure 4.2. It starts off with a slope lower than 1, indicating that the bottom percentiles of the population command less than 1 percent of total income each. But we know that ultimately 100 percent of the population must enjoy 100 percent of the income, so the curve must eventually reach the northeast corner. This requires that eventually the Lorenz curve become steeper, to a slope greater than 1, indicating that the top percentiles of the population command more than 1 percent of income each.

The shape of the Lorenz curve offers a thorough characterization of inequality in the distribution of income in the economy. It is silent, however, as to the mean level of income. Two economies can have vastly different mean and total levels of income but identical Lorenz curves.

In some cases the comparison of Lorenz curves across countries or time periods leads to a pretty obvious conclusion about where inequality is worse. In Figure 4.3a Lorenz curve

B lies everywhere further from the 45 degree line than Lorenz curve A. This indicates that no matter what X you pick, the income share of the lowest X percent of the population is smaller in country B than in country A; everyone would agree that inequality is worse in B than in A. The Lorenz curve for A is said to **dominate** that for B. In other cases, however, Lorenz curves alone are not enough to draw a conclusion about the relative significance of inequality. In Figure 4.3b the Lorenz curves cross. If we look at the income share of the lowest 10 percent of the population, country A looks more unequal; but if we look at the income share of the lowest 50 percent of the population, country B looks more unequal. In cases like this, we can only draw conclusions as to the relative severity of inequality if we introduce values that identify which particular features of the income distribution should be given most weight when making comparisons. The next section turns to this problem.

Figure 4.3
Comparing Lorenz Curves



Summarizing Inequality in a Distribution of Living Standards

When defining and measuring inequality, it will often be convenient to summarize salient features of the information presented in the Lorenz curve with a single number. For this we require a formula for a statistic that is a function of some or all of the information plotted in the diagram. We hope that when comparing the income distributions in two groups for which Lorenz curves do not cross, the statistic will always take a higher number for the Lorenz curve that lies further from the 45 degree line. When Lorenz curves cross, a formula will allow us to draw a conclusion about where inequality is worse, and we hope to construct the formula in a way that places emphasis on the features of the distribution that matter most to us.

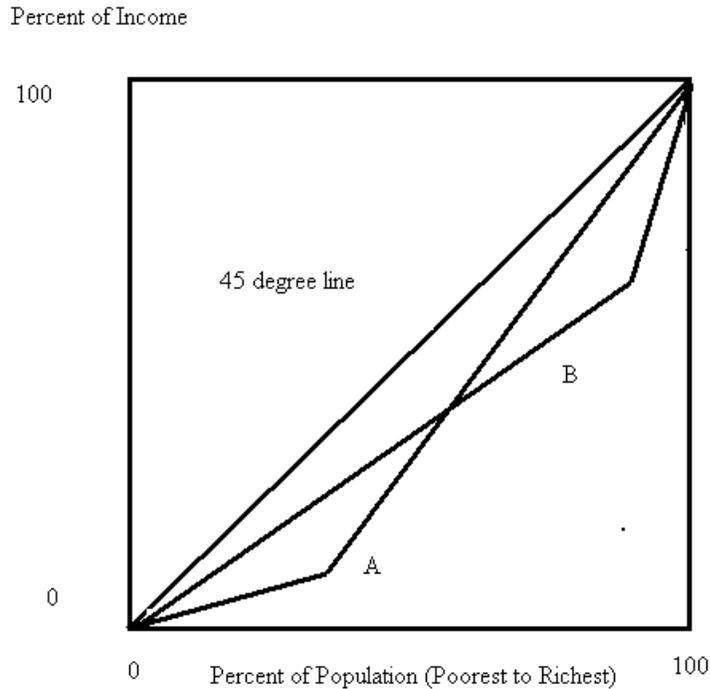
Income shares. Several very simple summary statistics make use of only a small portion of the information contained in the Lorenz curve, but are easy to interpret. It is common to describe income distributions by reporting the **shares of total income** enjoyed, for example, by the bottom 20 percent and top 10 or 5 percent of the population. We may also combine this information into a single statistic, by taking the ratio of the income share of the top 10 percent to the income share of the bottom 40 percent. The higher is this number, the greater the skew in the distribution of income toward the rich.

Gini coefficient. A popular measure of inequality that makes use of all the Lorenz curve information is the **Gini coefficient**, which may be defined as the ratio of two areas in the Lorenz curve diagram. The numerator is the area between the Lorenz curve and the 45 degree line. The denominator is the area of the entire triangle lying below the 45 degree line. If the distribution were perfectly equal, then the Lorenz curve would lie on the 45 degree line, the numerator would be zero, and the Gini coefficient would be zero. If the distribution were perfectly unequal, the Lorenz curve would follow the horizontal axis until it shoots up along the east end of the box, the numerator would be equal to the denominator, and the Gini coefficient would be equal to one. Thus the Gini coefficient lies between 0 and 1, with higher values indicating higher inequality.

When Lorenz curves don't cross, the areas between the 45 degree line and the Lorenz curve will be greater for the Lorenz curve lying further from the 45 degree line. Thus the Gini coefficient satisfies the logical requirement that it gives a higher measure of inequality to the distribution characterized by a Lorenz curve that is dominated by another. Even when Lorenz curves cross, the Gini coefficient calculation will often provide a clear ranking of the distributions. Whichever Lorenz curve cuts off a larger area between the curve the 45 degree line will earn a higher inequality ranking.

Will everyone agree with every ranking produced by the Gini coefficient? No. Consider the case of the two Lorenz curves drawn in Figure 4.4. They have been drawn so that the area between the 45 degree line and the Lorenz Curve is the same for both curves; thus their Gini coefficients are the same. But would everyone agree that inequality is equally severe in the two cases? No. Lorenz curve A describes a group in which the poorest 1/3 equally share 1/9 of total income, while the other 2/3's equally shares the remaining 8/9. Lorenz curve B describes a group in which the poorest 2/3 equally share 1/2 of total income, while the remaining 1/3 of the population equally shares the other half. Many (though perhaps not all) would consider inequality to be worse in population A than in population B, because many of us instinctively place greater emphasis on the circumstances of the poorest. The Gini coefficient formula does not build in that greater sensitivity.

Figure 4.4
Lorenz Curves with Equal Gini Coefficients



Despite its imperfections, the Gini coefficient is a frequently used measure of inequality. While in principal it can vary from 0 to 1, in practice, when describing income and wealth distributions, it tends to vary from about .2, reflecting very low inequality, to .8, reflecting a very high degree of inequality. Table 4.2 provides example Gini coefficient and income share reports drawn from diverse countries around the world.

Table 4.2
Select Income Distribution Statistics, 2004.

Country	Gini Coefficient	Share of income or consumption (Percent)	
		Poorest 20 percent of population	Richest 20 percent of population
Hungary	.24	7.7	37.5
Ethiopia	.30	9.1	39.4
Turkey	.40	6.1	46.7
United States	.41	5.4	45.8
Malawi	.51	4.9	56.1
Brazil	.59	2.0	64.4
Namibia	.71	1.4	78.7

Source: United Nations Development Program, Human Development Reports data site, <http://hdr.undp.org>.

Many other measures of inequality are possible, each embodying somewhat different values regarding the features of income distribution that matter the most. For more on

the theoretical development of inequality measures that satisfy desirable properties, see the readings at the end of the chapter.

Vulnerability

Compared to the terms poverty and inequality, “**vulnerability**” is a new entrant to the development lexicon. Specific quantitative measures of vulnerability (comparable to the aggregate measures of poverty and inequality discussed above) are unlikely ever to achieve wide acceptance, for reasons we will discuss below. It is nonetheless useful to identify vulnerability as a distinct development outcome of potential policy interest, because it acknowledges an important reality of life in developing countries: that the future is uncertain, giving rise to fears of serious reductions in well-being, especially for today’s poor and near poor, but for the non-poor as well. As we saw in Chapter 2, too much rain, too little rain, insect infestations, crop and animal diseases and sudden price changes may reduce farm income. Job loss and breakdowns in supplier and buyer networks may reduce non-farm income. Disease, crime, traffic accidents, domestic violence, conflict, and other hazards may deprive people of well-being along multiple dimensions. Thus many people whose current living standards are above the poverty line nonetheless live in the shadow of deprivation, and many of the poor live in the shadow of utter destitution. These groups may merit special policy attention.

Vulnerability is a more complex concept than poverty or inequality. In light of Chapter 2, we can understand vulnerability as arising out of the conjunction of two sets of circumstances: (1) households’ exposure to **risks or hazards** (of the shocks we described in Chapter 2), and (2) households’ lack of attractive access to **ex ante and ex post coping mechanisms**, through which they might mitigate the effects of the hazards on their consumption and well-being. Ex ante coping mechanisms include efforts to reduce exposure to risk (e.g. non-use of high-yielding but risky crops or use of insecticide-treated bed nets to reduce exposure to malaria), as well as precautionary saving, purchase of formal insurance and cultivation of informal insurance arrangements. Ex post coping mechanisms include the diversion of household labor from activities whose productivity has been reduced by the hazard into other activities and taking out loans to maintain consumption until times are better. Policymakers have reason to concern themselves not only with households’ ultimate vulnerability to fluctuations in consumption, but also with the underlying risks they face and the costly coping mechanisms they may employ to deal with them. Some people may appear invulnerable to shocks, in the sense that their consumption does not fluctuate much, but only because they have made choices that reduce their exposure to risk, at the cost of avoiding investments with the potential to raise them out of poverty.

Defining Vulnerability

In general terms, the vulnerability of concern to policymakers arises when people face “significant risk of significant future reductions in well-being.” This definition makes clear that vulnerability is a forward-looking concept: it pertains not to living standards today, but to hopes and fears regarding future living standards. It is also a probabilistic

concept, applying to entire groups of people facing some probability of loss all the time, even though the feared losses will be realized only for a small fraction of the group or only in a small fraction of time periods. The inclusion of the “significant” qualifiers captures the idea that while every person in the world is at some risk of loss at all times, there are socially unacceptable levels of vulnerability that particularly merit policy concern.

The general definition of vulnerability begs many questions that we would have to answer on route to creating more specific definitions and constructing useful measures. These include:

- To what hazard(s) must the risk relate (e.g. bad weather, job loss, tuberculosis, any hazard that might lead to reductions in consumption)?
- Is the relevant risk an objective probability or a subjective fear?
- How high must the probability be, and over what period of time, to be considered “significant”?
- To what kinds of loss must the hazards lead (e.g. in consumption expenditure, nutrition, disability, asset holdings)?
- Is the “significance” of the reduction in well-being a function of its size (relative, say, to the level of well-being before the reduction) or of the ultimate level to which it reduces well-being?
 - If it is a function of the size of the reduction, by how large a percentage must well-being drop?
 - If it is a function of the ultimate level to which well-being is reduced, what is the threshold level (something like a poverty line) below which well-being must drop?
- Does it matter whether the feared reductions in well-being are likely to be temporary or longer-lasting?

Different answers to these questions could easily generate hundreds of distinct notions of vulnerability. Among the most prevalent notions among social scientists and policymakers are: vulnerability of the non-poor to falling into poverty, vulnerability of the poor and near poor to falling into “destitution” (defined by per capita consumption expenditure falling below a very low threshold) or malnutrition, and vulnerability of anyone (including those with high incomes) to serious disruptions to their way of life, or to disability or death (see Alwang, et al., 2001; Dercon,2005).

Empirical study of vulnerability

If we were to attempt the construction of aggregate vulnerability measures comparable to aggregate poverty measures, we would need to answer all of the questions raised above, figure out how to quantify the resulting notion of vulnerability at the level of individuals or households, and then use formulas to turn all the micro measures into a single aggregate measure. Unfortunately, the forward-looking and probabilistic nature of vulnerability concepts renders such measurement extremely difficult. We cannot observe today the probabilities with which people will be hit by hazards in the future. In fact, we cannot even identify all the people in the past who were at risk of suffering serious

reductions in well-being, because the shocks will have been realized for only a fraction of them. If we wish to define vulnerability in subjective terms, we could ask people how likely they think it is that they will be hit by hazards or consumption reductions of various types, but such questions can be very hard to phrase and very hard to answer. If we wish to define vulnerability in more objective terms, we would need to estimate in some way the probabilities with which people will be hit by various hazards, and how likely they would be to mitigate the consumption effects of those hazards through the use of various coping mechanisms. This would require estimation or calibration of complicated quantitative models. It is difficult to imagine the achievement of widespread agreement regarding the many details of the models used to generate such estimated probabilities.

Even without constructing direct measures of vulnerability itself, however, researchers may undertake empirical work that sheds useful light on vulnerability, by helping to quantify the probabilities with which households in various groups or geographic locations are likely to be hit by bad weather, malaria or other hazards; the availability and feasibility of various coping mechanisms for responding to shocks; and the fractions of groups that have been hit by significant shocks in recent years.

An important source of information on vulnerability is **panel datasets**, which provide living standards information for a sample of people, not just at one point in time, but for several or many points in time. Panel datasets are constructed by administering the same questionnaire multiple times, over a series of seasons or years, to the same sample of households. Such data collection efforts are difficult, costly and rare, but the few that have been undertaken have proven useful for gaining insights into vulnerability. They demonstrate, for example, that the incomes of the poor and near poor tend to vary a great deal from year to year. This suggests that many who are not poor today are indeed vulnerable to becoming poor next year.

The figures in Table 4.3 are drawn from a paper by Baulch and Hoddinott (2001), which reviews extant studies about poverty dynamics. The figures for different locations are drawn from different studies performed by different researchers and employing different methods of calculating living standards and poverty, thus they are not strictly comparable across locations. Even so, they are useful for demonstrating that the share of households and individuals that are “sometimes poor” is often much larger than the share of households who are “always poor.” They suggest that in any one year there are many non-poor households who will fall into poverty in subsequent years.

Table 4.3
A Few Illustrative Results Regarding Poverty Dynamics

Location	Length of Study	“Always Poor”	“Sometimes Poor”	“Never Poor”
India	9 years	21.8	65.8	12.4
Zimbabwe	4 years	10.6	59.6	29.8
China	6 years	6.2	47.8	46.0

Pakistan	5 years	3.0	55.3	41.7
----------	---------	-----	------	------

Source: Baulch and Hoddinott (2001).

Panel data such as those reported in Table 4.3 are indicative of vulnerability, but fall short of identifying how many households are truly vulnerable today (to falling into poverty in the next month or year), for several reasons. First, just because one third of the non-poor in one year fall into poverty the next year, this does not mean that all members of the heterogeneous non-poor group faced the identical one-third probability of falling into poverty. Some (e.g. those with many assets and stable income sources) were nearly certain to remain non-poor, some (e.g. those with few assets who had only become non-poor for transitory reasons) were nearly certain their living standards would deteriorate, and other faced a range of probabilities of falling into poverty in between 0 and 1. More detailed data analysis might be useful for identifying the types of non-poor households for whom the probabilities of falling into poverty were particularly high. For example, declines into poverty may have been more frequent for households practicing a certain type of agriculture, with land holdings below a certain size and living in certain regions. An additional problem with using panel data to evaluate vulnerability is that the data will, of course, describe past income fluctuations. The next three years may be very different from the past three years.

Development ends and means

Most development actors would agree that the objective in economic development is to achieve some combination of economic growth and reductions in poverty, inequality and vulnerability. Economic growth creates the potential for sustained and widespread improvements in well-being. Reductions in income poverty or inequality ensure that the income gains associated with growth are distributed in a desirable manner, and reductions in non income-based measures of poverty and in vulnerability demonstrate that key non-income dimensions of living standards are improving along with incomes.

Development actors differ, however, in the priorities they place on these diverse sorts of improvement. Some would be more willing than others to trade off reduced growth over the next five years in exchange for more rapid and widespread poverty reduction during the same period. Some would also be more willing than others to trade off reduced improvements in income-based poverty measures for more rapid reductions in non-income measures of poverty or in vulnerability. Development actors differ even more in the priority they place on reductions in inequality (relative to growth and poverty reduction) when evaluating development success. For a few, reduced inequality is critical to the definition of the better world they pursue, while many others would consider development successful even when inequality is rising, as long as the economy is growing and poverty is falling. (Indeed, in this text we will tend to emphasize poverty reduction over inequality reduction, reflecting the much greater consensus on the importance of the former over the latter.) Naturally, such differences in **values** can lead to profound differences of opinion regarding development strategies.

Differences in values are not the only source of policy disagreement, however. Even actors with identical values, who are thus in complete agreement about the *ends* of development, might nonetheless disagree on the best *means* for achieving development, because they may differ in **beliefs** about how the world actually works. More specifically, they may disagree about the likely impacts on growth, poverty, inequality and vulnerability of the various policies they might pursue.

One broad disagreement of enduring importance in development debates concerns the inter-relationships between economic growth, inequality and poverty. If rapid economic growth almost always bring with it significant reductions in poverty, regardless of the policies employed to achieve the high growth, then the most effective way to pursue broad-based development may be simply to pursue policies that are good for growth. If, in addition, explicit policy efforts to reduce poverty and inequality tend to slow growth, then policy makers would have even stronger reason to concentrate on speeding growth, while letting poverty and inequality “take care of themselves.” If, however, economic growth is a highly variegated phenomenon that does not always deliver poverty reduction, and especially if explicit efforts to reduce poverty and inequality might even speed growth, then identifying the best development strategy is a more complicated task. Policymakers must then attempt to identify **pro-poor growth** policies from among the wider array of growth policy options, and must also contemplate the use of explicit **safety net policies** focused on reducing poverty, inequality and vulnerability. The next section begins examining these inter-relationships.

Growth, Inequality and Poverty

Mathematical relationships. If we know the average level of income in an economy, and we know the shares of that income going to each group in the economy, then we can work out the level of income enjoyed by any one group, including the poorest. Thus once we have measures of average income and income inequality, income-based poverty measures provide no new information (though they may present the information in a way that adds new insight). Similarly, if we had detailed information on economic growth and on changes in the distribution of income over a period (as well as information on the initial level and distribution of income), we would know what happened to income poverty over the period. More specifically, if we know that average income is rising (i.e. “the size of the pie” is increasing) and inequality is holding constant or falling (i.e. “the share of the pie going to the poor” is staying constant or increasing), then mathematical relationships tell us that income poverty (defined in absolute terms) is falling by at least some measure (because the “size of the piece of pie enjoyed by the poor” is rising). (To see that poverty need not fall by all measures, see problem 4 below.)

Ravallion (2001) offers a global summary of the historical relationships between growth and changes in inequality and poverty, drawing on data for 117 “spells”, or periods between pairs of comparable household surveys, pertaining to 47 developing countries in the 1980s and 1990s. He divides the spells up into four groups, depending on whether average income is rising or falling, and whether inequality as measured by the Gini coefficient is rising or falling, and then observes what was happening to poverty as

measured by the headcount ratio relative to the \$1 per day poverty line. His calculations are reported in Table 4.4. As we must expect, the largest reductions in poverty take place during spells in which growth is positive and inequality is falling.

Table 4.4
Median Annual Rates of Increase or Decrease in the \$1 per Day Headcount Ratio
For Four Types of “Spells”, Reported in Ravallion (2001)

Change in Gini coefficient:	Growth in average income:	
	Negative	Positive
Rising	+14.3 %	-1.3 %
Falling	+1.7 %	-9.6%

The mathematical relationships tell us that *if* we could increase economic growth while holding the distribution of income constant, or *if* we could speed reductions in inequality while holding economic growth constant, we would speed the rate of poverty reduction. But they tell us nothing about the likely effects on poverty of *actual* policy efforts to speed growth or reduce inequality. Not only might well-intentioned policy efforts to speed growth fail, but even successful policy efforts to increase growth might also increase inequality, leaving the net effect on poverty ambiguous. Similarly, efforts to redistribute income might slow growth, again leaving the net effect on poverty ambiguous. To understand these policy-relevant relationships between growth, inequality and poverty, we must move beyond mere mathematical relationships to socio-economic relationships.

Growth, inequality and poverty as joint outcomes of the socio-economic system. Our starting point for understanding the deeper socio-economic relationships is to recognize that growth and changes in inequality and poverty are joint outcomes of complex socio-economic systems. A country’s people, taking their current assets as given, and interacting with each other in markets and non-market institutions, make many choices that together determine the levels and growth rates of each person’s income, consumption expenditure, and enjoyment of other dimensions of well-being. The exogenous factors that shape these many decisions and outcomes (i.e. the circumstances that people within the system take as given at any point in time) include the level, mix and distribution of a country’s initial assets, the external economic conditions the country faces, and the policies pursued by the government and other development actors. Measures of growth, inequality, and poverty reduction simply summarize different features of the multidimensional outcomes that emerge from the system. They are outcomes jointly determined by initial conditions, policies and external circumstances.

When formulating questions about the relationships between growth, inequality and poverty reduction, then, we must be careful not to over-simplify. It does not make much sense to think of growth or reductions in inequality as “causing” poverty reduction, as if we could somehow meaningfully impose a growth or distribution outcome on the system

and then observe what this does to poverty. It does, however, make sense to pose questions about changes in the exogenous factors that shape growth and distribution outcomes. For example, we might ask: “Do policies that are successful in speeding economic growth, regardless of their specifics, have a strong tendency to reduce poverty as well?”

The disaggregated impacts of growth policies. The next step toward understanding policy-relevant relationships between growth, income poverty and inequality is to think practically about how policies to promote growth impact the socio-economic system. When policies succeed in speeding economic growth, they do so by speeding the economy’s accumulation of assets (broadly defined, as in Chapter 3). Diverse growth policies might promote the creation of quite different types of assets, owned by different members of society. Some growth policies promote the accumulation of buildings and machines owned by large companies producing consumer goods for the domestic market. Others promote private investment in export manufacturing, while yet others build roads that connect remote regions to markets or help small farmers increase their productivity through the adoption of advanced technologies. These increases in assets will tend to raise productivity and incomes for their owners, and will also lead new owners or users to an array of changes in their supplies and demands for labor, goods and services. Prices are likely to change in many markets, with each price increase hurting net buyers of the good in question while helping net sellers. Economic change is also likely to generate change in non-market institutions, as geographic movement of peoples and new economic relationships place stresses on old social norms and relationships. The net result of all these changes is likely to be a complicated set of increases in income and well-being for some groups and decreases for other groups. Policies that promote growth will reduce income poverty only if the poor are among those who enjoy income increases.

What must be true for the poor to share in the benefits of growth? Drawing on the Chapter 2’s discussion of the households’ pursuit of well-being, we observe that policies promoting asset creation and growth may improve the circumstances of the poor through four channels. First, they might increase the assets that the poor own or have access to, whether because the policies give the poor new assets directly or because they encourage the poor to undertake their own investments. Second, even if the assets created as a result of good growth policies are not owned by the poor, they may increase the incomes of the poor indirectly, by driving up wages for low-skill labor or driving up the prices of goods produced by household enterprises operated by the poor. Third, increases in assets owned by the non-poor might also lead to expansion of the supply of goods consumed by the poor, driving down their prices. Finally, with additional assets and income, the non-poor might become involved in public or private safety net institutions, which channel funds to the poor and vulnerable.

Why might the poor fail to share in the benefits of growth? Unfortunately, there are reasons to worry that each of the potential connections between growth and the incomes of the poor may fail without explicit policy attention. Unless growth policies explicitly direct new assets to the poor, the poor may fail to participate in asset creation, because

they may be unable to take advantage of attractive new investment opportunities generated by growth policies. As we will see in Chapters 9 and 10, their lack of collateral may render the poor unable to obtain the loans required to make investment feasible, and many of the investments of greatest potential value to the poor -- in infrastructure, education and improved crop technologies -- have public good characteristics that render investment difficult to achieve without intervention. Growth may also fail to increase the incomes the poor earn from their pre-existing assets in markets for labor, or for the produce of their small farm and non-farm enterprises, because the poor may live in remote areas un-integrated into the larger markets in which growth is taking place (as discussed in Chapter 6), or because the assets whose increase underlies growth lead their owners to increase demand for skilled labor and specialized inputs produced by large suppliers rather than the unskilled labor and products that the poor have to offer (as discussed in Chapter 7). Remoteness from markets may also prevent the poor from enjoying any consumer price reductions arising out of increased supply. Finally, growth may fail to stimulate greater transfers to the poor through private safety net institutions, both because some of the poor may always have been excluded from those networks and because the socio-economic change associated with growth may cause old networks to break down (see Chapter 10),

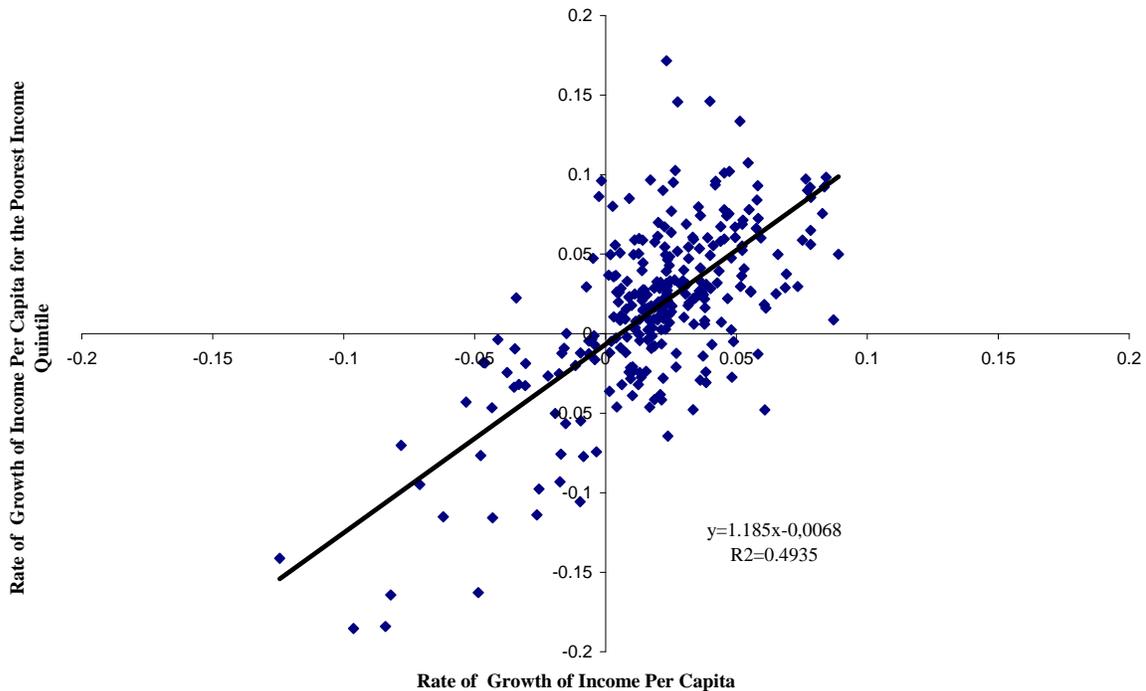
Why might growth policies differ in the extent to which they are “pro-poor”? While it is theoretically possible that the poor will tend to share in growth under a wide range of policies, the preceding paragraphs provide many reasons to suspect that growth policies will differ in the extent to which they are pro-poor. Growth policies that explicitly direct new assets to the poor or address their financial constraints, that extend the reach of markets into more remote areas, and that emphasize accumulation of assets that tend to stimulate the demand for unskilled labor are more likely to bring poverty reduction than other policies equally successful at producing growth in average income.

What do we know about the recent historical tendency for economic growth to bring poverty reduction? Several recent studies of aggregate cross-country data highlight two key observations. First, on average over the whole world (or at least over the countries for which adequate data can be mustered) economic growth has been associated with little change in income distribution; and more rapid economic growth has thus been associated with more rapid reduction in income poverty. Second, the experience of individual countries has, however, been highly dispersed around that world average, with poverty sometimes falling significantly even when growth is very slow, and sometimes failing to fall even when growth is rapid.

Figure 4.6 is drawn from a much-cited study by David Dollar and Aart Kraay (2006). The 285 data points plotted in the graph describe growth experiences for 92 countries over various time intervals (of at least 5 years in length) spanning a 40 year period. Each country-interval observation is represented by a point that indicates the rate of economic growth (measured along the horizontal axis) and the rate of growth of the income of the poorest 20 percent of the population (measured along the vertical axis).

Figure 4.6

Cross-Country Correlation between Income Growth for the Bottom Twenty Percent and Average Income Growth



The slope of the regression line drawn through the middle of the scatter plot in Figure 4.6 is very close to one, indicating that, on average, every one percentage point increase in the rate of economic growth is associated with a one percentage point increase in the rate of growth of average income for the poorest 20 percent. This implies that on average (over the whole world) the income share of the poor (a measure of inequality) has tended to be unchanged by growth, and the incomes of the poorest 20 percent have risen at the same rate as the incomes of the average members of the population. Authors like Dollar and Kraay (2006) emphasize this average relationship and conclude (as indicated by the title of their paper) that “Growth is good for the poor.”^{4,5} Many policy advocates have interpreted this average relationship as reflecting a strong, worldwide structural tendency for growth to translate into significant poverty reduction, regardless of the policies used to promote growth.

Notice, however, that the data points representing particular places and time intervals are highly dispersed around the regression line. That is, among episodes with similar growth rates, the income shares of the poor have risen in some cases and fallen in others, and rates of poverty reduction have varied a great deal! In light of the theoretical discussion

⁴ Their paper is part of a literature that claims to study the “causal effect of growth on poverty.” This chapter has argued that growth and poverty are joint outcomes of the socio-economic system, and that it is not even meaningful to think about “growth causing poverty reduction.” We might thus question whether this literature addresses a well posed research question.

⁵ For a discussion of how measurement errors might cause a regression like this one to produce a slope estimate close to one, even when the true average coefficient is greater or less than one, see Deaton (2005).

above, this is no surprise, given that the countries of the world are diverse in economic structure and have pursued quite different policies in recent decades.⁶ Indeed, it casts doubt on conclusions drawn by analysts who pay attention only to the average relationship. It suggests that we cannot count on a strong structural tendency for growth to bring poverty reduction in any particular country or time period. It seems likely that the details of growth policies matter, and that efforts to identify pro-poor growth policies will be required if policymakers are to have simultaneous success in speeding growth and reducing poverty.

More disaggregated case studies of economic growth episodes provide further reason to suspect that growth policy specifics – such as the extent to which they favor the regions and economic sectors where large fractions of the poor live and work, and the extent to which they favor creation of infrastructure connecting remote regions to markets – must play an important role in shaping their poverty reduction impacts. For example, across states of India, rates of economic growth over the last 25 years ranged from 2.2 percent in Bihar to 7.2 percent in Karnataka. More rapid growth in non-agriculture than in agriculture contributed to a tendency for growth to be more rapid in states starting at higher levels of manufacturing development and per capita income levels. As a result, growth was associated with rising inequality in average income across regions (Chaudhuri and Ravallion, 2006). In Ghana, export crop market liberalization and infrastructure building programs led to rapid growth in the rural forest zone, where most export crops are grown, but to stagnation and decline in other (poorer) rural areas, where food crops are grown (Aryeety and McKay, 2007). Studies of Ethiopia (Dercon, 2006) and Viet Nam (Klump, 2007) highlight how incomes in more remote regions tend to grow more slowly than incomes in places better connected to markets, and for households living at greater distances from markets within regions. For other interesting case studies, see Ravallion and Chen (2007) on China and Thurlow and Wobst (2006) on Zambia.

Case studies also suggest that even when growth translates into reduction in income-based poverty statistics, they need not translate into improvements in complementary non-income measures of poverty. For example, growth in Ghana brought a large decline in income poverty rates, as well as improvements in school enrollment and the share of the population with access to potable water, but also was accompanied by increases in malnutrition among children under five years old, and reductions in health care use and life expectancy (Aryeety and McKay, 2007). This provides further reason to pay attention to details when seeking growth accompanied by multidimensional reductions in poverty.

⁶ The authors attempt to study the role that policy differences might play in generating this dispersion, by looking for systematic tendencies for the income growth of the poorest 20 percent to rise more or less (while holding overall income growth constant) depending on the levels of several “policy” variables. They find no systematic relationship. Unfortunately, their policy variables include only a few broad macro outcome variables such as inflation rates and average years of primary education, which fall far short of measuring the detailed differences in policy that might determine the strength of the tendency for the poor to share in the benefits of growth.

What effects might successful growth policies have on vulnerability? Just as the details of growth policies shape the extent to which the resulting growth brings poverty reduction, so the details are likely to shape the effects on vulnerability. Growth policies will be more likely to reduce vulnerability to destitution if they raise the incomes of the poor (rendering it less likely that income reductions will drive them all the way into destitution), if they create assets (such as floodwalls or immunization programs) that reduce inherent risks of crop loss or ill health, or if they facilitate the development of market relationships through which the effects of some shocks can be spread out (Chapter 6). On the other hand, it seems possible that growth may even increase vulnerability, unless explicit efforts are undertaken to prevent this. Policies that speed growth often require tremendous change at the microeconomic level, with some firms and sectors being born and others dying. Risk of unemployment may rise, and the geographic movement of people in response to new opportunities may mean a breakdown of traditional social safety net institutions.

How might explicit efforts to reduce poverty, inequality or vulnerability slow or speed growth? Whether it makes sense to pursue broad-based development through single-minded efforts to increase growth – eschewing explicit efforts to reduce poverty, inequality or vulnerability – depends not only on whether growth has a strong tendency to bring poverty reduction (regardless of the policies pursued), but also on the likely growth impacts of explicit policy efforts to reduce poverty and vulnerability. If such policies tend to slow growth, then they might prove counter-productive in the long run, reducing poverty today but reducing growth and the potential for sustained future poverty reduction. If, however, such policies have the potential to speed growth, then they might be doubly valuable for stimulating broad-based development.

Explicit poverty reduction policies (such as those discussed in Part VII) tax the non-poor (or draw on external funding) to finance the distribution of benefits to the poor. In the early days of development, most theorists and policy analysts assumed that such efforts would slow economic growth, for two reasons. First, if the poor save a smaller fraction of their income than do the non-poor, then transferring income to the poor would reduce the rate of saving, tending to reduce the rate of investment and asset accumulation. (See the Lewis Model described in Appendix 3A.) Second, policies and programs that tax the incomes of the non-poor and provide the poor with “handouts” might weaken the incentive of both groups to work hard and innovate. (We will examine this more carefully in Chapter 19.)

In the 1990s, thought concerning the impact of poverty reduction policies on growth underwent a profound transformation. Motivated by cross-country empirical observations, theoretical research pointed to several channels through which efforts to reduce poverty and inequality might *increase* economic growth (Aghion, et al., 1999). A first line of theoretical research highlighted another implication of the financial market failures mentioned above. If the poor lack financing at reasonable interest rates, then they may be prevented from undertaking investment projects with high returns. The projects left undone might even have higher returns than the projects undertaken by the non-poor. By allowing the poor to undertake their high-return investments, policies that

increase the incomes or assets of the poor might increase aggregate rates of saving and investment, and increase the average return on investments, thereby increasing economic growth (Galor and Zeira, 1993). Recent innovations in the design of poverty reduction programs create even greater potential for poverty reduction programs to spur investment and growth, by strengthening the incentives of participants to invest in human capital formation. For example, new “conditional cash transfer programs” direct benefits to poor households conditional on the households’ compliance with requirements that they send their children to school and make use of public clinics, appear to be successful in increasing human capital investments (see Chapter 20).

A second channel through which explicitly re-distributive policies might improve economic growth is by insuring poor and near poor households against the worst potential consequences of risky investments. In the absence of insurance they may fail to undertake investments that would raise their average incomes. If safety net programs provide households with the assurance that they may obtain program support if their risky investments work out badly, then households might become more willing to undertake investments that contribute to economic growth as well as sustained poverty reduction. Some public employment programs guaranteeing employment at low wages produce such insurance (Chapter 21).⁷

Tentative conclusions and where we go from here

We have taken a first, broad look at the socio-economic relationships between growth, inequality, poverty, vulnerability and policy, and have found strong reason to suspect that successful growth efforts can have diverse implications for poverty, inequality and vulnerability, depending on the exact nature of policies pursued, and on differences in local and global conditions. We have also raised (in a general fashion) the possibility that explicit efforts to reduce poverty, inequality and vulnerability may even speed growth. This suggests that when analyzing and selecting policies to employ in the pursuit of multidimensional development objectives, it is important to consider not only their macro implications for the rate of economic growth, but also their more detailed micro implications for the pattern of changes in income and well-being across diverse poor and non-poor groups. It is also useful to consider not only policies that are traditionally associated with the promotion of economic growth, but also policies whose most obvious impacts are re-distributive.

Notice, however, that the discussion here sheds little practical light on the exact nature of growth policies that are likely to be good for poverty reduction, and sheds little light on the exact nature of safety net policies that would have the greatest potential to speed economic growth. Indeed, the exact nature of such policies is likely to differ from place to place. If we wish useful answers to practical policy questions, we must develop a

⁷ Another related literature explores political economy channels through which higher initial levels of income or wealth inequality in a society cause lower subsequent growth rates. For example, Alesina and Rodrik’s (1994) theoretical model suggests that in more equal societies, where the median voter derives a larger fraction of income from capital, government will tend to pursue policies that tax capital less and encourage growth more. Such models have more to say about what policies are likely to be pursued, than about the likely growth-inducing impacts of redistributive policies, should they be pursued.

much more detailed, disaggregated framework, within which to study how specific policies – defined in practically detailed terms – are likely to affect diverse socioeconomic groups, whether directly or indirectly through resulting changes in markets and non-market institutions. Part III of the text develops this framework, while later Parts apply the framework in the analysis of specific policies.

References and Further Reading

Aghion, P., E. Caroli, and C Garcia-Penalosa, 1999, “Inequality and Economic Growth: The Perspectives of New Growth Theories,” *Journal of Economic Literature* 37(4), p. 1615-1660.

Alwang, J., P. B. Siegel, and S.L. Jorgensen, 2001, “Vulnerability: A View from Different Disciplines,” Social Protection Discussion Paper Series No. 0115, World Bank.

Aryeety, E. and A. McKay, 2007, “Ghana: The Challenge of Translating Sustained Growth into Poverty Reduction,” in T. Besley and L.J. Cord (eds.), *Delivering on the Promise of Pro-Poor Growth: Insights and Lessons from Country Experiences*, Palgrave Macmillan and the World Bank.

Baulch, B. and J. Hoddinott (2000), “Economic Mobility and Poverty Dynamics in Developing Countries,” *Journal of Development Studies*, vol. 36, no. 6, August, pp. 1-24.

Besley, T. and R. Burgess, 2003, “Halving Global Poverty,” *Journal of Economic Perspectives* 17(3), p. 3-22.

Besley, T. and L. Cord, eds. 2007, *Delivering on the Promise of Pro-Poor Growth: Insights and Lessons from Country Experiences*, World Bank and Palgrave Macmillan.

Bidani, B. and M. Ravallion (1993), “A Regional Poverty Profile for Indonesia,” *Bulletin of Indonesian Economic Studies* Vol. 29, No. 3, December, 37-68.

Carter, M. and C. Barrett, 2006, “The Economics of Poverty Traps and Persistent Poverty: An Asset-Based Approach,” *Journal of Development Studies* 42(2), p. 178-99.

Chaudhuri, S. and M. Ravallion, 2006, “Partially Awakened Giants: Uneven Growth in China and India,” World Bank Policy Research Working Paper 4069.

Deaton, A., 2001, “Counting the World’s Poor: Problems and Possible Solutions,” *World Bank Research Observer* 16(2), p.125-47.

Deaton, A., 2006, “Measuring Poverty,” in Banerjee, A., R. Benabou, and D. Mookherjee (eds.), *Understanding Poverty*, Oxford University Press, p. 3-16.

Dercon, S., 2005, “Vulnerability: A Micro Perspective,” working paper, Oxford University.

Dercon, S., 2006, “Economic Reform, Growth and the Poor: Evidence from Rural Ethiopia,” *Journal of Development Economics* 81(1), 1-24.

Dollar, D. and A. Kraay, 2002, "Growth is Good for the Poor," *Journal of Economic Growth* 7, pp. 195-225.

Foster, J., J. Greer and E. Thorbecke, 1984, "A Class of Decomposable Poverty Measures," *Econometrica* 52(3), 761-66.

Galor, O. and J. Zeira, 1993, "Income Distribution and Macroeconomics," *Review of Economic Studies* 60(1), p.35-52.

Grosh, M. and P. Glewwe 2000, *Designing Household Survey Questionnaires for Developing Countries: Lessons from 15 Years of the Living Standards Measurement Study*, The World Bank.

Klump, R., 2007, "Pro-Poor Growth in Vietnam: Miracle or Model?" in T. Besley and L.J. Cord (eds.), *Delivering on the Promise of Pro-Poor Growth: Insights and Lessons from Country Experiences*, Palgrave MacMillan and the World Bank.

Ravallion, M. (1992), "Poverty Comparison: A Guide to Concepts and Methods," *Living Standards Measurement Working Paper No. 88*, World Bank, February.

Ravallion, M. and S. Chen, 2007, "China's (Uneven) Progress Against Poverty," *Journal of Development Economics* 82, p. 1-42.

Sen, A. (1981), "Concepts of Poverty" and "Poverty: Identification and Aggregation", Chapters 2 and 3 in *Poverty and Famines: An Essay on Entitlement and Deprivation*, Oxford: Clarendon Press.

Thurlow, J. and P. Wobst, 2006, "Not All Growth is Equally Good for the Poor: The Case of Zambia," *Journal of African Economies* 15(4), 603-625.

United Nations Development Program, 2007, *Measuring Human Development: A Primer*, New York.

World Bank, *World Development Report 2000/01: Attacking Poverty*, 2001.

Questions for Review

1. What is a general definition of "poverty"?
2. What three questions must be answered when constructing a more specific definition or measure of poverty?
3. Discuss the relative merits of using the following measures of living standards when evaluating a household's poverty: **real income per capita**, **real consumption expenditure per capita**, **direct measures** of access to and use of health care, **household assets**.
4. Define *poverty lines*. What is the distinction between **absolute**, **relative** and **hybrid** approaches to constructing poverty lines?
5. What are some of the questions that must be addressed in a constructing a poverty line for comparison to measures of per capita household income or consumption expenditure?
6. Why might it be useful to identify poverty lines for more than one level of poverty?
7. What is meant by the distinction between **persistent** and **transitory** poverty?
8. Define the following measures of an individual's poverty: **income gap**, and **proportional income gap**.
9. Discuss the strengths and weaknesses of the following **aggregate** poverty measures:

- **Headcount Ratio,**
 - **Total Income Gap,**
 - **Average Proportional Income Gap** among the poor,
 - **Poverty Gap Index** (or P_1),
 - **Squared Proportional Income Gap Index** (or P_2).
10. What is meant by the statement that “selecting a single poverty statistic for measuring poverty reduction success is a value-laden activity”?
 11. Compare and contrast the policy objectives of reducing poverty and reducing **inequality**.
 12. What two questions must be addressed when constructing specific definitions and measures of inequality?
 13. Draw a **Lorenz curve** diagram, label it carefully, and describe how it is constructed.
 14. Draw Lorenz curves for societies with perfect income equality and perfect income inequality, and explain why you drew them the way you did.
 15. What does it mean for one Lorenz curve to **dominate** another?
 16. Explain why it is difficult to rank two countries according to the degree of inequality found in them when their Lorenz curves cross, and discuss the role that values must play in determining the ranking.
 17. Describe some simple measures of income inequality based on **share of total income**.
 18. Define the **Gini coefficient** and discuss the values embodied in its construction.
 19. Provide a general definition of **vulnerability**.
 20. What two sets of circumstances must combine to render a household vulnerable?
 21. What are some of the questions that would have to be addressed in constructing more specific definitions and measures of vulnerability?
 22. What does it mean to say that vulnerability is a forward-looking concept?
 23. What does it mean to say that vulnerability is a probabilistic concept?
 24. Describe some of the difficulties of measuring vulnerability at the household level.
 25. What are **panel datasets**, and what are their strengths and weaknesses for shedding light on vulnerability?
 26. Why would most development analysts want to employ measures from all three of the following categories when measuring development success:
 - a measure of economic growth
 - measures of income poverty or inequality
 - measures of non-income poverty and vulnerability.
 27. Discuss the roles of **values** and **beliefs** in explaining disagreements regarding the best policies to employ in pursuing development.
 28. What is meant by the terms **pro-poor growth policies** and **safety net policies**?
 29. What is the “mathematical relationship” between growth, inequality and poverty?
 30. Why does this mathematical relationship not provide us with sufficient evidence for arguing that policies to speed economic growth (unaccompanied by efforts to change the distribution of income) will always speed poverty reduction?
 31. What is the significance of the statement that “growth, and changes in poverty and inequality are joint outcomes of complex socio-economic systems”?
 32. What are the four channels through which policies that promote asset accumulation (and thus economic growth) might lead to income increases for the poor?
 33. Why might we worry that each of these channels might fail to be activated by particular growth policies?
 34. Describe what we learn about the worldwide average relationship between growth, inequality and poverty from figure 4.6.
 35. Describe what we learn about the relationship between growth, inequality and poverty from the dispersion around the regression line in Figure 4.6.

36. What are some channels through which growth policies may reduce or increase vulnerability?
37. Through what channels might explicit efforts to reduce poverty, inequality or vulnerability bring increases or decreases in rates of economic growth?

Questions for Discussion

1. Which of the aggregate poverty measures defined in this chapter appeals to you as the measure that best captures your concerns and priorities? Might your answer to this question depend on circumstances? Explain.
2. Do you think that your own concerns about people with low incomes and poor living standards are driven more by concerns regarding poverty or inequality? If you had to choose between two policies, one of which would reduce poverty while increasing inequality, and the other of which would reduce inequality but increase poverty, which would you choose?
3. Just as we can construct statistics to summarize poverty or inequality, so can we construct indices to summarize the state of “development”. One such index is the “Human Development Index”, which is constructed by adding together three indices related to life expectancy, education and GDP per capita. The life expectancy and GDP per capita indices are constructed by taking statistics on life expectancy at birth and GDP per capita and re-scaling them so that they are expressed on a zero to one scale. The education index is constructed from data on the adult literacy rate and the school enrollment rate (in primary, secondary and tertiary schools), each re-scaled to a zero-one scale, with two-thirds weight placed on the adult literacy index and one-third on the school enrollment index. For a discussion of the issues involved in constructing such a measure, see United Nations Development Program (2007). What do you think are the strengths of this index as a measure of development? What do you think of as the weaknesses of this index? What might be the uses of such an index?
4. Pick a country you would like to learn more about. Access the most recent World Development Report, published by the World Bank and available on-line in the World Bank web pages. Find the “Selected Indicators” near the end of the text, and study the tables provided there on “Key development indicators” for countries around the world. What do you learn there about growth, poverty and inequality in the country you picked? How does your country’s performance in these areas compare to the performance of other countries with similar levels of GDP per capita? How does it compare to performance in the rest of the world?

Problems

1. This problem provides a brief review of summation notation, using an example related to the distribution of income in a population. Order the individuals in a group from 1 to N , with individual 1 being the poorest person and individual N being the richest. An individual’s index is his rank number in this ordering. For example, the fifth poorest person has person index 5. Let Y_i be the income of a person i . In summation notation, the Greek letter

Σ (capital sigma) denotes a sum. More specifically, the expression $\sum_{i=1}^N Y_i$, which is read as

“the sum from $i=1$ to N of Y -sub- i ”, can be defined as follows:

$$\sum_{i=1}^N Y_i = Y_1 + Y_2 + \dots + Y_N.$$

- a. Using summation notation, write down a formula for the mean of income in this population.
- b. Consider the expression

$$\frac{1}{q} \sum_{i=1}^q \frac{(z - Y_i)}{z},$$

where z is the income poverty line and q is the individual with the highest income who remains under the poverty line. State in plain language the calculation this expression describes and offer an intuitive interpretation of the statistic that results from this calculation.

2. The following table lists the incomes for all individuals in each of three very small countries (just 10 people each). Incomes are listed in Currency Units (CUs) per week. The official poverty line is 10 CUs per week.

Individual	Country 1	Country 2	Country 3
1	8	3	6
2	8	3	6
3	8	9	6
4	8	9	6
5	8	12	6
6	8	12	12
7	12	12	12
8	12	12	12
9	12	12	12
10	12	12	12

- a. Fill in the following table. (See example calculations below table.)

Poverty Measure	Country 1	Country 2	Country 3
P_0 (Headcount Ratio)	.6		
P_1	.12		
P_2	.024		

- b. Fill in the following table. For each poverty measure, enter into the table the country rankings from most poor (1) to least poor (3) according to that measure.

	Country 1	Country 2	Country 3
Ranking according to P_0			
Ranking according to P_1			
Ranking according to P_2			

- c. Write a brief essay on differences between the three poverty measures in the *values* underlying them, and how these differences in values lead to differences in poverty *rankings* in this case.

3. Draw a diagram like that in Figure 4.1, including an initial income locus. For each of the following cases, draw a new income locus such that the shift from the initial locus to the new locus is associated with the changes described:

- reduction in the headcount ratio but increase in average depth of poverty among the (remaining) poor
- constant headcount ratio but reduction in poverty gap index
- constant headcount ratio and poverty gap index but reduction in the P2 measure.

4. The table on the next page describes the distributions of income in two states (A and B) and in two sub-regions (rural and urban) of each state. Every individual in these states lives in a household of size one, and has an income of exactly 100, 200 or 10,000 dollars per year (so it is easy to describe the distributions and calculate poverty statistics). The first two sections of the table present the numbers and percentages of individuals in each region at each income level.

- a. Fill in the two rows of poverty statistics in each of the last two sections of the table. For the first of these sections, use a poverty line of \$201. For the last use a poverty line of \$101.
- b. Suppose attention is restricted to headcount ratio statistics employing a poverty line of \$201, and representatives of State A are attempting to argue that their state should be given priority in the allocation of poverty alleviation funds. Would they prefer to employ statistics calculated at the state level (columns 3 and 6) or sub-region level (1,2,4 and 5)? Why?
- c. Suppose attention is restricted to headcount ratio statistics calculated at the sub-region level, and representatives of State A are still attempting to make the same argument. Would they prefer to employ a poverty line of \$201 or \$101? Why?
- d. Suppose attention is restricted to poverty statistics calculated at the sub-region level and employing a poverty line of \$201, and representatives of State A are still at it. Would they prefer to use headcount ratio statistics or total income gap statistics?
- e. Given that no single statistic captures everything that matters about poverty, why might lawmakers, bureaucrats or NGO program officers still want to use well defined allocation rules based on a simple statistic? In particular, why not just say: “We know poverty when we see it, just give us *discretion* to allocate funds as we see fit.”?

	State A			State B		
	Urban	Rural	Total	Urban	Rural	Total
Population (# of people)	10,000	10,000	20,000	10,000	90,000	100,000
No. of people with Income of:						
\$100	0	0	0	0	45,000	45,000
200	2,000	6,000	8,000	2,000	0	2,000
10,000	8,000	4,000	12,000	8,000	45,000	53,000
Percent of Population with Incomes of:						
\$100	0	0	0	0	50	45
200	20	60	40	20	0	2
10,000	80	40	60	80	50	53

Poverty statistics Using poverty line Of \$201: • Headcount ratio(%) • Total Income Gap(\$)						
Poverty statistics Using poverty line Of \$101: • Headcount ratio (%) • Total Income Gap(\$)						

5. The first column in the following table describes the initial incomes of all 10 people in a very small country. The official poverty line in this country is 15.

- Fill in the table in the following way. For the first 10 rows, in the second and third columns, fill in the income each person would have if each person's income exactly doubled or tripled, respectively. Then fill in the remaining rows, making use of the income information from the same column.
- Plot a Lorenz curve describing the distribution of income in the first column. How would the Lorenz curves for the distributions in the second and third column compare to this one? Explain.
- Imagine that economic growth brings a change in incomes from the situation described by the first column to the situation described by the second column in just one year. What rate of economic growth would this imply? Describe in words what happens to the incomes of those who were officially poor at the beginning of the period. Is the improvement for the poor in this case picked up better by changes in the Headcount Ratio or Changes in the average income among the officially poor?
- Now imagine that economic growth brings a change in incomes from the situation described by the first column to the situation described by the third column in just one year. What rate of economic growth would this imply? Describe in words what happens to the incomes of those who were officially poor at the beginning of the period. Is the improvement for the poor in this case picked up better by changes in the Headcount Ratio or Changes in the average income among the officially poor?
- Study the absolute income increases enjoyed by various members of the economy over the course of growth from the first column to the second or third column. Comment on the statement: "If the distribution of income remains constant during growth, then everyone in the economy shared equally in the growth."

Person	Initial Income	Income After Distribution- Neutral Doubling	Income after Distribution- Neutral Tripling
1	1		
2	1		
3	1		
4	7		
5	7		
6	7		
7	19		
8	19		
9	19		

10	19		
Total Income	100		
Average Income			
Headcount Ratio			
Average Income Among the Officially Poor			