poor but also to the severity of their poverty—is attributable to Watts (1968):

$$P_{\text{Watts}} = \sum_{i=1}^{q} \log(z/y_i)/n.$$

Another early poverty measure that is also sensitive to the severity of poverty is the second measure of Clark, Hemming, and Ulph (1981)

$$P_{C\text{-}H\text{-}U} = \frac{1}{c}\sum_{i=1}^{q} [1 - (y_i/z)^c]/n, \text{ where } c \leq 1.$$

One finds these measures used only occasionally in empirical work.

4.5 Poverty Dominance

All of the poverty measures presented in section 4.4 permit complete binary comparisons for any two income distribution vectors $X, Y \in \Omega$. That is, given a particular poverty measure $P(.)$ and a particular poverty line z, the poverty measure $P(.)$ is able to say which distribution has more poverty than the other.

Uncertainty arises in poverty comparisons from two sources: We may not be sure where exactly to draw the poverty line z, and we may not be prepared to commit to a particular poverty measure to the exclusion of others. We could, of course, make numerical calculations for a wide range of poverty lines and poverty measures and see if they all rank X as having more (or less) poverty than Y. But even this would not be conclusive proof; it could be that one of the other poverty measure/poverty line combinations that we didn't try would give the opposite ranking.

Fortunately, dominance methods are now available to test whether one income distribution has more poverty than another for a broad class of poverty measures and a wide range of poverty lines. Unlike inequality measurement, in which dominance theory has a long tradition, dominance analysis in poverty measurement is still a new field. Important contributions have been made by Atkinson (1987), Foster and Shorrocks (1988), Ravallion (1994), and Jenkins and Lambert (1998). Poverty dominance may be nested within the broader concept of "deprivation dominance" (Jenkins and Lambert 1997; Xu and Osberg 1998; Shorrocks 1998).

We begin by noting that most of the poverty measures considered in this chapter—including the headcount ratio, the P_α measure, the Watts index, and the Clark-Hemming-Ulph measure, among others—are members of the general additive class

$$P = \sum_{i=1}^{n} p(z, y_i)/n \text{ such that}$$

$$\begin{aligned} &1.\ p(z, y_i) = 0 \text{ if } y_i \geq z, \text{ and} \\ &2.\ p(z, y_i) > 0 \text{ if } y_i \leq z \end{aligned} \tag{4.1}$$

The fact that many poverty measures belong to this general class (but not all do—in particular, the Sen index does not) leads us to consider whether there are dominance criteria for these measures. One may pose three related questions:

1. Are there circumstances under which *all* members of the class of poverty measures (4.1) would rank one income distribution as having more poverty than another for a *given* poverty line z?

2. Given the controversy that usually goes with setting a poverty line, dominance criteria encompassing many poverty lines would be very desirable. In particular, are there circumstances under which a *given* poverty measure belonging to class (4.1) would rank one income distribution as having more poverty than another for a *range* of poverty lines $\underline{z} < z < \bar{z}$?

3. Are there circumstances under which *all* poverty measures belonging to class (4.1) would rank one income distribution as having more poverty than another for a *range* of poverty lines $\underline{z} < z < \bar{z}$?

This section shows that the answers to all three questions are affirmative. The following notation will be used:

z = poverty line,

y_{min} = lowest income in population,

y_{max} = highest income in population,

z_{max} = highest possible poverty line,

$F_X(z)$ = cumulative density function for distribution X when the poverty line is z, also called the "poverty incidence curve."

For the class of poverty measures given by (4.1), we now consider three levels of poverty dominance.

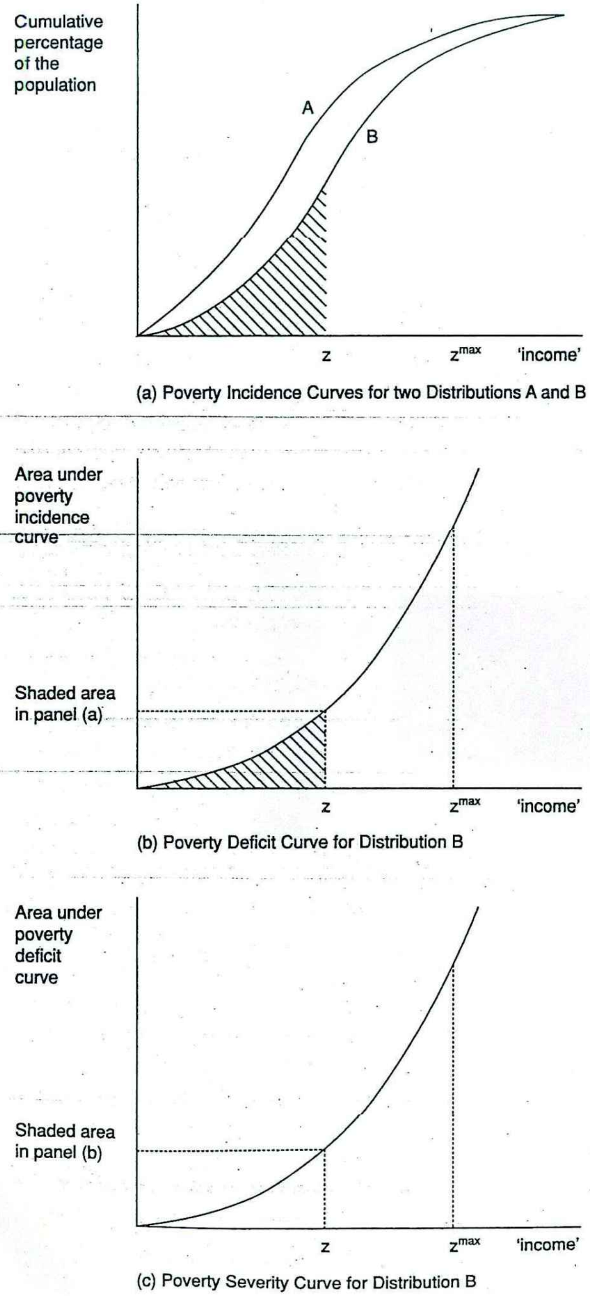

First-Order Dominance (FOD)

First-order dominance in poverty measurement is defined as follows: If the cumulative density function for distribution A (F_A) is everywhere at least as high as that for distribution B (F_B) for all z between y_{min} and z_{max}, A *first-order-dominates* B, written A *FOD* B. (See figure 4.2.a.) An increase in some poor person's income, holding other poor persons' incomes constant, is sufficient but not necessary for first-order dominance.

The relationship between *FOD* and poverty dominance using measures belonging to the class (4.1) is given by the following theorem:

THEOREM 4.1 (ATKINSON 1987; FOSTER AND SHORROCKS 1988): *A FOD $B \Leftrightarrow pov_A > pov_B$ for all poverty measures belonging to the class (4.1), or for any monotonic transformation thereof, and for all poverty lines between y_{min} and z_{max}.*

As applied to the P_α class, theorem 4.1 tells us that A *FOD* $B \Rightarrow pov_A > pov_B$ for all P_α, $\alpha \geq 0$ and for all poverty lines between y_{min} and z_{max}.

Theorem 4.1 gives a condition for ranking two income distributions when their poverty incidence curves do not cross. If they do cross, rankings may be still be possible using the results of:

Second-Order Dominance

Second-order dominance in poverty measurement is defined as follows: Take the areas under the F_A and F_B curves. Call these *poverty deficit curves*, $D(z)$, as in figure 4.2.b. If the poverty deficit curve for A is somewhere above and never below the curve for B for all z between y_{min} and z_{max}, A *second-order-dominates* B, written A *SOD* B.

A subset of the measures belonging to the class (4.1) will give identical poverty rankings in cases of second-order dominance. Specifically:

THEOREM 4.2 (ATKINSON; FOSTER AND SHORROCKS): *A SOD $B \Leftrightarrow pov_A > pov_B$ for the subset of poverty measures belonging to the class (4.1) which are strictly decreasing and at least weakly convex in the incomes of the poor, or for any monotonic transformation thereof, and for all poverty lines between y_{min} and z_{max}.*

Figure 4.2
Three orders of Poverty Dominance
Source: Ravallion 1994 (pp. 67–68).

For the P_α class, theorem 4.2 tells us that $A\ SOD\ B \Rightarrow pov_A > pov_B$ for all P_α, $\alpha \geq 1$ and for all poverty lines between y_{min} and z_{max}.

If poverty deficit curves do not cross, theorem 4.2 provides a criterion for ranking the poverty of different income distributions. However, the poverty deficit curves may cross, in which case poverty rankings may still be possible by turning to:

Third-Order Dominance

Third-order dominance in poverty measurement is defined analogously to second-order dominance. Take the areas under the poverty deficit curves for distributions A and B and call these *poverty severity curves*, $S(z)$. (See figure 4.2.c.) If the poverty severity curve for A is somewhere above and never below the poverty severity curve for B for all z between y_{min} and z_{max}, *A third-order dominates B*, written *A TOD B*.

Before stating the next theorem, recall the definition of distributional sensitivity given above. We then have:

THEOREM 4.3 (ATKINSON, FOSTER, AND SHORROCKS): $A\ TOD\ B \Leftrightarrow pov_A > pov_B$ for the subset of poverty measures belonging to the class (4.1) which are distributionally-sensitive, or for any monotonic transformation thereof, and for all poverty lines between y_{min} and z_{max}.

For the P_α class, theorem 4.3 tells us that $A\ TOD\ B \Rightarrow pov_A > pov_B$ for all P_α, $\alpha \geq 2$ and for all poverty lines between y_{min} and z_{max}.

Comparing the Different Levels of Dominance

When we examine the definitions of the various levels of dominance, we see a hierarchical relationship among them: If *A FOD B*, then the area under the poverty incidence curve for A is necessarily greater than the area under the poverty incidence curve for B, which is the definition of second-order dominance. Likewise, given second-order dominance, the area under the poverty deficit curve for A is necessarily greater than the area under the poverty deficit curve for B, which is the definition of third-order dominance. Thus:

$$FOD \Rightarrow SOD \Rightarrow TOD. \tag{4.2}$$

Fact (4.2) makes precise exactly how far you can go in defining the variety of poverty measures and range of poverty lines for which one income distribution has more poverty than another. For instance, if neither *A FOD B* nor *B FOD A* but *A SOD B*, then you can conclude that A has more poverty than B only for those poverty measures in class (4.1) which are strictly decreasing and at least weakly convex in the incomes of the poor, or for any monotonic transformation thereof, and for all poverty lines between y_{min} and z_{max}. You would know too that because the poverty headcount ratio is not strictly decreasing in the incomes of the poor, that that measure would not necessarily show greater poverty in A than in B.

Suppose you have decided that you like the P_α class of poverty measures, but you are not sure how far these measures can take you in making ordinal poverty comparisons. Pulling together the preceding results, here is your answer:

1. $A\ FOD\ B \Rightarrow pov_A > pov_B$ for all P_α, $\alpha \geq 0$ and for all poverty lines between y_{min} and z_{max}.
2. $A\ SOD\ B \Rightarrow pov_A > pov_B$ for all P_α, $\alpha \geq 1$ and for all poverty lines between y_{min} and z_{max}.
3. $A\ TOD\ B \Rightarrow pov_A > pov_B$ for all P_α, $\alpha \geq 2$ and for all poverty lines between y_{min} and z_{max}.

Finally, it bears mention that your search for dominance results may be frustrated, because you have chosen a wider than necessary range for possible poverty lines. For example, you may have initially set z_{max} at a very high value (e.g., \$50,000 per capita per year) and then found that the poverty incidence curves cross at a lower income amount z^* (e.g., \$40,000), as shown in figure 4.3. In such cases, you might want to reset z_{max} to z^*, which would enable you to conclude that *A FOD B* for all poverty lines in the restricted range $[y_{min}, z^*]$.

4.6 The Concept of Relative Poverty

Thus far in this chapter, we have dealt only with absolute poverty. Some authors (e.g., Fuchs 1967, 1969; Ruggles 1990; Citro and Michael 1995; Ali 1997; Ali and Thorbecke 1998) take exception to this approach to determining poverty lines, preferring to measure relative poverty instead.[18] For rigorous discussions of absolute versus relative poverty, see Foster and Sen 1997 and Foster 1998.

Actually, "relative poverty" embodies two separate ideas, and the relative poverty measures therefore fall into two categories. In the first type of relative poverty measure, a group that is relatively

18. Other authors measure both relative and absolute poverty; for instance, Atkinson (1983a) and Blackburn (1994).

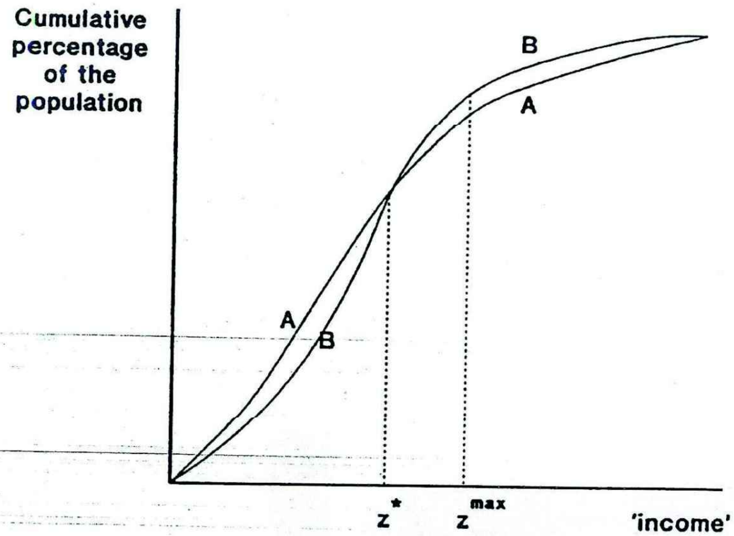

Figure 4.3
Intersecting Poverty Incidence curves
Source: Ravallion 1994 (p. 70).

the poorest (e.g., the poorest 40 percent) is defined, and the poverty measure used is then taken to be the average real income of this poorest group. Consider the example in chapter 1, in which the income distribution changes from

$$\text{Initial} = (\underbrace{1,1,1,1,1,1,1,1,1}_{9},\ \underbrace{2}_{1})$$

to

$$\text{A-C-E-G} = (\underbrace{1,1,1,1,1,1,1,1}_{8},\underbrace{2,2}_{2})$$

to

$$\text{B-D-F-H} = (\underbrace{1,1,1,1,1,1,1}_{7},\underbrace{2,2,2}_{3}).$$

The average absolute income of the poorest 40 percent of the population shows no change in this process. If you see no change in poverty in this process, this type of relative poverty measure might be a

reasonable one for you. But if you judge that poverty has decreased in this process, then you are assuredly *not* a relative poverty adherent, at least in this first sense.

There is, however, a second sense in which you might wish to move in the direction of relative poverty, and that is to use a higher poverty line the richer is the country in which poverty is being measured. Ravallion, Datt, and van de Walle (1991) have found empirically that the poverty lines used in countries tend to increase with their per capita consumption levels, and Ali (1997) regards the desirability of raising the poverty line as the mean increases as "obvious to us, Africans living amidst poverty." While there are different ways of adjusting your poverty line z as a function of the mean income or consumption, μ, the easiest such adjustment is to raise z in proportion to increases in μ, producing a thoroughgoing relative poverty measure. This procedure applies either when z has been set "scientifically" to begin with (e.g., as the cost of purchasing the minimal basket of goods and services) or when z has been set relatively from the beginning. Examples of such relative poverty lines are half the median income (Fuchs 1969), two-thirds of the median income, as is done by the Luxembourg Income Study (Atkinson et al. 1995), half the mean income, as is done by the European Union (O'Higgins and Jenkins 1990; Atkinson 1998) or two-thirds of the median income, as is used on occasion by the World Bank.

Now let us examine what happens to relative poverty when the poverty line z increases proportionately with the mean μ. We may start with a given income distribution X and then increase everybody's real income by the same proportion, producing the new income distribution λX, $\lambda > 1$. When z increases proportionately with μ, the number with incomes below such a relative poverty line is unchanged. So too are the average (normalized) income shortfall of the poor $\bar{I}$ and the Gini coefficient of income inequality among the poor. This means that the poverty headcount (H), the poverty headcount ratio (H), the Sen index of poverty (P_{sen}), and the P_α class will all show *no change* in poverty when z changes in proportion to μ.

You now need to ask yourself whether this is what you want. Personally, I would want poverty in a country to *fall* when everyone experiences a given percentage increase in income. If you feel the same, then these relative poverty approaches are not for you.[19]

19. Be careful not to confuse poverty and inequality here. It is perfectly consistent for you to maintain that a proportionate increase in all incomes leaves relative *inequality* unchanged while reducing absolute *poverty*.

One alternative is to increase the poverty line when economic growth takes place, but by a smaller percentage than the growth rate. Because this is neither fully absolute nor fully relative, it has been termed a "hybrid" approach by Foster (1998). See Citro and Michael 1995 and Atkinson and Bourguignon 1999 for detailed proposals along these lines.

I prefer a different alternative, which is to choose an absolute poverty line, relatively defined. That is, you can set your z higher in relatively rich countries than in relatively poor ones. This is in fact done: The poverty lines used around the world (in 1985 Purchasing Power Parity [PPP] dollars, per person per day) range from $1 for developing countries as a whole to $2 in Latin America to $4 in Eastern Europe and the Commonwealth of Independence States (CIS) nations to $14.40 in the United States. Once these lines are set, they should be adjusted by the respective countries' rates of inflation and nothing more. The best problem that any country could have would be for its economy to grow so fast for so long that its current poverty line is rendered obsolete!

4.7 Summary

"Poverty" has been defined as the inability of an individual or a family to command sufficient resources to satisfy basic needs. A number of technical considerations go into setting a country's poverty line. If we can agree on where to set the poverty line, we can gauge the amount of poverty in a population by measuring the extent of poverty for each constituent individual and then totaling these using a suitable aggregator function. Among the axioms that might be desirable for poverty measures are anonymity, population-homogeneity, monotonicity or strong monotonicity, and distributional sensitivity. The poverty headcount possesses only the first of these properties and the poverty headcount ratio only the first two. However, all four properties are satisfied by the Sen poverty index and the P_α index for $\alpha > 1$. If we are uncertain which poverty measure or which poverty line to use, we may under certain conditions be able to use dominance results to obtain ordinal poverty rankings for a broad class of poverty measures and a broad range of poverty lines; see section 4.5 for details. Finally, and alternatively, if you prefer relative poverty notions to absolute poverty ones, you may prefer one of the types of measures described in section 4.6.

5 Does Economic Growth Reduce Absolute Poverty? A Review of the Empirical Evidence

***

***Except pages 100 - 102**

Most economists accept without question that economic growth reduces absolute poverty. Some of the phrases in our profession reflect this: "trickle-down," "a rising tide lifts all boats," "the flying geese," and so on. I shall refer to this view as the "shared growth" position, in that when economic growth takes place, the poor and others share the fruits of it, to a greater or lesser degree.[1]

On the other side is the distinctly less popular view that economic growth might make the poor poorer. To take just one example, Nobel Prize-winning economist Arthur Lewis (1983) gave six reasons why development of enclaves may lower incomes in the traditional sector: The development enclave may be predatory on the traditional sectors; products of the enclaves may compete with and destroy traditional trades; the wage level of the enclave may be so high that it destroys employment in other sectors; the development of the enclave may result in geographical polarization; development of the enclave may lead to generalized improvements in public health and therefore lower death rates; and development of the enclave may stimulate excessive migration from the countryside.

The most respectable present incarnation of this view is in the work of immiserizing growth theorists, who have proven rigorously that it is possible that economic growth *might* make the poor poorer.[2] Whether growth *does* make the poor poorer or not is an empirical question, to which we now turn.

1. Whether it is to greater or to a lesser degree is, of course, what inequality is all about. Chapter 3 reviewed the growth-inequality evidence.

The evidence reviewed in this chapter is on "absolute poverty," namely, the extent of poverty when a fixed real poverty line is used.

2. A good summary may be found in Bhagwati and Srinivasan 1983 (chapter 25). Bhagwati himself now dismisses these theoretical possibilities as the work of "ingenious economists (properly) making their mark by

Economic growth would be expected to alleviate poverty, at least to some extent. How much depends on at least two factors. One is the growth rate itself. A study by Squire (1993) used an internationally comparable poverty line and regressed the rate of poverty reduction in a country against its rate of of economic growth. His results show that a one percent increase in the growth rate reduces the poverty headcount ratio by 0.24 percentage points. Similarly, Ravallion and Chen (1997) found that the larger the rate of change in log mean consumption or income in a country, the larger the decline in a country's log poverty rate. Furthermore, Roemer and Gugerty (1997) found that in the economic growth experiences of twenty-six developing countries, the rate of increase of the incomes of the poorest 20 percent was 92 percent of the rate of GDP growth, while for the poorest 40 percent, the rate of increase was essentially identical to the rate of GDP growth. Thus, in all these studies, faster economic growth has been found to lead to greater poverty reduction.

The rate of poverty reduction would also be expected to depend on the extent of economic inequality. In a very straightforward statistical sense, we would expect to find that economic growth reduces poverty by more if inequality falls than if it does not. This expectation is confirmed in cross-country analysis carried out by Bruno, Ravallion, and Squire (1998). For twenty countries during the period 1984–1993, they regressed the rate of change in the proportion of the population living on less than $1 per person per day against the rate of change in real mean income and obtained a regression coefficient of −2.12 with a t ratio of −4.67. This means that a 10 percent increase in the mean can be expected to produce a roughly 20 percent drop in the proportion of people in poverty. And when the P_2 measure is used instead, the effect is even greater: −3.46 with a t ratio of −2.98. They conclude: "Absolute poverty measures typically respond quite elastically to growth, and the benefits are certainly not confined to those near the poverty line."

To see what role inequality plays in poverty reduction, these authors also ran a multiple regression of the change in poverty (using the proportion of the population living on less than $1 per person per day) as a function of both the change in the mean and the change in inequality as measured by the Gini coefficient and found a coefficient of −2.28 (t = −6.07) on the former and 3.86 (t = 3.20) on the latter. In their words (p. 11): "Measured changes in inequality do have a strong independent explanatory power; indeed, rates of poverty reduction respond even more elastically to rates of change in the Gini index than they do to the mean."

The Bruno-Ravallion-Squire findings show that the change in poverty is related both to economic growth and to changing inequality. First, holding the dispersion of income the same, the faster is the rate of economic growth, the larger is the reduction in poverty. Second, for any given growth rate, the more dispersed the distribution is becoming, the smaller is the reduction in poverty.[4]

The relative importance of these growth and redistributive factors for poverty reduction can be gauged using the following equation devised by Ravallion and Huppi (1991); Datt and Ravallion (1992); and Kakwani (1993). Let poverty in a country at time t be denoted by

$$P_t = P(z/\mu_t, D_t),$$

where z is the poverty line, μ_t is mean expenditure per capita, and D_t is the inequality in the distribution of expenditure per capita. Then the change in poverty between a base year B and a terminal year T can be written as

$$P_T - P_B = \underbrace{P(z/\mu_T, D_B) - P(z/\mu_B, D_B)}_{\text{Growth component}} + \underbrace{P(z/\mu_B, D_T) - P(z/\mu_B, D_B)}_{\text{Redistribution component}} + \text{residual}.$$

Applying this methodology to developing countries in Africa, Latin America, and East and Southeast Asia, Demery, Sen, and Vishwanath (1995) came to two conclusions. First, poverty change is largely determined by economic growth. Second, changes in inequality are of secondary importance in the great majority of cases.

An important qualifier needs to be added. In countries' actual experiences, it has proved far easier to generate economic growth than to change the Gini coefficient. In the developing world, GDP per capita grew by 26 percent between 1985 and 1995 (World Bank 1997), while Gini coefficients in the world barely changed over the same period (Deininger and Squire 1998, table 5). In a similar vein, Adelman and Robinson's simulation results showed that even huge changes in policy parameters (such as a doubling of the tax rate,

4. In cases of economic decline, poverty rises by more the more negative is the rate of economic growth and the larger is the rate of increase of dispersion.

increasing agricultural capital stocks by 30 percent, fixing all agricultural prices at world prices, and subsidizing the consumption of food, housing, and medical services for the poorest 60 percent of households) would change the Gini coefficient in Korea by only one or two Gini points in most cases. The point is that in comparing the elasticities of poverty with respect to growth and with respect to inequality, one should not fall into the trap of thinking that it is as easy to lower inequality by 10 percent as it is to achieve 10 percent growth; the former is far more difficult than the latter.

5.3 Poverty Reduction and Growth: Individual Country Experiences

The data presented in the previous section showed that the predominant tendency is for poverty to fall when economic growth takes place. This happens in the great majority of cases. Consequently, the shared growth view is clearly a better general description of the growth/poverty relationship than is the immiserizing growth position.

Not only does this conclusion hold in general but it holds for every region of the world. In Asia, the "big five" countries—China, India, Indonesia, Pakistan, and Bangladesh, which together have three-fifths of the developing world's people and two-fifths of the poor—have all made "impressive progress" in reducing income poverty (United Nations 1997). Ahuja et al. (1997) report that "poverty has been declining in every East Asian economy for which we have data except Papua New Guinea."[5] Unfortunately, the crisis of the late 1990s in a number of East and Southeast Asian countries has demonstrated that the relationship between growth and poverty holds in times of economic decline as well. Although direct data on poverty are not yet available for the post-crisis period, the dramatic increase in unemployment caused by the crisis is surely leading to an increase in poverty (ILO 1998a; Manuelyan-Atinc and Walton 1998).

In Latin America and the Caribbean (LAC), data compiled by Londoño and Székely (1998) show that poverty rose during the 1980s and fell slightly during the 1990s. The 1980s was a "lost decade" for

5. Papua New Guinea's per capita GDP grew at about a 2 percent annual rate during this period of time.

The countries covered in their study are China, Malaysia, Thailand, Indonesia, the Philippines, Papua New Guinea, Lao PDR, Vietnam, and Mongolia.

Latin America, the growth rate in the region being a negative 1.2 percent (Inter-American Development Bank 1991). In the 1990s, the region experienced a modest recovery: a 6 percent increase (total, not per annum) in real GDP per capita and a 4 percent increase in private consumption. Changes in poverty within countries mirror this trend: During the negative growth decade, the poverty headcount ratio rose in ten of the thirteen countries in Latin America and the Caribbean for which we have data, whereas in the recovery, it fell in nine of the thirteen.[6]

In Africa, the experiences that have been documented also show that when growth has taken place, poverty has fallen. This was the case in Morocco and Tunisia in the latter half of the 1980s (Chen, Datt, and Ravallion 1993), in Ghana from 1987/1988 to 1991/1992 (World Bank 1995; Ghana Statistical Service 1995), in Nigeria from 1985 to 1992 (Canagarajah, Nwafon, and Thomas 1995), and in rural Ethiopia from 1989 to 1994 (Dercon, Krishnan, and Kello 1994; Dercon and Krishnan 1995). However, there were also negative results. In Kenya, lack of economic growth resulted in a constant poverty headcount ratio and continued poverty for an increasing number of people (Mukui 1994; World Bank 1996a). In Tanzania, real per capita income of the poorest 40 percent fell by 28 percent between 1983 and 1991 (Ferreira 1993, cited in Wangwe 1996)—a time during which the economy was contracting at the rate of 1 percent per year. And in Côte d'Ivoire, poverty increased during the 1985–1988 recession, using the poverty headcount ratio, the normalized poverty deficit (P_1) and the squared poverty gap (P_2) for two alternative poverty lines (Grootaert 1994). Note that in each of these cases where poverty did not fall, growth did not take place.

Finally, Eastern Europe and the countries of the former Soviet Union experienced dramatic economic contractions of the early 1990s. While there has been some growth since then in the Eastern European countries, the decline has continued in the countries of the Former Soviet Union. Along with the decline in income, the transition from socialism has also brought increasing inequality. Together,

6. Poverty rose in the 1980s and fell in the 1990s in Brazil, Guatemala, Honduras, Panama, Peru, and Venezuela, in each case reflecting the negative economic growth of the 1980s and the positive economic growth of the 1990s.

Another study that also reaches the conclusion that economic growth usually reduced poverty in Latin American countries is that by Ganuza, Morley, and Taylor (1998).