1

The Growth Elasticity of Poverty Reduction: Explaining Heterogeneity across Countries and Time Periods

François Bourguignon

1.1 Introduction

Part of the ongoing debate on poverty reduction strategies bears on the issue of the actual contribution of economic growth to poverty reduction. There is no doubt that faster economic growth is associated with faster poverty reduction. But what is the corresponding elasticity? If it is reasonably high, then poverty reductions strategies almost exclusively relying on economic growth are probably justified. If it is low, however, ambitious poverty reduction strategies might have to combine both economic growth and some kind of redistribution. Ravallion and Chen (1997) estimated that, on average on a sample of developing countries, the growth elasticity of poverty, as measured by the number of individuals below the conventional $1-a-day threshold, was around 3—namely, a 1 percent increase in mean income or consumption expenditures in the population reduces the proportion of people living below the poverty line by 3 percent. As emphasized in *Attacking Poverty, World Development Report 2000/2001*, however, there is very much cross-country heterogeneity behind this average figure—which, as a matter of fact, was found there closer to 2 than 3.[1] Several countries knew only limited changes in poverty despite satisfactory growth performances whereas poverty fell in some countries where growth had yet been disappointing. Understanding the causes of that heterogeneity is clearly crucial for the design of poverty reduction strategies.

To get some idea of the actual heterogeneity in the relationship between changes in poverty and changes in income, and of the ambiguity of average cross-sectional data, figure 1.1 plots observations that come from a sample of growth spells taken over various periods in selected countries. These spells are essentially defined by

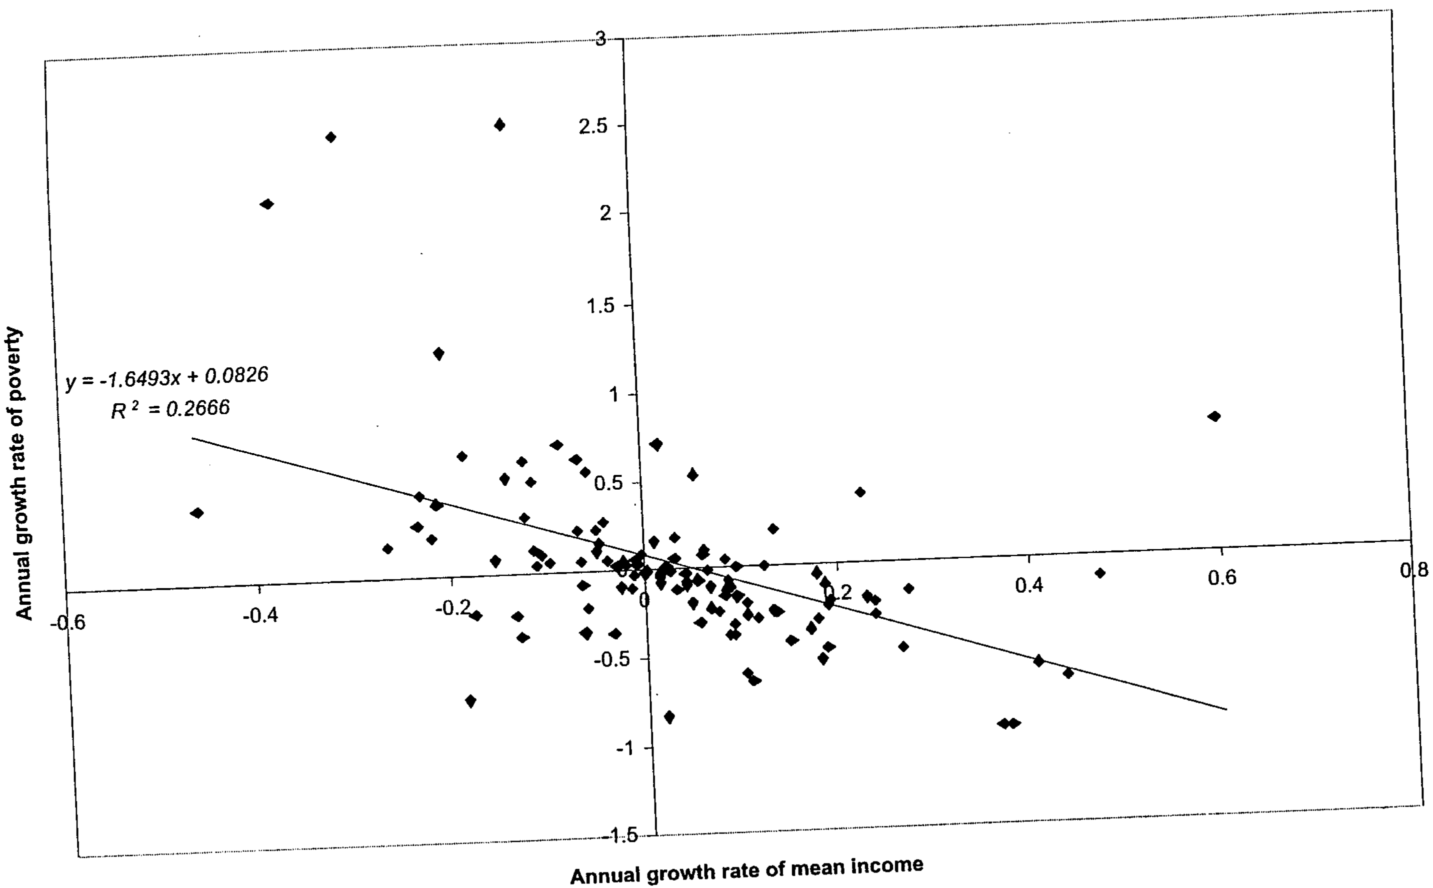

Figure 1.1
The relationship between poverty reduction and growth in a sample of growth spells

the availability of fully comparable household surveys at the two ends of the spell, only the most distant ends being considered in the case of adjacent annual observations. The sample comprises 114 spells covering approximately 50 countries—see the list in the appendix. The common poverty line that is used is the $1-a-day line after local currency expenditures data have been translated into dollars and PPP correction.[2] The poverty measure being used is the headcount ratio, that is the proportion of the population below the poverty line. As expected, the scatter of observations shows a declining relationship between the change in poverty and the change in the mean income. The fitted OLS straight line has a slope equal to −1.6, suggesting an average elasticity somewhat below the value of 2 reported in World Bank (2000). Differences in the mean elasticity in different samples of countries or growth spells is not the issue, however. The issue is the difference across countries or growth spells. From that point of view figure 1.1 is disappointing. Changes in the mean income of the population explain only 26 percent of the variance of observed changes in poverty headcounts. If this figure is taken seriously, would it make sense to base poverty reduction so much on growth strategies, as is often suggested? Wouldn't it be better to identify first the nature of the remaining 74 percent and then the reason why poverty in various countries tend to react very differently to the same increase in the mean income of the population?

Analytically, an *identity* links the growth of the mean income in a given population, the change in the distribution of relative incomes, and the reduction of poverty. Formally, the relationship between poverty and growth may be obtained from that identity in the case where there would be no change in the distribution of relative individual incomes, or, in other words, if income growth were the same in all segments of society. Even in that case, however, the growth-poverty relationship is not simple and the corresponding elasticity is certainly not constant across countries and across the various ways of measuring poverty. In effect, the growth elasticity of poverty is a decreasing function of the development level of a country and of the degree of inequality of the income distribution, this function depending itself on the poverty index that is being used. A rather precise characterization of that relationship is offered in this chapter under some simplifying assumption about the underlying distribution of income.

The second source of heterogeneity is of course the change in the distribution of relative incomes over time. Measuring the actual contribution of that source to the observed evolution of poverty is of utmost importance since this should give an indication of the practical relevance of distributional concerns in comparison with pure growth concerns in poverty reduction policies. Because the actual contribution of growth to poverty reduction is not precisely identified by the practice that consists of assuming a constant elasticity, it follows that the contribution of the distributional component is also imprecisely estimated. The methodology proposed in this chapter permits correcting that imprecision too and has implications for the interpretation of available evidence on the growth-poverty relationship.

Many recent papers focused on the statistical relationship between economic growth and poverty reduction across countries and time periods. Many of them—see, for instance, de Janvry and Sadoulet (1995, 2000), Ravallion and Chen (1997), Dollar and Kraay (2000)—are based on linear regressions where the evolution of some poverty measure between two points of time is explained by the growth of income or GDP per capita and a host of other variables, the main issue being the importance of GDP and these other variables in determining poverty reduction. By adopting a linear regression framework, or by investing too little in functional specification testing, however, these papers miss the earlier point, that is that of a complex but yet identity-related relationship between mean income growth and poverty change. On the contrary, other authors —for instance, Ravallion and Huppi (1991), Datt and Ravallion (1992), Kakwani (1993)[3]—fully take into account the poverty/mean-income/distribution identity in studying the evolution of poverty and its causes. In particular, they are all quite careful in distinguishing precisely the effects on poverty reduction of growth on the one hand and distributional changes on the other. At the same time, their analysis is generally restricted to a specific country or a limited number of countries or regions: Indonesia, regions of Brazil and India, the Cote d'Ivoire, respectively.

This chapter stands midway between these two approaches. Ideally, international comparisons in the evolution of poverty should all rely on the methodology based on the poverty/mean-income/distribution identity. But, because it requires using the full microeconomic information on the distribution of income or expenditures in each country or region, it may seem cumbersome. Instead, this chapter proposes a methodology that is less demanding. It relies on functional approximations of the identity, and in particular on an approximation based on the assumption that the distribution of income or consumption expenditure is lognormal.

A simple application to the sample of growth spells shown in figure 1.1 shows that these approximations fit the data extremely well and do incomparably better than the linear model that is generally used. It also suggests that only half of the observed changes in poverty in the sample may be explained by economic growth, the remaining half being the result of changes in the distribution of relative incomes.

The chapter is organized as follows. Section 1.2 introduces and discusses the analytical identity that links poverty, growth, and distribution. Closed-form formulae for the growth elasticity of poverty when the distribution is lognormal are derived. They are used to analyze in some detail the theoretical relationship that exists among poverty reduction, the level of development, and the inequality of the distribution when the distribution of income is remaining constant over time. Section 1.3 tests the empirical validity of various approximations to the preceding identity, including the lognormal approximation, and compares the results with the standard linear specification. A concluding section 1.4 draws several implications of the main argument in the chapter for the empirical analysis of the relationship between growth and poverty and for the design of poverty reduction strategies.

1.2 The Arithmetic of Distributional and Poverty Changes

Any poverty index may be seen as a statistics defined on all individuals in a population whose standard of living lies below some predetermined limit. In what follows, it is assumed that there is no ambiguity on the definition of that "poverty line"; that is, whether it is defined in terms of income or consumption, the kind of equivalence scale being used to account for heterogeneity in household composition, and indeed the level of that poverty line. This poverty line will also be assumed to be constant over time—at least during the period being analyzed. Since the argument will implicitly refer to international comparisons, it makes also sense to assume that this poverty line is defined in "absolute" terms and the same across countries, for instance, the familiar $1 or $2 a day after correction for purchasing power parity.

Given this definition of poverty, let y be a measure of individual living standard—say, income per adult equivalent—and let z be the poverty line. In a given country, the distribution of income at some point of time, t, is represented by the cumulative distribution function $F_t(Y)$, which stands for the proportion of individuals in the population with living standard, or income, y, less than Y. The most widely used poverty index is simply the proportion of individuals in the population below the poverty line, z. This index is generally referred to as the "headcount." With the preceding notations, it may be formally defined as

$$H_t = F_t(z). \tag{1}$$

For the sake of simplicity, the analysis will be momentarily restricted to that single poverty index.

The definition of the headcount poverty index implies the following definition of change in poverty between two points of time, t and t':

$$\Delta H = H_{t'} - H_t = F_{t'}(z) - F_t(z).$$

To show the contribution of growth to the change in poverty, it is convenient to define the distribution of *relative income* at time t as the distribution of incomes after normalizing by the population mean. This is equivalent with defining the distribution of income in a way that is independent of the scale of incomes. Let $\tilde{F}_t(X)$ be that distribution. With this definition, any change in the distribution of income may then be decomposed into (a) a proportional change in all incomes that leaves the distribution of relative income, $\tilde{F}_t(X)$, unchanged; and (b) a change in the distribution of relative incomes, which, by definition, is independent of the mean. For obvious reasons, the first change will be referred to as the "growth" effect whereas the second one will be termed the "distributional" effect.

This decomposition was discussed in some detail by Datt and Ravallion (1992) and Kakwani (1993). It is illustrated in figure 1.2. This figure shows the density of the distribution of income—that is, the number of individuals at each level of income represented on a logarithmic scale on the horizontal axis. In that figure, the function $F(\ )$ appears only indirectly as the area under the density curves. The move from the initial to the new distribution goes through an

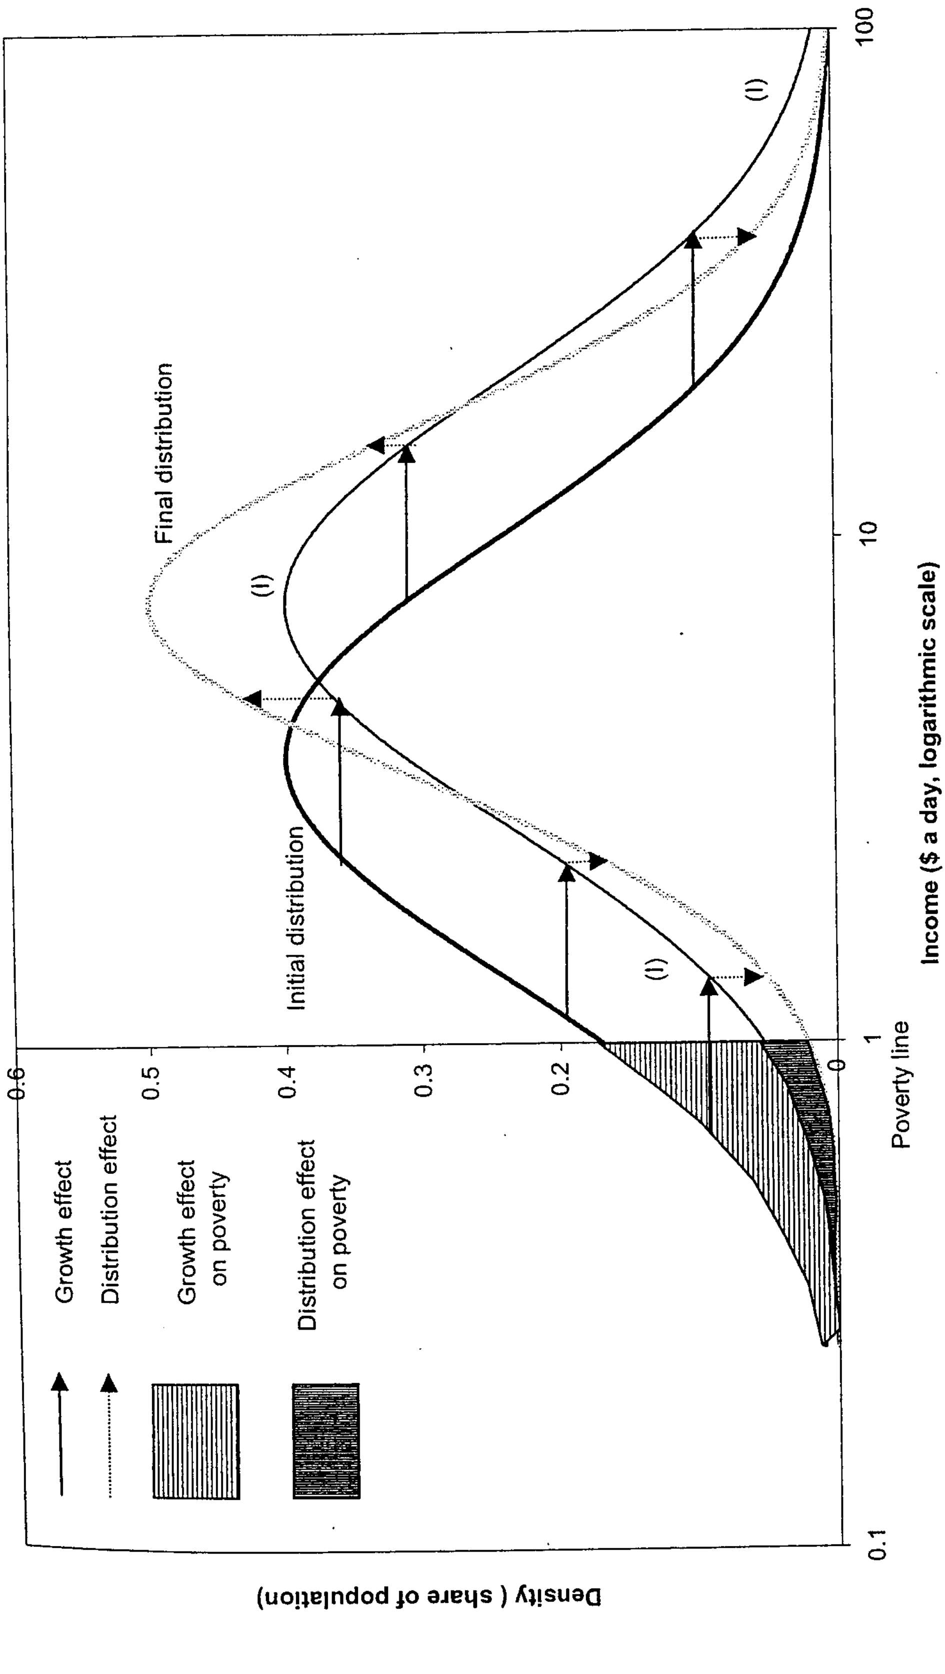

Figure 1.2
Decomposition of change in distribution and poverty into growth and distributional effects

intermediate step, which is the horizontal translation of the initial density curve to curve (I). Because of the logarithmic scale on the horizontal axis, this change corresponds to the same proportional increase of all incomes in the population and thus stands for the "growth effect." Then, moving from curve (I) to the new distribution curve occurs at constant mean income. This movement thus corresponds to the change in the distribution of "relative" income, or the "distribution" effect. Of course, there is some path dependence in that decomposition. Instead of moving first rightward and then up and down as in figure 1.2, it would have been possible to move first up and down and to have the distribution effect based on the mean income observed in the initial period, and then to move rightward. Presumably, these two paths are not necessarily equivalent except for infinitesimal changes. This issue is ignored in what follows, assuming in effect that all changes are sufficiently small for path dependence not to be a problem.[4]

Figure 1.2 illustrates the natural decomposition of the change in the whole distribution of income between two points of time. Things are somewhat simpler if the focus is exclusively on poverty, as measured by the headcount. In figure 1.2, the poverty headcount is simply the area under the density curve at the left of the poverty line, arbitrarily set at $1 a day. The move from the initial curve to the intermediate curve (I) and then from that curve to the new distribution curve have natural counterparts in terms of changes in this area. More formally, this decomposition may be written as

$$\Delta H = H_{t'} - H_t = \left[\tilde{F}_t\left(\frac{z}{\bar{y}_{t'}}\right) - \tilde{F}_t\left(\frac{z}{\bar{y}_t}\right)\right] + \left[\tilde{F}_{t'}\left(\frac{z}{\bar{y}_{t'}}\right) - \tilde{F}_t\left(\frac{z}{\bar{y}_{t'}}\right)\right]. \quad (2)$$

This expression is the direct application in the case of the headcount poverty index of the general formula proposed by Datt and Ravallion (1992), which they applied to Brazilian and Indian distribution data. It is indeed a simple *identity* since it consists of adding and subtracting the same term $\tilde{F}_t(z/\bar{y}_{t'})$ in the original definition of the change in poverty. The first expression in square bracket in (2) corresponds to the growth effect at "constant" relative income distribution, $\tilde{F}_t(\ )$, that is the translation of the density curve along the horizontal axis in figure 1.2, whereas the second square bracket formalizes the distribution effect—that is, the change in the relative income distribution, $\tilde{F}_{t'}(X) - \tilde{F}_t(X)$, at the new level of the "relative" poverty line, that is the ratio of the absolute poverty line and the mean income, $X = z/\bar{y}_{t'}$.

Shifting to elasticity concepts, the growth elasticity of poverty may thus be defined as

$$\varepsilon = \lim_{t' \to t} \frac{\left[\tilde{F}_t\left(\frac{z}{\bar{y}_{t'}}\right) - \tilde{F}_t\left(\frac{z}{\bar{y}_t}\right)\right] \Big/ \tilde{F}_t\left(\frac{z}{\bar{y}_t}\right)}{(\bar{y}_{t'} - \bar{y}_t)/\bar{y}_t}. \qquad (3)$$

The distribution effect is more difficult to translate in terms of elasticity because it generally cannot be represented by a scalar.

The terms entering the decomposition identity (2) may be evaluated as long as one observes some continuous approximation of the distribution functions $F(\ )$ at the two points of time t and t'. Continuous kernel approximations of the density and cumulative relative distribution functions may be computed from available microeconomic data. With this kind of tool, evaluating the decomposition identity for any growth spell for which distribution data are available at the two ends of the spell should not be difficult. In a cross-country framework, however, this might require manipulating a large number of microeconomic data sets and may be found cumbersome.

Interestingly enough, a very simple approximation of (2) may be obtained in the case where the distributions may be assumed to be lognormal, probably the most standard approximation of empirical distributions in the applied literature. The relative income distribution writes in that case:

$$\tilde{F}_t(X) = \Pi\left[\frac{\log(X)}{\sigma} + \frac{1}{2}\sigma\right],$$

where $\Pi(\ )$ is the cumulative distribution function of the standard normal and σ is the standard deviation of the logarithm of income. Substituting this expression in (2) shows that the change in poverty headcount between time t and t' depends on the level of mean income at these two dates, $\bar{y}/z$, expressed as a proportion of the poverty line, and on the standard deviation, σ, of the logarithm of income at the two dates. Allowing t' to be close to t and taking limits as in (3) then leads to

$$\frac{\Delta H}{H_t} = \lambda\left[\frac{\log(z/\bar{y}_t)}{\sigma} + \frac{1}{2}\sigma\right] \cdot \left[-\frac{\Delta \log(\bar{y})}{\sigma} + \left(\frac{1}{2} - \frac{\log(z/\bar{y}_t)}{\sigma^2}\right)\Delta\sigma\right], \qquad (2')$$

where $\lambda(\ )$ stands for the ratio of the density to the cumulative function—or hazard rate—of the standard normal, $\Delta \log(\bar{y})$ is the growth rate of the economy and $\Delta\sigma$ is the variation in the standard deviation of the logarithm of income. Based on that expression and following (3), the growth elasticity of poverty, ε, may be defined as the relative change in the poverty headcount for 1 percent growth in mean income, for constant relative inequality, σ:

$$\varepsilon = \frac{\Delta H}{\Delta \log(\bar{y})H_t} = \frac{1}{\sigma}\lambda\left[\frac{\log(z/\bar{y}_t)}{\sigma} + \frac{1}{2}\sigma\right]. \tag{3'}$$

In that expression, the growth elasticity of poverty appears explicitly as an increasing function of the level of development, as measured by the inverse of the ratio $z/\bar{y}_t$, and a decreasing function of the degree of relative income inequality as measured by the standard deviation of the logarithm of income, σ.[5]

The preceding relationship is represented in figure 1.3 by curves in the development inequality space along which the growth elasticity of poverty is constant. The inverse of the development level is measured along the horizontal axis by the ratio of the poverty line to the mean income of the population. Inequality is measured along the vertical axis by the Gini coefficient of the distribution of relative income rather than the standard deviation of logarithm. This measure of inequality is more familiar than σ but is known to be an increasing function of it.[6]

Figure 1.3 is useful to get some direct and quick estimate of the growth elasticity of poverty. Consider, for instance, the case of a poor country where the mean income is only twice the poverty line at the right end of the figure. Reading the figure, it may be seen that the growth elasticity is around 3 if inequality is low—namely, a Gini coefficient around 0.3—but it is only 2 if the Gini coefficient is around the more common value of 0.4. If the economy gets richer, then the elasticity increases. But at the same time it becomes more sensitive to the level of inequality. For instance, when the mean income of the population is four times the poverty line, the growth elasticity of poverty is 5 for low-income inequality—namely, a Gini equal to 0.3—but 2 again if inequality is high—namely, a Gini coefficient equal to 0.5. These various combinations are illustrated by the example of a few countries in the mid-1980s. The ratio of the poverty line to the mean income is taken to be the \$1-a-day line related to

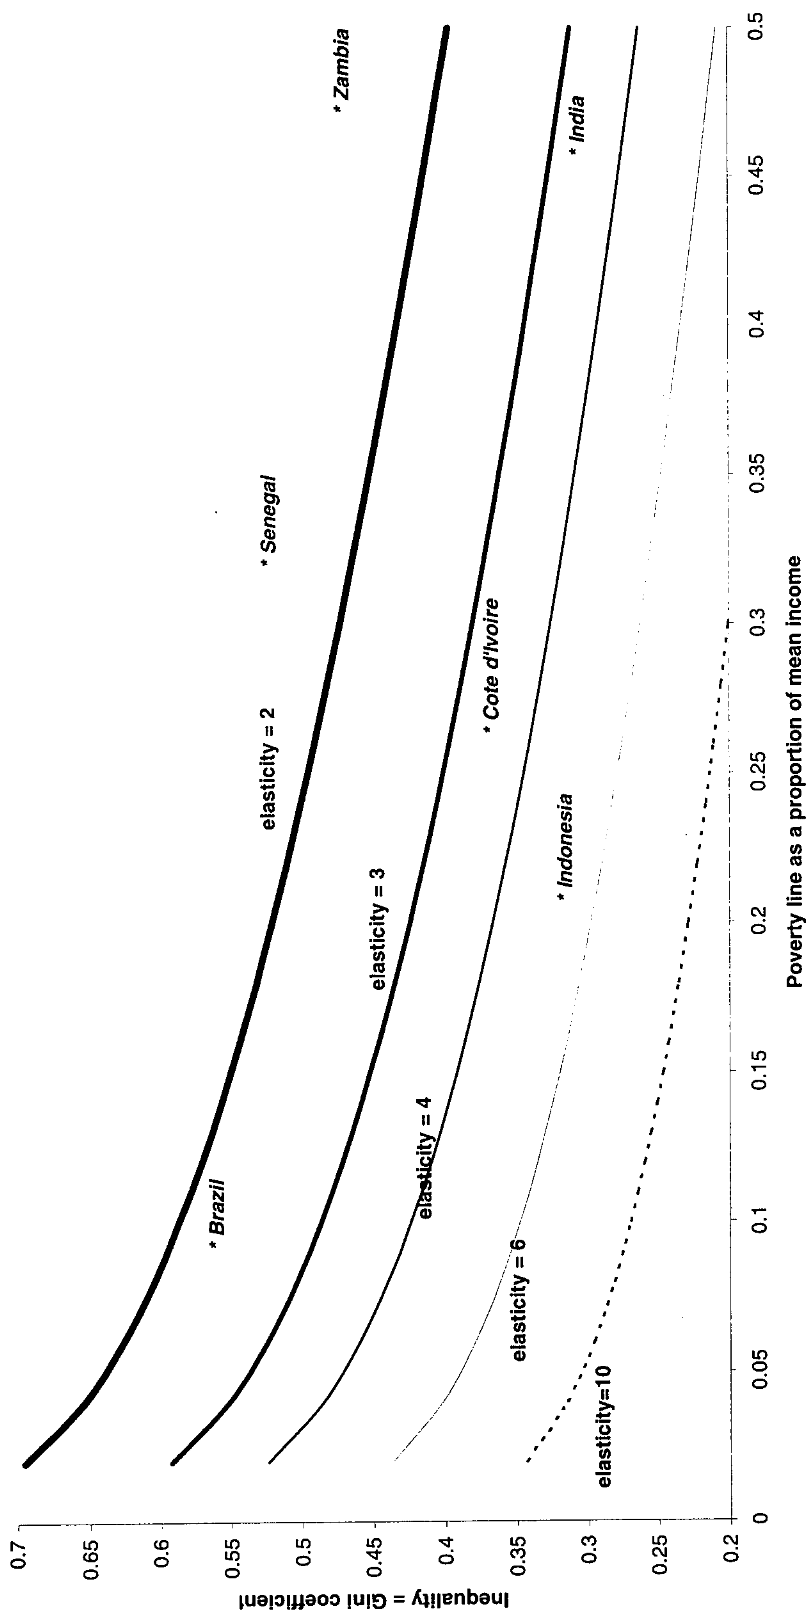

Figure 1.3
Poverty (headcount)/growth elasticity as a function of mean income and income inequality, under the assumption of constant (lognormal) distribution

GDP per capita, whereas Gini coefficients are taken from the Deininger and Squire (1996) database. Growth elasticities of poverty reduction consistent with the lognormal assumption vary between 5 for a country like Indonesia to 3 for India and the Cote d'Ivoire, which were poorer and/or had a higher level of inequality. The elasticity is around 2 for Brazil, despite the fact that it is considerably richer than the other countries. Income inequality makes the difference. Finally, the elasticity is much below 2 in the case of Senegal and Zambia, both of which are poor and unequal.

Figure 1.3 and the underlying lognormal approximation to the growth elasticity of poverty in (5) refers to the headcount as the poverty index. Similar expressions and curves may be obtained using alternative indices. For instance, the poverty gap is a measure of poverty obtained by multiplying the headcount by the average relative distance at which the poor are from the poverty line. The advantage of that measure over the headcount is obviously that it takes into account not only the proportion of people being poor but also the intensity of poverty. Under the lognormal assumption, it may be shown that the growth elasticity of the poverty gap is given by the following formula:[7]

$$\varepsilon_y^{PG} = -\frac{\Pi[\log(z/\bar{y}_t)/\sigma - \sigma/2]}{(z/\bar{y}_t) * \Pi[\log(z/\bar{y}_t)/\sigma + \sigma/2] - \Pi[\log(z/\bar{y}_t)/\sigma - \sigma/2]}, \qquad (4)$$

whereas the elasticity with respect to the standard deviation of the logarithm of income is given by

$$\varepsilon_\sigma^{PG} = -\frac{\sigma \cdot \pi[\log(z/\bar{y}_t)/\sigma - \sigma/2]}{(z/\bar{y}_t) * \Pi[\log(z/\bar{y}_t)/\sigma + \sigma/2] - \Pi[\log(z/\bar{y}_t)/\sigma - \sigma/2]},$$

where $\Pi(\ )$ and $\pi(\ )$ are respectively the cumulative distribution and the density functions of the standard normal variable. As before, the growth elasticity depends on the level of development as measured by the ratio $z/\bar{y}_t$ and the inequality of the distribution of income as given by the standard deviation of the logarithm of income. Although not shown here, the iso elasticity contours in the development/inequality space corresponding to function (4) are very similar to what was shown in figure 1.3 for the headcount index. Equivalent formulas and similar figures could be derived for other poverty measures, in particular those belonging to the well-known P_α family—see Foster, Greer, and Thorbecke (1984).[8]

1.3 Revisiting the Empirical Evidence on the Growth-Poverty Relationship

The identity linking growth, poverty reduction, and distributional changes should be readily apparent from data on growth-poverty spells. This section shows that this is indeed the case, provided the adequate specification is used. In particular the lognormal approximation proves to be extremely precise.

The data being used are the same as those used for figure 1.1. The sample includes mostly developing countries but a few transition economies are also present. The only modification made to the original data was to eliminate all spells where the percentage change in the poverty headcount was abnormally large in relative value. Cases left aside essentially correspond to situations where the poverty headcount went from zero or an almost negligible figure to some positive value, or the opposite, in a few years. This was made in order to comply with the definition of the elasticity concepts given earlier and the corresponding requirement of "small" changes.[9] The mean annual growth in mean income over all these spells is 2.7 percent, the mean relative change in the poverty headcount—with a $1-a-day poverty line—is close to zero and the mean change in the Gini coefficient is −.0022.

Based on this data set, four different models are compared against each other. The first corresponds to the naive view alluded to above that there is a constant elasticity between poverty reduction and growth. It consists of regressing observed changes in the poverty headcount on observed changes in mean income. The second model, which is termed the "standard" model in table 1.1 includes the observed change in income inequality, as measured by the Gini coefficient as an additional explanatory variable. It is thus consistent with a decomposition of type (2) above, except for the fact that both the growth elasticity and the Gini elasticity of poverty reduction are taken to be constant. The third model improves on the previous one by allowing the growth elasticity to depend on the inverse level of development, as measured by the poverty-line/mean-income ratio, and on the initial degree of inequality as measured by the Gini coefficient. Nothing is imposed a priori on that relationship, though. Finding the optimal functional specification is left to econometrics. To do so, two additional variables are introduced in the regression: the interaction between growth and the preceding two variables. The final model relies on the lognormal approximation discussed earlier. The explanatory variables are the theoretical elasticity defined in (3′) times the observed growth rate of mean income[10] and the change in the Gini coefficient. If the lognormal approximation discussed earlier is not too unsatisfactory, then one should find that the coefficient of the theoretical elasticity in that regression is not significantly different from unity.

The estimation of these four models is reported in the first four columns of table 1.1. Results fully confirm the identity relationship discussed in this chapter. First, one can see that the naive model suggests a significant negative elasticity of poverty with respect to growth but its explanatory power is low. This strictly corresponds to the fitted line in figure 1.1 with a R^2 equal to 26 percent. Things improve quite substantially when one shifts to the standard model by adding distributional changes in the regression equation. Adding the change in the Gini coefficient in the linear specification practically doubles the R^2 coefficient, suggesting that the heterogeneity in distributional changes is as much responsible for variation in poverty reduction across growth spells as the heterogeneity in growth rates itself.

Interacting growth with the initial poverty-line/mean-income ratio and the initial Gini coefficient in the "improved standard model 1" yields still another significant improvement in explanatory power. The two interaction terms are very significant and go in the right direction. As expected, both a lesser level of development and a higher level of inequality reduce the growth elasticity of poverty. Both effects are significant and sizable. At the mean point of the sample—namely, for a growth spell leading to an annual 2.7 percent rise in mean income—an increase of the initial level of development by one standard deviation of the poverty-line/mean-income variable increases poverty reduction by some 3 percentage points annually. In the same conditions, an increase in the initial Gini coefficient by one standard deviation diminishes poverty reduction by a little less than 1 percentage point.

No assumption is made in the preceding regression on the way income growth, the development level and the initial degree of inequality interact to determine poverty reduction. In the fourth column of table 1.1, on the contrary, it is assumed that the joint effect of these three variables is in accordance with the theoretical elasticity

Table 1.1
Explaining the evolution of poverty across growth spells[a]
(dependent variable = percentage change in poverty headcount during growth spell)

Explanatory variables	Naive model (1)	Standard model (2)	Improved standard model 1 (3)	Identity check: Lognormal model 1 (4)	Improved standard model 2 (5)	Identity check: Lognormal model 2 (6)
Intercept	0.0826	0.0972	0.0837	0.0752	0.0977	0.0692
	0.0434	*0.0364*	*0.0349*	*0.0325*	*0.0321*	*0.0281*
Y = percentage change in mean income	−1.6493	−2.0124	−6.3518		−7.8706	
	0.2585	*0.2223*	*1.2451*		*1.1310*	
DGini = Variation in Gini coefficient		4.7178	5.2863	5.0769	21.5606	
		0.6731	*0.6529*	*0.6118*	*4.1205*	
Y * poverty-line/mean-income			3.9678		3.9481	
			1.1662		*1.0286*	
Y * initial Gini coefficient			7.0039		9.6869	
			2.4586		*2.2104*	
DGini * poverty-line/mean-income					−16.3898	
					2.8255	
DGini * initial Gini coefficient					−20.3600	
					7.4388	
Y * theoretical value of growth elasticity under lognormal assumption				−0.8727		−0.9261
				0.0778		*0.0679*
DSigma * theoretical value of poverty inequality elasticity under lognormal assumption[b]						0.6824
						0.0601
R^2	0.2666	0.4916	0.555	0.5857	0.6651	0.6892

[a] Ordinary least squares estimates, standard errors in italics. The sample includes the 114 growth spells listed in the appendix. All coefficients are significantly different from zero at the 1 percent probability level except the intercept.
[b] Switching from the coefficient of Gini to the standard deviation of the logarithm of income is done through the formula in note 6.

derived in the preceding section under the assumption that the underlying distribution of relative income is lognormal. The resulting explanatory variable in the poverty reduction regression thus is this theoretical elasticity times the observed growth of the mean income. This test of the identity that links poverty reduction and growth is very successful. On the one hand, the R^2 coefficient proves to be substantially higher than when the various explanatory variables are entered without functional restriction as in the preceding regression, despite less degrees of freedom being used. On the other hand, the coefficient of the theoretical value of the elasticity proves to be only slightly below unity. Overall, it must then be concluded that the best "single" explanation of observed poverty reduction in a sample of growth spells is indeed provided by the identity that logically links poverty and growth under the lognormal assumption.

In all the preceding regressions, no care is taken of the fact that, according to the identity discussed above, the role of the change in inequality on poverty reduction is unlikely to be linear. The last two columns of table 1.1 correct the preceding results for this. In column (5) the effect of the distributional change on poverty reduction is assumed to depend both on the initial level of development and the initial level of inequality. In column (6), the two explanatory variables are the growth and inequality elasticities, as defined in (2′) under the lognormal assumption, multiplied respectively by the relative change in mean income and the change in the Gini coefficient. The "improved standard model 2" in column (5) proves to do much better than the improved standard model 1 in column (3) where distributional changes were only captured by the change in the Gini coefficient. The R^2 coefficient gains more than 10 percentage points, solely because of the interaction terms between the change in the Gini coefficient and the two key level variables, the initial inequality level and the initial development level. Surprisingly enough, still a better score is achieved in column (6) with the lognormal approximation with only two explanatory variables that stand for the growth and inequality elasticity of poverty. It may be seen, moreover, that the coefficient of the growth-elasticity term is not significantly different from unity, something that would seem to be in favor of the lognormal approximation. Unfortunately, this is not true of the inequality-elasticity variable, which in effect would lead to reject that approximation. Note, however, that an additional hypothesis that cannot be tested independently of the lognormality is the infinitesimal approximation behind all types of elasticity calculation. As may be seen in figure 1.1, some observations are such that the annual absolute relative change in the mean income of the population is greater then 25 percent. It is not clear that the standard elasticity calculation would apply in this case. Yet, no attempt has been made in this chapter to use nondifferential expressions.

That the lognormal approximation plays an effective role in the R^2 of the last column of table 1.1 being still far from unity is confirmed when the preceding exercise is repeated with the poverty gap, rather than the poverty headcount. Table 1.2 is the equivalent of table 1.1 with the poverty index being the poverty gap rather than the poverty headcount. The striking feature there is of course that all R^2 are much lower than what was found with the poverty headcount, and also that the improvement in the fit of poverty reduction observations due to the approximations discussed in this chapter is less dramatic.

This last set of results may be somewhat disappointing but they are easily understandable. Using the headcount leads somehow to make predictions or simulations that concerns the value of the cumulative distribution function of income, $\tilde{F}_t(X)$, at a single point, namely, $X = z/\bar{y}_t$. Using the poverty gap, on the contrary, requires using predictions on the mean value of income for all people below $z/\bar{y}_t$, that is about the full range of values of $F_t(X)$ below the value X. This is much more demanding. The lognormal approximation may not perform too badly at a specific point close to the poverty line. It may do much worse if all incomes below that value are to be taken into account too.

This experiment is interesting because it suggests that, if one wants to go beyond the poverty headcount in poverty measurement, then functional approximations to growth and distribution elasticities of poverty reduction may simply be unsatisfactory. Dealing with the issue of the determinants of poverty reduction will then require working with the full distribution of income or living standards rather than a few summary measures. This will probably prove to be the only satisfactory solution in the long run and the sooner poverty specialists will get used to dealing systematically with distribution data, rather than inequality or poverty summary measures, at the national level, the better it will be.

Table 1.2
Explaining the evolution of poverty across growth spells (poverty gap)[a]
(dependent variable = percentage change in poverty gap during growth spell)

Explanatory variables	Naive model (1)	Standard model (2)	Improved standard model 1 (3)	Identity check: Lognormal model 1 (4)	Improved standard model 2 (5)	Identity check: Lognormal model 2 (6)
Intercept	0.1434	0.1660	0.1683	0.1412	0.1823	0.1272
	0.0857	*0.0777*	*0.0796*	*0.0787*	*0.0812*	*0.0759*
Y = percentage change in mean income	−1.2388	−1.8003	−0.4101		−2.3059	
	0.5101	*0.4747*	*2.8388*		*2.8658*	
Variation in Gini coefficient		7.2962	7.1231	7.1781	27.8647	
		1.4375	*1.4885*	*1.4806*	*10.4411*	
Y * poverty-line/mean-income			−0.9647		−0.9486	
			2.6588		*2.6065*	
Y * initial Gini coefficient			−2.4774		0.8514	
			5.6055		*5.6010*	
DGini * poverty-line/mean-income					−20.0988	
					7.1598	
DGini * initial Gini coefficient					−26.8577	
					18.8496	
Y * theoretical value of growth elasticity under lognormal assumption				−0.4961		−0.5750
				0.1661		*0.1627*
DSigma * theoretical value of poverty inequality elasticity under lognormal assumption[b]						0.7035
						0.1218
R^2	0.05	0.229	0.2308	0.1938	0.2849	0.249

[a] Ordinary least squares estimates, standard errors in italics. The sample includes the 114 growth spells listed in the appendix. All coefficients are significantly different from zero at the 1 percent probability level except the intercept.
[b] Switching from the coefficient of Gini to the standard deviation of the logarithm of income is done through the formula in note 6.

1.4 Some Implications

This being said, the identity that links poverty reduction, mean income growth, and distributional change has several implications for policymaking and economic analysis in the field of poverty that are worth stressing.

On the policy side, it was shown in this chapter that this identity permits identifying precisely the potential contribution of growth and distributional change to poverty reduction. However, it also introduces a point that is often overlooked in the debate of growth vs. redistribution as poverty reduction strategies—an exception being Ravallion (1997). To the extent that growth is sustainable in the long run whereas there is a natural limit to redistribution, it may reasonably be argued that an effective long-run policy of poverty reduction should rely primarily on sustained growth. According to the basic identity analyzed in this chapter, however, income redistribution plays essentially two roles in poverty reduction. A permanent redistribution of income reduces poverty instantaneously through what was identified as the "distribution effect." But, in addition it also contributes to a permanent increase in the elasticity of poverty reduction with respect to growth and therefore to *an acceleration of poverty reduction* for a given rate of economic growth. This is quite independent of the phenomenon emphasized in the recent growth-inequality literature according to which growth would tend to be faster in a less inegalitarian environment.[11] If this were true, there would then be a kind of "double dividend" associated with redistribution policy since it would at the same time accelerate growth and accelerate the speed at which growth spills over onto poverty reduction.

From the point of view of economic analysis, the basic argument in this chapter has clear implications for the understanding of poverty reduction. The common practice of trying to explain the evolution of some poverty measure over time, or across various countries, as a function of a host of variables including economic growth has something tautological. The preceding section has shown that the actual growth elasticity of poverty reduction in a given country could be estimated with considerable precision, even under the log-normal approximation. Running a regression like the "standard model" above where growth appears among the regressors would thus make sense only in the case where one does not observe the actual determinants of the theoretical value of the growth-poverty elasticity, that is the level of development and inequality. This would seem very unlikely, though, and it must be admitted that the way poverty reduction depends on growth is in effect perfectly known. Under these conditions, the only thing that remains to be explained in the basic poverty reduction decomposition formula is the pure "distributional change" effect. Somehow, all the variables that may be added in the regression after the growth effect has been rigorously taken into account should track the change in the distribution and its effects on poverty. In other words, the only thing that poverty change regressions should try to do is really to identify the causes of distributional changes and their effects on poverty indices.

This last remark relates to an earlier point about the nature of poverty being analyzed. In an international context, it seems natural that cross-country comparisons of poverty reduction bears on an "absolute" concept of poverty—namely, $1 or $2 a day. However, most of the argument in this chapter can be reinterpreted as saying that changes in absolute poverty may be decomposed into changes in the mean income of the population and changes in "relative poverty," as measured for instance by the number of people below some fixed proportion of the mean or median income, or simply the poorest x percent of the population. Viewed in this way, the argument in the preceding paragraph basically says that, in understanding the evolution of absolute poverty, the main object of analysis should really be the evolution of relative poverty, as the effect of a change in the mean income on absolute poverty is practically tautological. Although they do not formulate it in this way, this is what Dollar and Kraay (2000) attempt to do by focusing on the mean income or the income share of the bottom 20 percent of the population. When they show that this income share does not seem to move in any systematic direction with growth across countries, they in effect validate the use of the identity relationship to compute the effect of growth on absolute poverty reduction.

1.5 Appendix

The countries and growth spells used in the empirical part of this chapter appear in the following table. Spells are defined by the initial and terminal years.

Table 1.A.1
Countries and spells in database

Country	Spell	Country	Spell
Algeria	88–95	Honduras	94–96
Bangladesh	84–85	India	83–86
Bangladesh	85–88	India	86–87
Bangladesh	88–92	India	87–88
Bangladesh	92–96	India	88–89
Brazil	85–88	India	89–90
Brazil	88–89	India	90–92
Brazil	89–93	India	92–94
Brazil	93–95	India	94–95
Brazil	95–96	India	95–96
Chile	87–90	India	96–97
Chile	90–92	Indonesia	84–87
Chile	92–94	Indonesia	87–90
China	92–93	Indonesia	90–93
China	93–94	Indonesia	93–96
China	94–95	Indonesia	96–99
China	95–96	Jamaica	89–90
China	96–97	Jamaica	93–96
China	97–98	Jordan	92–97
Colombia	88–91	Kazakhstan	93–96
Colombia	95–96	Kenya	92–94
Costa Rica	86–90	Kyrgyz Republic	93–97
Costa Rica	90–93	Lesotho	86–93
Costa Rica	93–96	Madagascar	80–93
Cote D'Ivoire	87–88	Malaysia	84–87
Cote D'Ivoire	88–93	Malaysia	87–89
Cote D'Ivoire	93–95	Malaysia	89–92
Dominican Rep	89–96	Malaysia	92–95
Ecuador	88–94	Mauritania	88–93
Ecuador	94–95	Mauritania	93–95
Egypt	91–95	Mexico	84–92
El Salvador	89–95	Mexico	89–95
El Salvador	95–96	Morocco	85–90
Estonia	93–95	Nepal	85–95
Ethiopia	81–95	Niger	92–95
Ghana	87–89	Nigeria	85–92
Ghana	89–92	Nigeria	92–97
Guatemala	87–89	Pakistan	87–90
Honduras	89–90	Pakistan	90–93
Honduras	92–94	Pakistan	93–96

Table 1.A.1
(continued)

Country	Spell	Country	Spell
Panama	89–91	Thailand	88 (2)–92
Panama	91–95	Thailand	92–96
Panama	95–96	Thailand	96–98
Panama	96–97	Trinidad and Tobago	88–92
Paraguay	90–95	Tunisia	85–90
Peru	94–96	Turkey	87–94
Philippines	85–88	Uganda	89–92
Philippines	88–91	Ukraine	95–96
Philippines	91–94	Venezuela	81–87
Philippines	94–97	Venezuela	87–89
Romania	92–94	Venezuela	89–93
Russia	93–96	Venezuela	93–95
Russia	96–98	Venezuela	95–96
Senegal	91–94	Yemen?	92–98
Sri Lanka	85–90	Zambia	91–93
Sri Lanka	90–95	Zambia	93–96
Thailand	81–88 (1)		

Notes

A preliminary version of this chapter was presented at the conference World Poverty: A Challenge for the Private Sector, organized by the Amsterdam Institute for International Development, Amsterdam, October 3–4, 2000. The present version benefited from comments made by participants to a seminar at CESIfo (Munich, May 2001) and by participants to LACEA 2001 in Montevideo. I am particularly indebted to Daniele Checchi, Gylfi Zoega, and an anonymous referee. I am also indebted to Martin Ravallion for making the data set used in this chapter available to me. Of course, I remain solely responsible for any remaining error. Views expressed in this chapter are those of the author and should not be attributed to the World Bank or any affiliated organization.

1. See World Bank (2000, 47).

2. For a full description of these data, see Chen and Ravallion (2000).

3. See also the presentation of that methodology in the short literature survey provided by Fields (2001).

4. Figure 1.2 provides a very handy general representation of distributional changes and has been used in a number of circumstances. See, for instance, Quah (chapter 2).

5. These properties derive from the fact that $\lambda(\ )$ is known to be an increasing function in the case of the standard normal distribution. Indeed it may be shown that $\lambda(x)$ is the incomplete mean of the standard normal variable over the range (α, x).

6. In the case of a lognormal distribution, both magnitudes are related by the following relationship (see Aitchinson and Brown 1966): $G = 2\Pi(\sigma/2^{1/2}) - 1$.

7. It is also known that $\varepsilon_y^{PG} = (PG - H)/PG$, where PG is the poverty gap.

8. Instead of considering poverty measures, it would also be interesting to consider aggregate measures of social welfare. From that point of view, the measure $W = \bar{y} \cdot (1 - G)$, where $\bar{y}$ is the mean income and G is the Gini coefficient that was originally proposed by Sen, lends itself to the same simple decomposition into growth and distribution effects as the poverty indices considered here.

9. Of course, it would have been possible to keep these observations if the original decomposition formula (2) with the lognormal approximation, rather than (2′), had been used in the econometric analysis that follows.

10. Expression (3′) relies on the standard deviation of the logarithm of income whereas inequality is measured by the Gini coefficient in the data set. However, a one-to-one relationship exists between these two magnitudes when the underlying distribution is lognormal (see note 5). That relationship was used to derive σ from the observed value of the Gini.

11. On this see the survey by Aghion, Caroli, and García-Peñalosa. (1999).

References

Aghion, P., E. Caroli, and C. García-Peñalosa. 1999. Inequality and economic growth: The perspective of new growth theories. *Journal of Economic Literature* 37(4): 1615–1660.

Aitchison, J., and J. Brown. 1966. *The Log-Normal Distribution*. Cambridge: Cambridge University Press.

Chen, S., and M. Ravallion. 2000. How did the world's poorest fare in the 1990s? Mimeo., The World Bank.

Datt, G., and M. Ravallion. 1992. Growth and redistribution components of changes in poverty measures: A decomposition with application to Brazil and India in the 1980s. *Journal of Development Economics* 38(2): 275–295.

de Janvry, A., and E. Sadoulet. 1995. Poverty alleviation, income redistribution and growth during adjustment. In *Coping with Austerity: Poverty and Inequality in Latin America*, ed. N. Lustig. Washington, DC: Brookings Institution.

de Janvry, A., and E. Sadoulet. 2000. Growth, poverty, and inequality in Latin America: A causal analysis, 1970–94. *Review of Income and Wealth* 46(3): 267–287.

Deininger, K., and L. Squire. 1996. A new data set measuring income inequality. *The World Bank Economic Review* 10: 565–591.

Dollar, D., and A. Kraay. 2000. Growth is good for the poor. Mimeo., The World Bank.

Fields, G. 2001. *Distribution and Development: A New Look at the Developing World*. Cambridge: The MIT Press.

Foster, J., J. Greer, and E. Thorbecke. 1984. A class of decomposable poverty measures, *Econometrica* 52: 761–766.

Kakwani, N. 1993. Poverty and economic growth with application to Cote d'Ivoire. *Review of Income and Wealth* 39: 121–139.

Ravallion, M. 1997. Can high-inequality developing countries escape absolute poverty? *Economic Letters* 56: 51–57.

Ravallion, M., and S. Chen. 1997. What can new survey data tell us about recent changes in distribution and poverty? *The World Bank Economic Review* 11: 357–382.

Ravallion M., and M. Huppi. 1991. Measuring changes in poverty: A methodological case study of Indonesia during an adjustment period. *The World Bank Economic Review* 5: 57–82.

World Bank. 2000. *Attacking Poverty, World Development Report 2000/2001*. Oxford: Oxford University Press.