

When Is Growth Pro-Poor? Evidence from a Panel of Countries

Aart Kraay
The World Bank

First Draft: December 2003

Revised: December 2004

Abstract: Growth is pro-poor if the poverty measure of interest falls. According to this definition there are three potential sources of pro-poor growth: (a) a high growth rate of average incomes; (b) a high sensitivity of poverty to growth in average incomes; and (c) a poverty-reducing pattern of growth in relative incomes. I empirically decompose changes in poverty in a sample of developing countries during the 1980s and 1990s into these three components. In the medium- to long-run, most of the variation in changes in poverty can be attributed to growth in average incomes, suggesting that policies and institutions that promote broad-based growth should be central to the pro-poor growth agenda. Most of the remainder of the variation in changes in poverty is due to poverty-reducing patterns of growth in relative incomes, rather than differences in the sensitivity of poverty to growth in average incomes. Cross-country evidence provides relatively little guidance as to the policies and institutions that promote these other sources of pro-poor growth.

The World Bank, 1818 H Street N.W., Washington, DC, 20433,
akraay@worldbank.org. This paper has been prepared in the context of the pro-poor growth program sponsored by the World Bank's PREM-Poverty group. I am grateful to Louise Cord, Roberta Gatti, Francisco Ferreira, Tamar Manuelyan-Atinc, Martin Ravallion, and an anonymous referee for helpful comments; and to Shaohua Chen and Diane Steele for providing data. I would also like to thank the research department at the International Monetary Fund for its hospitality while parts of this paper were written. The opinions expressed here are the author's, and do not reflect those of the World Bank or the IMF, their Executive Directors, or the countries they represent.

1. Introduction

The term “pro-poor growth” has recently become pervasive in discussions of development policy. Despite widespread use of the term, there is much less consensus as to what exactly pro-poor growth means, let alone what its determinants are. According to one view, growth is pro-poor if the accompanying change in income distribution by itself reduces poverty (Kakwani and Pernia (2000)). However, this definition is rather restrictive, as it implies that, for example, China’s very rapid growth and dramatic poverty reduction during the 1980s and 1990s was not pro-poor because the poor gained relatively less than the non-poor. A broader and more intuitive definition is that growth is pro-poor if the poverty measure of interest falls. Ravallion and Chen (2003) propose this definition and apply it to a particular poverty measure, the Watts index.

In this paper, I adopt the broader definition, and then apply standard poverty decomposition techniques to identify three potential sources of pro-poor growth: (a) a high growth rate of average incomes; (b) a high sensitivity of poverty to growth in average incomes; and (c) a poverty-reducing pattern of growth in relative incomes. I implement this decomposition for several poverty measures, using household survey data for a large sample of developing countries in the 1980s and the 1990s. I then use variance decompositions to summarize the relative importance of these different sources of pro-poor growth. Finally, I investigate the correlates of the sources of pro-poor growth in this panel of observations on changes in poverty.

The main results of this paper are the following. First, regarding the relative importance of the three potential sources of pro-poor growth, I find that most of the variation in changes in poverty is due to growth in average incomes. In contrast, changes in relative incomes account for only 30 percent of the variance of changes in the headcount measure of poverty in the short run, and only three percent in long run. Growth in average incomes accounts for virtually all of the remaining 70 percent of the variance in the short run, and 97 percent of the variance in the long run, while cross-country differences in the sensitivity of poverty to growth are very small. The share of the variance of changes in poverty due to relative income changes is somewhat larger

for more bottom-sensitive poverty measures, reflecting the fact that changes in these measures place less weight on growth in average incomes.

Second, I find some evidence that growth in average household survey incomes is correlated with several of the usual determinants of growth from the empirical growth literature, including institutional quality, openness to international trade, and size of government. The evidence documented here for the cross-country correlates of growth in household survey incomes is not especially compelling, given various limitations of the dataset. However, I find almost no evidence that poverty-reducing patterns of growth in relative incomes are significantly correlated with a set of explanatory variables that the empirical growth literature has identified as significant determinants of growth in per capita GDP. The same is true for a number of other variables, that while not generally significant for growth, have been suggested in the literature as potentially reducing inequality.

Taken together, these results underscore the importance of growth in average incomes for poverty reduction. This in turn suggests that a policy package focusing on determinants of growth in average incomes, such as the protection of property rights, sound macroeconomic policies, and openness to international trade should be at the heart of pro-poor growth strategies. Moreover, the absence of compelling cross-country evidence that these factors are systematically correlated with the changes in income distribution that matter most for poverty reduction suggests that there are no obvious tradeoffs – policies that lead to growth in average incomes are unlikely to systematically result in adverse effects on poverty through their effects on relative incomes. This does not mean that growth in average incomes is sufficient for poverty reduction. Rather, the results presented here suggest that cross-country evidence is unlikely to be very informative about the policies and institutions that are likely to lead to poverty-reducing patterns of growth in relative incomes. This suggests that more micro-level and case-study research may be useful in shedding light on the determinants of poverty-reducing distributional change.

This paper is related to a growing empirical literature on growth, inequality, and poverty. Most immediately, this paper builds on Dollar and Kraay (2002). In that paper, we defined the poor as those in the bottom quintile of the income distribution, and

empirically investigated the determinants of growth in incomes of the poorest quintile. In a large panel of countries, we found that growth in incomes of the poor tracked growth in average incomes roughly one-for-one. Since the growth rate of average incomes of the poor is the sum of the growth rate of average incomes and the growth rate of the first quintile share, our paper showed that neither average incomes, nor a large set of other control variables, were significantly correlated with changes in the first quintile share. That paper contributed to a growing literature on the determinants of inequality, including Li, Squire and Zhou (1998), Gallup, Radelet and Warner (1998), Spilimbergo, Londono and Szekely (1999), Leamer, Maul, Rodriguez and Schott (1999), Easterly (1999), Barro (2000), Foster and Szekely (2001), and Lundberg and Squire (2003).

This paper differs from Dollar and Kraay (2002), as well as much of the existing literature on determinants of inequality, in two respects. First, instead of looking at relative poverty measures or inequality, here I focus on changes in absolute poverty measures as the dependent variable.¹ As is well understood, changes in absolute poverty measures are complicated non-linear functions of underlying changes in average income and measures of income inequality. The second contribution of this paper is to empirically construct the exact measures of distributional change that matter for changes in various poverty measures for a large sample of countries, rather than simply looking at common summary statistics of inequality such as the Gini coefficient or quintile shares. This means that I can empirically study the contributions of growth and distributional change to changes in poverty, with less restrictive assumptions about the shape of the underlying income distribution.²

Despite these differences, the main conclusions of this paper are similar to those in Dollar and Kraay (2002). In particular, both papers find that growth in average incomes matters a great deal for reductions in both relative and absolute poverty. Both papers also find little evidence that common determinants of growth, as well as a

¹ A notable early exception is Ravallion and Chen (1997), who estimate regressions of changes in absolute poverty on changes in mean incomes using a panel of household surveys from developing countries.

² For example, Lopez (2003) investigates the determinants of growth and change in the Gini coefficient, and then draws conclusions regarding the likely effects on poverty by assuming that the distribution of income is lognormal, so that there is a one-to-one mapping between the Gini coefficient and the Lorenz curve. In this paper I use a more flexible three-parameter approximation to the Lorenz curve, rather than the one-parameter approximation implicit in the lognormal assumption.

number of other variables, are robustly correlated with patterns of distributional change that matter for poverty reduction.

The rest of this paper proceeds as follows. Section 2 reviews standard poverty decomposition techniques and uses them to illustrate the channels through which growth and distributional change matter for changes in a number of poverty measures. Section 3 describes the dataset of changes in poverty in a large sample of developing countries on which the empirical analysis is based. Section 4 provides evidence on the relative importance of the sources of pro-poor growth, as well as evidence on some of the correlates of these sources. Section 5 concludes.

2. Empirical Framework

In this section I use standard techniques to decompose changes in poverty into three components: (a) growth in average incomes; (b) the sensitivity of poverty to growth in average incomes; and (c) changes relative incomes. Let $y_t(p)$ denote the income of the p^{th} percentile of the income distribution at time t . This can be written as a function of average income, μ_t , and the Lorenz curve, $L_t(p)$, i.e. $y_t(p) = \mu_t \cdot \frac{dL_t(p)}{dp}$. Let

P_t denote the following generic additive poverty measure:

$$(1) \quad P_t = \int_0^{H_t} f(y_t(p)) \cdot dp$$

where $H_t = y_t^{-1}(z)$ denotes the fraction of the population below the poverty line, z . This notation captures a number of different poverty measures. For example, if

$f(y_t(p), \theta) = \left(\frac{z - y_t(p)}{z} \right)^\theta$ we have the Foster-Greer-Thorbecke class which includes the

headcount ($\theta=0$), the poverty gap ($\theta=1$), and the squared poverty gap ($\theta=2$). If

$f(y_t(p)) = \ln\left(\frac{z}{y_t(p)}\right)$, we have the Watts poverty index.

Next, differentiate this poverty measure with respect to time to obtain the following expression for the proportionate change in poverty:³

$$(2) \quad \frac{dP_t}{dt} \cdot \frac{1}{P_t} = \int_0^{H_t} \eta_t(p) \cdot g_t(p) \cdot dp$$

Equation (2) expresses the proportional change in poverty as the average across all percentiles of the income distribution of the growth rate of each percentile multiplied by the sensitivity of the poverty measure to growth in that percentile. In particular,

³ Differentiating under the integral sign in Equation (1) requires the application of Leibnitz's rule. Note that the term involving the derivative of the upper limit of integration is zero, since the poverty measures are zero when evaluated at the poverty line.

$\eta_t(p) \equiv \frac{df(y_t(p))}{dy_t(p)} \cdot \frac{y_t(p)}{P_t}$ is the elasticity of the poverty measure with respect to the income of the p^{th} percentile. This term captures the effect on poverty of a small change in incomes of individuals at the p^{th} percentile of the income distribution. This sensitivity is multiplied by the growth rate of each percentile, $g_t(p) \equiv \frac{dy_t(p)}{dt} \cdot \frac{1}{y_t(p)}$, which Ravallion and Chen (2003) refer to as the “growth incidence curve”. The overall proportional change in poverty then consists of the average across all percentiles of the product of these two terms.

In order to separate out the effects of growth in average incomes, re-write Equation (2) as:

$$(3) \quad \frac{dP_t}{dt} \cdot \frac{1}{P_t} = \left(\frac{d\mu_t}{dt} \cdot \frac{1}{\mu_t} \right) \cdot \int_0^{H_t} \eta_t(p) \cdot dp + \int_0^{H_t} \eta_t(p) \cdot \left(g_t(p) - \left(\frac{d\mu_t}{dt} \cdot \frac{1}{\mu_t} \right) \right) \cdot dp$$

Equation (3) identifies the three sources of pro-poor growth discussed above: (a) growth in average incomes; (b) the sensitivity of poverty to growth in average incomes; and (c) growth in relative incomes. The first term in Equation (3) captures the first two sources of pro-poor growth. It consists of growth in average incomes, $\left(\frac{d\mu_t}{dt} \cdot \frac{1}{\mu_t} \right)$, multiplied by a term summarizing the sensitivity of the poverty measure to changes in average incomes, $\int_0^{H_t} \eta_t(p) \cdot dp$. This is simply the average across all percentiles of the sensitivity of poverty to growth in each percentile of the income distribution. The second term in Equation (3) captures the remaining source of pro-poor growth: changes in relative incomes. This third source of pro-poor growth is the average across all percentiles of the income distribution of the product of (a) the growth rate of income in the p^{th} percentile *relative to average income growth*, and (b) the sensitivity of poverty to growth in that percentile. For example, if the poverty measure of interest is very sensitive to growth among the poorest, and if the income of the poorest grows faster than average incomes, then poverty will fall faster.

Equation (3) is useful for thinking about the various definitions and sources of pro-poor growth. The Kakwani and Pernia (2000) definition of pro-poor growth states that growth is pro-poor if and only if the second term in Equation (3) is negative, i.e. the pattern of growth in relative incomes is such that the poverty measure falls. A broader definition of pro-poor growth suggested by Ravallion and Chen (2003) is that growth is pro-poor if the poverty measure of interest falls. According to this definition, there are three potential sources of pro-poor growth: (a) rapid growth in average incomes; (b) a high sensitivity of poverty to growth in average incomes; and (c) a poverty-reducing pattern of growth in relative incomes. In the empirical section of this paper, I will use data on income distributions and average incomes for a large sample of developing countries to construct these three sources of pro-poor growth, document their relative importance, and investigate their determinants. Before doing so, however, it is useful to examine the key ingredients in Equation (3) in more detail: the pattern of growth in relative incomes, $g_t(p) - \left(\frac{d\mu_t}{dt} \cdot \frac{1}{\mu_t} \right)$, and the function summarizing the sensitivity of poverty to growth in each percentile, $\eta_t(p)$.

Figure 1 graphs two examples of the pattern of growth in relative incomes, for China over the period 1990-1998, and for Indonesia over the period 1996-1999. In China, according to the household survey average incomes grew at 14 percent per year, and the dollar-a-day headcount measure of poverty fell from 51 percent to 33 percent of the population. However, there was also a sharp increase in inequality during this period, with the Gini coefficient rising from 34 to 40. The pattern of relative income growth rates shown in the relative growth incidence curve highlights this pattern of increased inequality. Growth in the poorest 80 percentiles of the population was below average, while only the richest 20 percent of the population saw above-average growth. In Indonesia, survey mean income fell dramatically between 1996 and 1999 at nearly 9 percent per year during the East Asian financial crisis. Yet during this period, the pattern of growth in relative incomes was poverty-reducing. Inequality as measured by the Gini coefficient fell from 36.5 to 31.5. The relative growth incidence curve is downward sloping, indicating that incomes of the richer percentiles of the income distribution fell faster than incomes of poorer percentiles. In fact, below-average growth was recorded only for the richest 20 percent of the population. Despite this pro-poor pattern of relative

income growth, the headcount measure of poverty increased from 8 percent to 13 percent of the population, driven by the large negative growth effect.⁴

Consider next the sensitivity of poverty to growth in different percentiles of the income distribution. In the case of the Foster-Greer-Thorbecke class,

$$\eta_t(p) = -\frac{\theta}{P_t} \cdot \frac{y_t(p)}{z} \cdot \left(1 - \frac{y_t(p)}{z}\right)^{\theta-1}, \text{ while for the Watts index, } \eta_t(p) = -\frac{1}{P_t}.$$

Note that these sensitivities in general depend not only on the poverty measure of interest, but also on the entire distribution of income as summarized by $y_t(p)$. Figure 2 graphs these sensitivities, using the actual distribution of income in China in 1990 as an example, to show how different poverty measures are sensitive to growth in different percentiles of the income distribution.

In the case of the headcount, this sensitivity is zero everywhere except just below the poverty line where it spikes down to minus infinity. This is because the headcount simply adds up the number of people below the poverty line – small increases in the incomes of inframarginal poor people that do not bring them above the poverty line will not reduce the headcount. The same is true for increases in incomes of those above the poverty line, including the “near-poor” just above the poverty line. The case of the headcount already illustrates the broader point of Figure 2: the extent to which a given pattern of growth is pro-poor depends crucially on the poverty measure of interest. In particular, if the objective of pro-poor growth is to reduce the headcount measure of poverty, then a pro-poor growth strategy should focus exclusively on raising the incomes of those just at the poverty line, and should ignore everyone else.

This strong -- and well-understood to be absurd -- conclusion is driven by the choice of the headcount as the poverty measure of interest. Consider next the poverty gap and the squared poverty gap. The poverty gap is most sensitive to growth in incomes of those at the poverty line, but is also sensitive to growth in incomes of everyone below the poverty line. The intuition for this is the following: the poverty gap

⁴ It is important to note that the growth incidence curves here adjust only for average inflation. In the case of Indonesia food price inflation during the crisis was much higher than non-food price inflation (see Suryahadi et. al. (2003) for details). To the extent that food represents a larger share of the consumption basket of the poor, the pattern of growth in real incomes was less pro-poor than depicted in Figure 1.

reflects a social welfare function which is indifferent to the distribution of income among poor people. In this case a given rate of average growth results in a larger absolute increase in income for a person near the poverty line, and so the poverty measure is most sensitive to those nearest the poverty line, but is non-zero for all poor people.

The squared poverty gap is also sensitive to growth in the incomes of all those below the poverty line, but the sensitivity is now U-shaped. Growth in incomes of the richest and poorest of those below the poverty line matters least, and the squared poverty gap is most sensitive to growth in incomes of poor people somewhere in between these two extremes. The intuition for this again depends on the underlying social welfare function, which now values absolute transfers from richer to poorer poor people. This however is offset by the fact that a given average growth rate results in a larger absolute increase in income for richer poor people. This is why the sensitivity of the poverty measure to growth is a non-monotonic function of the income percentile.

The Watts index has the property that it is equally sensitive to growth in all percentiles below the poverty line. This is why Ravallion and Chen (2003) argue that a good measure of pro-poor growth is the average (across all percentiles) growth rate of those below the poverty line, i.e. the average growth rate of incomes of the poor. In this paper I go further and decompose the average growth rate of incomes of the poor into growth in average incomes and the average growth rate of the poor *relative* to growth in average incomes. This allows me to distinguish between the effects of growth in average incomes and growth in relative incomes on the Watts measure, and all the other measures considered here. This distinction is not trivial, as we will see in the empirical section of the paper that, across countries, growth in average incomes accounts for a much greater share of the variation in changes in poverty than do changes in relative incomes.

Finally consider the average across all percentiles of the sensitivity of poverty to growth in the incomes of percentile p , $\eta_i(p)$. Recall from Equation (3) that this average sensitivity measures the effect of growth in average incomes on the poverty measure. High values of this average sensitivity of poverty to growth in average incomes are one of the three potential sources of pro-poor growth. For the Foster-Greer-Thorbecke class of poverty measures, this average sensitivity can be expressed in terms of the poverty

measure itself when θ is not equal to zero, $\int_0^{H_t} \eta_t(p) \cdot dp = \theta \cdot \left(1 - \frac{P_t(\theta - 1)}{P_t(\theta)} \right)$, where $P_t(\theta)$

denotes the FGT measure with parameter θ .⁵ In the case where θ is zero, the sensitivity of the headcount to growth in average incomes is:

$$\int_0^{H_t} \eta_t(p) \cdot dp = -\frac{1}{P_t(\theta)} \cdot \frac{L_t'(H)}{\mu \cdot L_t''(H)}$$

which can be expressed as the slope of the density of income at the poverty line. For the Watts measure, the average elasticity is simply minus one times the ratio of the headcount index to the Watts index. While these results are useful for analytically characterizing the sensitivity of the different poverty measures to growth in average incomes, we will see shortly that cross-country differences in the sensitivity of poverty to growth in average incomes are not empirically very important, in the sense that they explain little of the cross-country variation in the first term in Equation (3).

⁵ This result can be found in Kakwani (1993).

3. Data

In the rest of this paper I use the analytic framework of the previous section to decompose observed changes in poverty into the three terms discussed above: (a) growth in average incomes; (b) the sensitivity of poverty to growth in average incomes; and (c) changes in relative incomes. After constructing these three terms for a large sample of developing countries, I use them to identify the relative importance of, and factors correlated with, these sources of pro-poor growth.

I use household survey data on average incomes and ten points on the Lorenz curve for a large number of surveys, as compiled by Martin Ravallion and Shaohua Chen at the World Bank. Their data comes directly from primary sources, and is available at <http://www.worldbank.org/research/povmonitor>.⁶ Depending on the country, the surveys measure either the distribution of income or the distribution of consumption. Average income or consumption is measured in 1993 dollars and is adjusted for cross-country differences in purchasing power parity. Since I am interested in changes in poverty over time, I take only countries with at least two household surveys. This results in a total of 285 surveys covering 80 developing countries. Most of the survey dates are in the 1990s, with some countries extending back to the 1980s. I use the World Bank's "dollar-a-day" poverty line which in 1993 dollars is \$1.08 per day, or \$393 per year.

Using these surveys, I construct two datasets of spells of changes in poverty. In the first dataset, I consider all possible spells for each country, discarding only those few cases where the survey changes from an income to an expenditure survey or vice versa within a country. This results in 185 spells of poverty changes. The length of these spells is quite short, averaging 3.4 years and ranging from one to 13 years. In order to be able to look at changes over longer horizons, I also construct a dataset consisting of one spell per country, where the initial and final years are chosen so as to maximize the length of the spell given available data. This results in a set of 77 spells, with an average length of 8 years, and ranging from two to 19 years. I eliminate all spells where the headcount measure of poverty is negligible in either the initial or final period, i.e. below two percent. I also discard a small proportion of spells for which the average

⁶ I am grateful to Shaohua Chen for kindly providing key data from all of the household surveys, including some that was not at the time available on the poverty monitoring website.

annual growth rate of mean income exceeds 15 percent in absolute value, or for which the average annual growth rate of the headcount exceeds 30 percent in absolute value.⁷ Finally, for the dataset of long spells, I discard those countries for which the longest possible spell is shorter than five years. This reduces the first dataset to 110 spells covering 49 countries with an average length of 3.5 years, and the second dataset to 41 spells with an average length of 9.6 years. These spells are listed in Appendix 1.

In order to construct the different poverty measures and their decompositions discussed in the previous section, I need the full Lorenz curve and not just the 10 points provided in the Ravallion-Chen data. To obtain this, I assume that the Lorenz curve has the following functional form:

$$(4) \quad L(p) = p^\alpha \cdot (1 - (1-p)^\beta)^\gamma, \alpha \geq 0, 0 < \beta \leq 1, \gamma \geq 1$$

This particular parameterization is a member of a family of ordered Lorenz curves proposed by Sarabia, Castillo, and Slottje (1999). In Appendix 2, I discuss in more detail the quality of this parametric approximation to the Lorenz curve, using record-level data from Ghana as an example. The appendix shows that this approximation to the Lorenz curve is quite good, but that the associated quantile function tends to understate incomes of the poorest. This means that poverty measures based on this approximation are likely to be biased upwards, and more so for more bottom-sensitive poverty measures. I also show that these biases will lead to an underestimation of the sensitivity of poverty to growth in average incomes.

I estimate the three parameters of the Lorenz curve for each survey using an algorithm suggested by the same authors. This involves selecting all possible combinations of three points on the Lorenz curve, and then for each combination finding values of α , β , and γ such that the Lorenz curve passes through these three points. The final estimates of α , β , and γ are then found by averaging across all the resulting estimates of these parameters, discarding those for which the parameter restrictions

⁷ This cutoffs roughly correspond to the 5th and 95th percentiles of the distribution of growth rates of mean income and the headcount. In at least some of these cases growth rates of variables are sufficiently extreme as to be implausible, and likely reflect problems of comparability in surveys over time. Appendix 1 lists the discarded observations and summarizes some of the key results including these extreme observations.

indicated in Equation (4) that are required for the Lorenz curve to have positive first and second derivatives do not hold. I then obtain the quantile function by analytically differentiating the Lorenz curve and multiplying by average income. Using this, I can construct $\eta_t(p)$ for each poverty measure of interest, as well as the growth incidence

curve over the observed discrete interval, $g_t(p) = \frac{y_t(p)}{y_{t-1}(p)} - 1$.

4. Results

I begin by constructing four poverty measures of interest (the headcount, the poverty gap, the squared poverty gap, the Watts index) for the initial and final years of each spell. I then compute the average annual growth rates of each of these measures over each spell. Table 1 reports the simple correlations of the levels and average annual growth rates in these poverty measures with the corresponding log-levels and growth rates of survey mean income. These simple correlations are all negative, and are large in absolute value, especially those in levels and those for the long spells. Figure 3 graphs the proportional change in the headcount against the growth rate of average incomes, using the sample of long spells. There is a strong and highly significant negative relationship between changes in poverty and change in average incomes. Table 1 and Figure 3 confirm the widely-understood empirical regularity that poverty on average falls as average incomes increase. In the rest of this section I go beyond this basic observation to document the relative importance of the different sources of pro-poor growth discussed above, and their correlates.

The Relative Importance of Sources of Pro-Poor Growth

I first decompose the change in poverty in each spell into a “growth component” and a “distribution component” using the decomposition suggested by Datt and Ravallion (1992), which is the discrete analog of the infinitesimal decomposition in Equation (3). Let $P(\mu_t, L_t)$ denote a poverty measure based on mean income at time t , μ_t , and the Lorenz curve at time t , L_t . The proportional change in the poverty measure over the discrete interval between time t and $t-1$ is:

$$(5) \quad \frac{P(\mu_t, L_t) - P(\mu_{t-1}, L_{t-1})}{P(\mu_{t-1}, L_{t-1})} = \frac{P(\mu_t, L_{t-1}) - P(\mu_{t-1}, L_{t-1})}{P(\mu_{t-1}, L_{t-1})} + \frac{P(\mu_{t-1}, L_t) - P(\mu_{t-1}, L_{t-1})}{P(\mu_{t-1}, L_{t-1})} + \varepsilon_t$$

The first term on the right-hand side is the growth component of the change in poverty, and is constructed as the proportional difference between the initial poverty measure and a hypothetical poverty measure computed using the second period mean but the first period Lorenz curve. The second term is the distribution component which is computed as the proportional difference between the initial poverty measure and a

hypothetical poverty measure constructed using the first period mean but the second period Lorenz curve. These two components are the discrete-time analogs of the two terms in Equation (3). Unlike Equation (3), however, there is also a residual term because the decomposition is done over a non-infinitesimal interval. Empirically however these residuals will turn out to be unimportant on average. I measure the proportional changes on the left- and right-hand side of Equation (5) as log differences and normalize by the length of the interval to obtain average annual percent changes in poverty and the growth and distribution components for each spell. I also divide the first term in Equation (5) by growth in average incomes to obtain the sensitivity of poverty to growth in average incomes.

Tables 2 and 3 report the results of applying this decomposition to the two datasets of spells. Throughout these two tables, I use the following variance decomposition to summarize the relative importance of the various sources of pro-poor growth. For two correlated random variables X and Y, I define the share of the variance of X+Y due to variation in X as $\frac{\text{VAR}(X)+\text{COV}(X,Y)}{\text{VAR}(X)+\text{VAR}(Y)+2\cdot\text{COV}(X,Y)}$.⁸ The top panel of each table documents the importance of the residual relative to the sum of the growth and distribution components of the change in poverty. The first column shows the variance of the sum of the growth and distribution components, the second column the variance of the residual, and the third the covariance between the two. The final column reports the share of the variance of changes in poverty due to the growth and distribution components, which is virtually one for all poverty measures. This simply reflects the fact that the variance of the residual term is tiny relative to the variance in measured changes in poverty. This can also be verified visually from the top panel of Figure 4, which graphs the change in the headcount measure of poverty on the horizontal axis, and the sum of the growth and distribution components on the vertical axis, using the dataset of long spells. The slope of the OLS regression line is the share of the variance in poverty changes due to the growth and distribution components, and one minus the slope is the share due to the residual term. It is clear from this graph that changes in poverty are

⁸ When X and Y are normally distributed, this variance decomposition has a very natural interpretation. It measures how much the conditional expectation of X increases for each unit that the sum (X+Y) is above its mean value. See Klenow and Rodriguez (1997) for details.

largely accounted for by the sum of the growth and distribution components, with very little of the variation due to the residual.

The middle panels of Tables 2 and 3 report the same variance decomposition, but now to assess the importance of the growth component relative to the distribution component of changes in poverty. For the sample of all spells, between 43 and 70 percent of the variation in changes in poverty is due to the growth component, with the remainder due to changes in relative incomes. For the long spells, between 69 and 97 percent of the variation in changes in poverty is attributable to the growth component, depending on the poverty measure of interest. In both tables, the growth component is relatively less important for bottom-sensitive poverty measures such as the poverty gap and the squared poverty gap. The middle panel in Figure 4 graphically summarizes this second decomposition for the long spells sample, plotting the growth component of changes in the headcount on the vertical axis, and the sum of the growth and distribution components on the horizontal axis. Again, the slope of the OLS regression line can be interpreted as the share of the variation on the horizontal axis due to the growth component. Visually inspecting this graph, it is clear that if the headcount declines substantially, it is mostly because the growth component of poverty reduction is large.

The bottom panels of Tables 2 and 3 further disentangle the growth component into growth in average incomes, and the sensitivity of poverty to growth in average incomes, i.e. they separate the first term in Equation (3) into its two components. Since the variance decomposition used here applies to sums of random variables, I take the logarithm of the absolute value of the growth component, which then becomes the sum of the logarithm of the absolute value of growth, and the logarithm of the absolute value of the average sensitivity of poverty to growth, and apply the decomposition to this sum. Tables 2 and 3 show that around 90 percent of the variation in the growth component of changes in poverty is due to differences in average income growth, and very little is due to differences in the sensitivity of poverty to average income growth. The bottom panel of Figure 4 illustrates this, but without the log transform required to do the variance decomposition. On the horizontal axis I graph the growth component of the change in poverty, while on the vertical axis I graph growth in average incomes. While the slope of this regression cannot be interpreted as a variance share, it nevertheless is very clear that cross-country differences in the growth component of changes in poverty are

overwhelmingly accounted for by cross-country differences in growth itself. Put differently, it is clear from this graph that if the growth component of poverty reduction is large, it is most likely that growth itself was large, rather than that the sensitivity of poverty to growth was large.⁹

Two striking features of Tables 2 and 3 merit further discussion: (a) the share of the variation in poverty measures due to growth declines as the poverty measures become more bottom-sensitive, i.e. when we move from the headcount to the poverty gap to the squared poverty gap; and (b) the share of the variance due to growth is smaller over the short horizons represented in the dataset of all spells, and is larger in the dataset of long spells.

Consider first the observation that the variation in changes in poverty due to growth declines as the poverty measures become more bottom-sensitive. This finding should not be interpreted as evidence that the poorest percentiles of the income distribution are more likely to experience slower-than-average growth. Rather, it primarily reflects the fact that more bottom-sensitive poverty measures place relatively less weight on changes in average incomes than they do on changes in relative incomes.¹⁰ Recall that the sensitivity of poverty to growth in average incomes is

$$\int_0^{H_t} \eta_t(p) \cdot dp = \theta \cdot \left(1 - \frac{P_t(\theta - 1)}{P_t(\theta)} \right) \text{ for the FGT family of poverty measures. Differentiating}$$

this sensitivity with respect to θ and using the fact that $\frac{\partial P_t(\theta)}{\partial \theta} < 0$ and $\frac{\partial^2 P_t(\theta)}{\partial \theta^2} < 0$, it is

straightforward to see that the sensitivity of poverty to growth in average incomes is strictly declining (in absolute value) as the poverty measure becomes more bottom-sensitive, i.e. as θ increases. This means that when relative incomes do not change, the proportional change in poverty associated with a given average growth rate will be

⁹ At first glance this result seems inconsistent with Ravallion (1997), who documents that the sensitivity of poverty to growth varies significantly with initial inequality. However, using either sample of spells I can replicate the result that the interaction of growth with the initial Gini coefficient is significantly correlated with the change in headcount measures of poverty. Although there are cross-country differences in the sensitivity of poverty to growth which are significantly correlated with initial inequality, in the data these differences are dominated by the much larger cross-country differences in growth itself, and this is what the variance decompositions show.

¹⁰ In Appendix 2 I also show that the approximation errors associated with the parameterization of the Lorenz curve reduce the sensitivity of poverty to growth in average incomes for more bottom-sensitive poverty measures.

smaller the more bottom-sensitive is the poverty measure. In other words, even for a purely distribution-neutral growth process, growth will appear to be less pro-poor (in the sense that the proportionate change in poverty is smaller) the more bottom-sensitive is the poverty measure.

When relative incomes also change, it is no longer possible to sign the derivative of the change in poverty with respect to θ for an arbitrary shift in the Lorenz curve. Empirically, however, in the majority of spells in this dataset, proportional changes in poverty are larger in absolute value the more bottom-sensitive are the poverty measures. This is true even though the growth component of changes in poverty is unambiguously smaller in absolute value in all spells, implying that, on average, the distribution component of changes in poverty becomes larger in absolute value the more bottom-sensitive are the poverty measures. This in turn accounts for a smaller share of the variance of changes in poverty due to growth for more bottom-sensitive poverty measures.¹¹

It is also possible to document directly that the incomes of the very poorest on average do not grow more slowly than average incomes, using the estimated growth incidence curves for each spell. In the dataset of long spells, I calculate the average annual growth rate of incomes relative to the survey mean, at the percentiles corresponding to 100 percent, 50 percent, and 25 percent of the initial-period headcount, using only the 22 spells for which the initial headcount is more than 10 percent. The average across spells of these growth rates are -1.2 percent, -1.4 percent, and -1.5 percent respectively, but with very large standard deviations of 3.6 percent, 5.0 percent, and 6.6 percent respectively. Based on this I cannot reject the null hypothesis that the growth rate of incomes at and below the poverty line do not differ significantly from growth in average incomes.

¹¹ For the particular case of equiproportionate shifts in the Lorenz curve, Kakwani (1993) shows that the elasticity of poverty with respect to the Gini coefficient is $\theta \cdot \left(1 + \left(\frac{\mu}{z} - 1 \right) \cdot \frac{P_i(\theta - 1)}{P_i(\theta)} \right)$ It is

straightforward to verify by simple differentiation that this elasticity is strictly increasing in θ when the poverty line is below the mean, i.e. $z < \mu$. Thus for this special case we can unambiguously show that the more bottom-sensitive the poverty measure, the growth component of changes in poverty will be smaller and the distribution component of changes in poverty will be larger.

The second striking feature of Tables 3 and 4 is that the share of the variance of changes in poverty due to growth is larger in the sample of long spells. In order to understand this finding it is useful to examine in more detail the sources of variation in the growth and distribution components of changes in poverty. In the long spells, the standard deviation of the growth component of the headcount is 0.88 percent, which reflects purely cross-country variation in the long-run average growth component of changes in poverty. In the sample of all spells, the standard deviation of the growth component rises to 1.04 percent, reflecting the addition of the within-country variation in the growth component. This fairly modest increase indicates that most of the variation in the growth component of changes in poverty in the sample of all spells is due to cross-country variation, and relatively little reflects within-country variation. Since we have already seen that cross-country differences in the sensitivity of poverty to growth are small, it is also the case in this dataset that average income growth itself varies relatively more across countries than it does within countries over time.

The distribution component of changes in poverty is a weighted average of the growth rates of each percentile of the income distribution relative to average growth. In both the long spells and the full sample of all spells, the relative growth rates of each percentile are on average near zero. However, in the sample of all spells, the relative growth rates of each percentile vary much more across spells than they do in the sample of long spells. This is shown in Figure 5, which reports the average (across spells) of the growth rate of percentile p relative to average growth, as well as the standard deviation (across spells) of this relative growth rate. In both samples the relative growth rates of all percentiles are close to zero on average across spells, consistent with the well-documented fact that growth tends to be distribution-neutral on average. However, the variation around this average is nearly twice as large in the sample of all spells as it is in the sample of long spells, reflecting substantial additional within-country variation over time in relative income growth rates.

Putting together these observations, we can now account for the difference in the relative importance of the growth and distribution components of changes in poverty in the sample of long spells and the sample of all spells. In the long spells sample, relative income changes over long periods are fairly closely clustered around zero for all percentiles of the income distribution. The weighted average of these relative growth

rates, i.e. the distribution component of changes in poverty, is therefore also near zero on average with relatively little dispersion across countries, and contributes relatively little to the cross-country variation in changes in poverty. In the sample of all spells, the variance of the growth component increases only modestly, since growth in average incomes tends to vary much more across countries than it does within countries over time. However, the variance of relative income changes increases much more, reflecting the much greater within-country variation over time in relative income growth rates. As a result, the distribution component of changes in poverty varies more across spells and accounts for a greater share of the overall variation in changes in poverty.

In summary, the results in this subsection tell us that, over longer horizons, between 69 and 97 percent of cross-country differences in poverty changes can be accounted for by growth. The relative importance of the growth component of changes in poverty is smaller for more bottom-sensitive poverty measures. This does not mean that the poorest on average experience slower growth. Rather, it simply reflects the fact that more bottom-sensitive poverty measures assign a smaller weight to changes in average incomes. The relative importance of the growth component is also smaller in the sample of shorter spells than in the sample of long spells. This reflects the fact that, in this dataset, there are non-trivial fluctuations in relative incomes from one survey date to the next, but these tend to average to zero both across countries and over time. Finally, cross-country differences in the sensitivity of poverty to growth in average incomes are empirically relatively unimportant in accounting for cross-country differences in rates of poverty reduction.

What Drives the Sources of Pro-Poor Growth?

I now turn to the question of what drives the various sources of pro-poor growth. In light of the results of the previous section that cross-country differences in the sensitivity of poverty to growth in average incomes are relatively unimportant, I focus primarily on the first and third sources of pro-poor growth: growth in average incomes, and changes in relative incomes. I measure growth in average incomes as the average annual growth rate over the spell of household average income or consumption. I use two different summary measures of changes in relative incomes. The first is simply the average annual proportional change in the Gini index, for comparability with existing

results on the determinants of changes in summary statistics of inequality. I also use the discrete-time distribution component of the change in the headcount measure of poverty. Results for the distribution component of the other three poverty measures are very similar to those for the headcount, and are not reported to save space.

There are many limitations to this dataset which complicate the identification of causal determinants of growth or change in relative incomes. The sample of observations is quite small, especially when we consider the long spells dataset where the determinants of longer-term growth and distributional change are more likely to be apparent. There is also substantial measurement error in the data on growth in survey means, and for measures of distributional change. While classical measurement error in these dependent variables will not necessarily lead to biases in coefficient estimates, it will inflate standard errors and reduce the significance of estimated coefficients. Because there are relatively few spells per country in the dataset consisting of all spells, and only one per country in the long spells dataset, it is not possible to meaningfully base identification on the within-country variation in the data. This raises the possibility that any observed partial correlations may be driven by unobserved country-specific characteristics excluded from the regressions. The small number of spells per country also means that it is not possible to rely on internal instruments to achieve identification.¹²

In light of these difficulties, my more modest objective here is to simply document some partial correlations between these sources of pro-poor growth and a number of right-hand-side variables of interest, and to interpret them with an appropriate abundance of caution. I consider the same list of right-hand-side variables as in Dollar and Kraay (2002). In that paper, we considered a small number of variables that are frequently found to be robustly correlated with real GDP growth in the cross-country growth literature: institutional quality as proxied by a measure of property rights protection (the “rule of law” indicator from Kaufmann, Kraay and Mastruzzi (2003)) as well as the World Bank’s Country Policy and Institutional Assessment (CPIA) indicator;

¹² This is of course especially problematic for the regressions below that involve a lagged dependent variable, which, together with unobserved country-specific effects, will make estimates of the coefficient on the lagged dependent variable inconsistent, and can bias the coefficients on the other variables in different directions depending on their correlation with the lagged dependent variable. See Caselli, Esquivel, and Lefort (1996) for details.

openness to international trade (the constant-price local currency ratio of exports plus imports to GDP); inflation as a proxy for stable monetary policy (measured as the logarithm of one plus the CPI inflation rate); the size of government (measured as the share of government consumption in GDP in local currency units); and a measure of financial development (the ratio of M2 to GDP in local currency units).

We also considered a number of variables that are generally less robustly correlated with growth, but that some studies have found to be correlated with inequality, either in levels or in differences. These include a measure of democracy (the “voice and accountability” indicator from Kaufmann, Kraay and Mastruzzi (2003)); relative productivity in agriculture (measured as the ratio of value added per worker in agriculture relative to overall value added per worker, both in current local currency units); and primary educational attainment. In that paper we found little evidence that any of these variables were robustly correlated with changes in a particular measure of inequality, the first quintile share. The results for the other measures of relative income change considered here will be quite similar.

This list of variables is clearly not an exhaustive list of the potential determinants of growth in average incomes or changes in relative incomes. However, it is a useful starting point in the search for the correlates of growth and distributional change that matter for poverty reduction. I begin by estimating a number of very parsimonious regressions for each of the dependent variables of interest. I regress growth in average incomes on the log-level of initial period income (to pick up convergence effects) plus each of the control variables described above, one at a time. I do the same for the change in the Gini coefficient, instead including the initial level of the Gini coefficient to pick up convergence in this variable. For the distribution component of the change in the headcount, I simply estimate univariate regressions on each of the right-hand-side variables.¹³

Table 4 reports the results, with each entry corresponding to a different regression. The rows correspond to each of the indicated right-hand-side variables.

¹³ Ravallion (2001) documents the empirical importance of inequality convergence using the Gini coefficient. I have experimented with alternative initial inequality measures in the regressions involving the distributional change components of the various poverty measures, but I find that none are robustly significant.

The columns correspond to the different dependent variables and different samples. For the regressions including either initial log income or the initial Gini, I do not report the coefficients on these variables to save space, but they generally enter negatively and usually significantly in all specifications, consistent with available evidence on convergence in both of these variables. Remember also that the distribution component of the change in the headcount is oriented such that a reduction corresponds to a reduction in poverty.

A first glance at Table 4 shows that very few of the explanatory variables of interest are significantly correlated with the dependent variable of interest at conventional significance levels. In fact, in the 54 regressions in these two tables, there are only two coefficients that are significant at the 5 percent level, and only two others at the 10 percent level. One possible explanation for the lack of significant results is that the measures of growth and distributional change on the right-hand-side are contaminated by substantial measurement error. It is difficult to judge however by how much standard errors should be adjusted to reflect this measurement error.¹⁴ Rather than try to judge the statistical significance of the partial correlations documented in Table 4, I simply describe some of the qualitative patterns that emerge.

Consider first *institutional quality*, as proxied by the rule of law indicator. This tends to be positively correlated with growth, but also positively correlated the measures of distributional change, suggesting that distributional change tends to raise poverty in countries with good institutional quality. The *voice and accountability* measure follows the same pattern, although less strongly so, likely because it is quite highly correlated with rule of law in this sample.

Openness to international trade, is positively correlated with growth, but negatively correlated with the Gini coefficient, indicating that distributional change tends to be poverty-increasing in countries that trade more. *Inflation* and *financial development* tend to be extremely weakly correlated with both growth and distributional change in this sample, and with differing signs across specifications. *Government consumption* is

¹⁴ There is an additional factor which likely biases standard errors upward in the sample of all spells. For countries with multiple spells of growth or distributional change, there is likely to be by construction a negative correlation between the errors of successive spells. Correcting for this will likely reduce standard errors somewhat.

negatively correlated with growth, but interestingly is also associated with reductions in inequality in three of the four specifications.

Relative productivity in agriculture is essentially uncorrelated with growth, but tends to be positively correlated with distributional change measures. Somewhat surprisingly the sign of the correlation suggests that countries with higher relative productivity in agriculture are more likely to experience poverty-increasing changes in relative incomes. Finally, *primary education* is also virtually uncorrelated with growth, and also is essentially uncorrelated with most of the distributional change measures, with the exception of the Gini in the long spells regression. Somewhat surprisingly, the correlation is positive, suggesting increases in inequality are more likely in countries with higher education.

Overall, while most of the partial correlations documented in Table 4 are not statistically significant, the qualitative pattern suggests that there may be some tradeoffs. For example, rule of law and trade are positively correlated with growth but also with poverty-increasing shifts in relative incomes, while the opposite is true for government consumption. Table 5 explores these possible tradeoffs in a slightly richer empirical specification, using the dataset of long spells. The first column reports a more fully-specified growth regression with initial income, and initial values of institutional quality, trade openness, and size of government as right-hand-side variables. Despite the likely noise in the data, it is encouraging that it is possible to find a plausible specification in which some of the determinants of growth from the growth literature are also reasonably significantly correlated with growth in the household survey mean. These variables enter with signs consistent with the broader growth literature. Initial income enters negatively, picking up convergence effects, although not significantly. Institutional quality and trade are both significantly positively correlated with growth, and larger government size is significantly associated with slower growth. I do not want to claim that these results are a robust feature of this particular dataset. However, the results are broadly consistent with the findings of the empirical growth literature, which uses per capita GDP growth rates for a much larger sample of countries, and so it seems reasonable to consider this particular specification.

In the second column of Table 5, I show the same regression, but instead using the change in the Gini coefficient as the dependent variable. None of the correlates of growth are significantly correlated with changes in this summary statistic of inequality. It is however difficult to move from the results in these first two columns to conclusions about the effects on poverty, without making particular assumptions on the shape of income distributions. Since I have already constructed the growth and distribution components of changes in poverty, I can simply use these as dependent variables to investigate how these correlates of growth matter for changes in poverty. The remaining two columns of Table 5 do this for the headcount index. Given the high correlation between the growth component of poverty changes and average income growth documented above, it is not surprising that the regressions for the growth components of poverty are very similar to the growth regression in the first column. The signs, however, are switched, because the sensitivity of poverty to growth is negative. Institutional quality, trade, and government size are all correlated with the growth component of the change in the headcount in the expected direction, although the significance is slightly less than before. In contrast, only trade appears to be significantly associated with increases in poverty through changes in relative incomes, at the 10 percent level.

The last two regressions permit quantifying the potentially offsetting impacts of these variables on poverty through the growth and distribution channels. Since the observed change in poverty is essentially equal to the sum of the growth and distribution components (with a relatively unimportant residual as we have seen), the overall effect on poverty of each of these variables is just the sum of the two coefficients. The estimates in Table 5 suggest that the magnitude of the growth effects of these three variables is substantially larger (in absolute value) than the distribution effects. For Rule of Law, the growth effect lowers poverty, while the distribution effect raises it, but the growth effect is an order of magnitude larger than the distribution effect. For trade, the growth effect is poverty-reducing while the distribution effect is poverty-increasing, but again the former dominates the latter. Finally, the growth and distribution effects work in opposite directions for government size, but with the adverse growth effect more than twice as large as the poverty-reducing distribution effect.

Conclusions

In this paper I have used standard decomposition techniques to identify three potential sources of pro-poor growth: (a) a high rate of growth of average incomes; (b) a high sensitivity of poverty to growth in average incomes; and (c) a poverty-reducing pattern of growth in relative incomes. Empirically implementing these decompositions for a large sample of changes in poverty, we have seen that only the first and third sources of pro-poor growth are empirically relevant. Moreover, especially in the medium- to long run, cross-country differences in growth in average incomes are the dominant factor explaining changes in poverty. Together, these decomposition results indicate that the search for pro-poor growth should begin by focusing on determinants of growth in average incomes. At some level, this is an encouraging conclusion, because by now a large body of empirical results exists on the policies and institutions that drive growth in average incomes.

Nevertheless, the empirical results shown here on the correlates of growth and distributional change are rather unsatisfactory. Most of the simple correlations between these dependent variables and a number of right-hand-side variables of interest are far from significant at conventional levels. It is possible to find multivariate specifications for growth in survey means over longer horizons that yield sensible results consistent with the empirical growth literature. At most, this provides some comfort that the results on partial correlates of growth in survey mean income documented here are more broadly robust and may even have causal interpretations. However, there is much more to be learned about why per capita GDP growth, whose determinants are well-documented, translates so imperfectly into growth in survey means.¹⁵

In contrast, in this sample it is difficult to find significant correlates of either changes in summary statistics of inequality such as the Gini, or distributional shifts that matter for a variety of poverty measures of interest such as the ones I have constructed here. Moreover, some of the partial correlations with distributional change documented here do not appear to be consistent with those uncovered in other papers. The wide range of signs and significance of results from the cross-country literature should caution us against drawing particularly strong conclusions about the determinants of pro-poor

¹⁵ See for example Deaton (2003) for a discussion of some of the relevant issues.

changes in relative incomes from any one cross-country study. Country-specific research using household level data is likely to shed more light on the forces driving relative income changes that matter for poverty reduction.

References

- Barro, Robert J. (2000). "Inequality and Growth in a Panel of Countries". *Journal of Economic Growth*. 5:5-32.
- Caselli, Francesco, Gerardo Esquivel and Fernando Lefort (1996). "Reopening the Convergence Debate: A New Look at Cross-Country Growth Empirics". *Journal of Economic Growth* 1: 363-389.
- Chen, Shaohua and Martin Ravallion (1997). "What Can New Survey Data Tell Us about Recent Changes in Distribution and Poverty?" *The World Bank Economic Review*, 11(2):357-382.
- Deaton, Angus (2003). "Measuring Poverty in a Growing World, or Measuring Growth in a Poor World". *NBER Working Paper No. 9822*.
- Dollar, David and Aart Kraay (2002). "Growth is Good for the Poor". *Journal of Economic Growth*. 7:195-225.
- Easterly, William (1999). "Life During Growth." *Journal of Economic Growth*, 4:239-276.
- Foster, James and Miguel Székely (2001). "Is Economic Growth Good for the Poor? Tracking Low Incomes Using General Means". Interamerican Development Bank Research Department Working Paper No. 453.
- Gallup, John Luke, Steven Radelet and Andrew Warner (1998). "Economic Growth and the Income of the Poor". Manuscript, Harvard Institute for International Development.
- Kakwani, Nanak (1993). "Poverty and Economic Growth, With Application to Cote d'Ivoire". *Review of Income and Wealth*. 39(2):121-139.
- Kakwani, Nanak (2000). "What Is Pro-Poor Growth?". *Asian Development Review*. 18(1): 1-16.
- Klenow, Peter J. and Andrés Rodríguez-Clare (1997), "The Neoclassical Revival in Growth Economics; Has It Gone Too far?" in Ben S. Bernanke and Julio Rotemberg eds. *NBER Macroeconomics Annual 1997*. Cambridge, MIT Press.
- Leamer, Edward, Hugo Maul, Sergio Rodriguez, and Peter Schott (1999). "Does Natural Resource Abundance Increase Latin American Income Inequality?". *Journal of Development Economics*. 59:3-42.
- Li, Hongyi, Lyn Squire and Heng-fu Zou (1998). "Explaining International and Intertemporal Variations in Income Inequality." *The Economic Journal*, 108:26-43.
- Lopez, Humberto (2003). "Pro Growth, Pro Poor: Is There a Tradeoff". Manuscript, The World Bank.

- Lundberg, Mattias and Lyn Squire (2000). "The Simultaneous Evolution of Growth and Inequality." *Economic Journal*. 113:326-344.
- Milanovic, Branko (2003). "Can We Discern the Effect of Globalization on Income Distribution? Evidence from Household Surveys". Manuscript, The World Bank.
- Ravallion, Martin (1997). "Can High-Inequality Developing Countries Escape Absolute Poverty?". *Economics Letters*. 56(1):51-57.
- Ravallion, Martin (2001). "Inequality Convergence". World Bank Policy Research Department Working Paper No. 2645.
- Ravallion, Martin and Shaohua Chen (2003). "Measuring Pro-Poor Growth". *Economics Letters*. 78:93-99.
- Sarabia, J-M, Enrique Castillo, and Daniel Slottje (1999). "An Ordered Family of Lorenz Curves". *Journal of Econometrics*. 91:43-60.
- Spilimbergo, Antonio, Juan Luis Londono, and Miguel Szekely (1999). "Income Distribution, Factor Endowments, and Trade Openness". *Journal of Development Economics*. 59:77-101.
- Suryahadi, Asep, Sudarno Sumarto, and Lant Pritchett (2003). "The Evolution of Poverty During the Crisis in Indonesia". SMERU Research Institute Working Paper.

Table 1: Correlations of Poverty Measures and Survey Mean Income or Consumption

	<i>Correlations with Survey Mean</i>	
	Levels	Growth Rates
<i>All Spells (110 Observations)</i>		
Headcount	-0.859	-0.673
Poverty Gap	-0.751	-0.595
Squared Poverty Gap	-0.640	-0.535
Watts	-0.672	-0.559
<i>Long Spells (41 Observations)</i>		
Headcount	-0.835	-0.717
Poverty Gap	-0.721	-0.688
Squared Poverty Gap	-0.628	-0.649
Watts	-0.651	-0.657

Table 2: Decomposing Changes in Poverty: All Spells

Growth, Distribution and Residual Components of Change in Poverty: $dP = G + D + R$

Growth and Distribution Components vs Residual (G+D vs R)

	<u>V(G+D)</u>	<u>V(R)</u>	<u>COV(G+D,R)</u>	<u>Share of Variance Due to G+D</u>
Headcount	0.0110	0.0004	-0.0001	0.9727
Poverty Gap	0.0206	0.0006	-0.0004	0.9937
Squared Poverty Gap	0.0320	0.0008	-0.0009	1.0017
Watts	0.0263	0.0006	-0.0006	0.9998

Growth vs Distribution Components (G vs D)

	<u>V(G)</u>	<u>V(D)</u>	<u>COV(G,D)</u>	<u>Share of Variance Due to G</u>
Headcount	0.0109	0.0065	-0.0032	0.7011
Poverty Gap	0.0143	0.0130	-0.0033	0.5304
Squared Poverty Gap	0.0169	0.0214	-0.0031	0.4292
Watts	0.0154	0.0175	-0.0033	0.4602

Average Growth and Sensitivity to Average Growth in Growth Component: $\ln|G| = \ln|d\ln\mu| + \ln|\eta|$

	<u>V(dlnμ)</u>	<u>V(η)</u>	<u>COV(dlnμ , η)</u>	<u>Share of Variance Due to dlnμ </u>
Headcount	1.0221	0.1494	-0.0473	0.9052
Poverty Gap	1.0221	0.1236	-0.0381	0.9201
Squared Poverty Gap	1.0221	0.1304	-0.0403	0.9159
Watts	1.0221	0.1175	-0.0366	0.9241

Table 3: Decomposing Changes in Poverty: Long Spells

Growth, Distribution and Residual Components of Change in Poverty: $dP = G + D + R$

Growth and Distribution Components vs Residual (G+D vs R)

	<u>V(G+D)</u>	<u>V(R)</u>	<u>COV(G+D,R)</u>	<u>Share of Variance Due to G+D</u>
Headcount	0.0061	0.0005	-0.0003	0.9636
Poverty Gap	0.0116	0.0008	-0.0010	1.0154
Squared Poverty Gap	0.0176	0.0012	-0.0018	1.0396
Watts	0.0142	0.0009	-0.0011	1.0166

Growth vs Distribution Components (G vs D)

	<u>V(G)</u>	<u>V(D)</u>	<u>COV(G,D)</u>	<u>Share of Variance Due to G</u>
Headcount	0.0078	0.0021	-0.0019	0.9738
Poverty Gap	0.0108	0.0041	-0.0017	0.7885
Squared Poverty Gap	0.0132	0.0066	-0.0011	0.6863
Watts	0.0118	0.0056	-0.0016	0.7182

Average Growth and Sensitivity to Average Growth in Growth Component: $\ln|G| = \ln|d\ln\mu| + \ln|\eta|$

	<u>V(dlnμ)</u>	<u>V(η)</u>	<u>COV(dlnμ , η)</u>	<u>Share of Variance Due to dlnμ </u>
Headcount	1.2199	0.1806	-0.0306	0.8880
Poverty Gap	1.2199	0.1591	-0.0516	0.9157
Squared Poverty Gap	1.2199	0.1633	-0.0588	0.9174
Watts	1.2199	0.1511	-0.0549	0.9237

Table 4: Correlates of Pro-Poor Growth

<i>RHS Variable is:</i>	All Spells <i>Dependent Variable Is:</i>				Long Spells <i>Dependent Variable Is:</i>			
	<u>Growth</u>	<u>Percent Change in Gini</u>	<u>Distribution Component of Headcount</u>	<u># Obs</u>	<u>Growth</u>	<u>Percent Change in Gini</u>	<u>Distribution Component of Headcount</u>	<u># Obs</u>
CPIA	0.008 (1.23)	0.000 (0.07)	0.007 (0.68)	107	0.005 (0.48)	0.001 (0.21)	0.000 (0.04)	39
KK Rule of Law	0.016 (1.76)*	0.000 (0.06)	0.011 (0.83)	110	0.024 (1.73)*	0.007 (1.06)	0.025 (2.00)**	41
Trade/GDP	0.018 (1.21)	-0.018 (1.55)	0.004 (0.20)	110	0.037 (1.68)	-0.002 (0.18)	0.006 (0.28)	41
ln(1+Inflation)	0.002 (0.09)	-0.002 (0.12)	-0.012 (0.38)	103	-0.001 (0.06)	0.001 (0.07)	-0.014 (0.61)	38
Government Consumption/GDP	-0.140 (1.04)	0.002 (0.02)	-0.087 (0.46)	109	-0.108 (0.72)	-0.023 (0.28)	-0.161 (1.16)	41
M2/GDP	0.014 (0.64)	-0.007 (0.40)	-0.007 (0.23)	110	-0.028 (0.63)	-0.010 (0.44)	-0.030 (0.71)	41
KK Voice and Accountability	0.008 (0.88)	0.005 (0.75)	0.007 (0.55)	110	0.022 (1.68)	0.010 (1.51)	0.019 (1.56)	41
Relative Productivity in Agriculture	0.007 (0.33)	0.022 (1.24)	0.040 (1.21)	109	0.013 (0.53)	0.021 (1.38)	0.042 (1.52)	40
Average Years of Primary Education	0.009 (1.42)	0.004 (1.01)	0.001 (0.10)	92	0.000 (0.02)	0.009 (3.08)***	0.007 (1.00)	30

Note: t-statistics reported below coefficient estimates. * (**) (***) denote significance at the 1 (5) (10) percent levels.

Table 5: Multivariate Growth and Distributional Change Regressions

	Annual Percent Change in:		Annual Percent Change in Headcount	
	<u>Survey Mean</u>	<u>Gini</u>	<u>Growth Component</u>	<u>Distribution Component</u>
Initial Income	-0.015 (1.36)	-0.007 (1.29)	0.035 (1.66)	-0.013 (1.17)
KK Rule of Law	0.047 (1.98)*	-0.008 (0.66)	-0.058 (1.55)	0.005 (0.31)
Trade/GDP	0.023 (2.08)**	0.010 (0.95)	-0.048 (2.23)**	0.030 (1.89)*
Government Consumption	-0.279 (2.23)**	-0.045 (0.40)	0.588 (2.10)**	-0.223 (1.06)
R-Squared	0.24	0.08	0.26	0.19
# Observations	41	41	41	41

Note: Heteroskedasticity-consistent t-statistics reported below coefficient estimates.

Figure 1: Relative Growth Incidence Curves

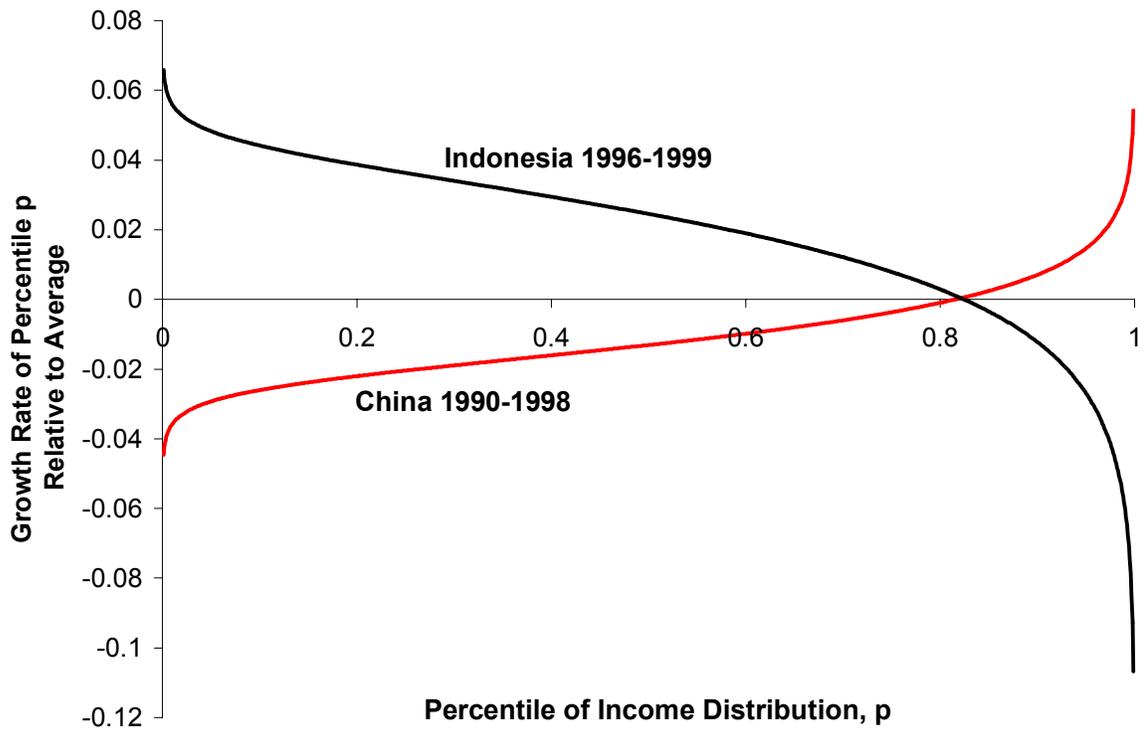
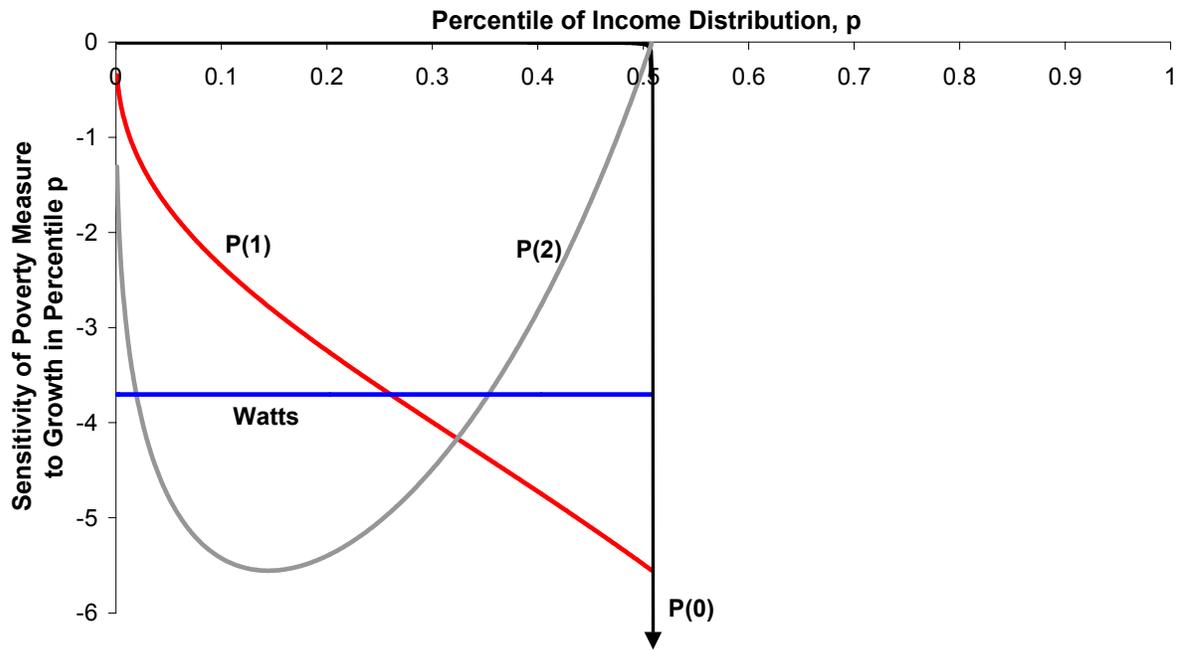


Figure 2: Sensitivity of Poverty to Growth in Percentile p



**Figure 3: Growth and Poverty Reduction
(Long Spells, Headcount)**

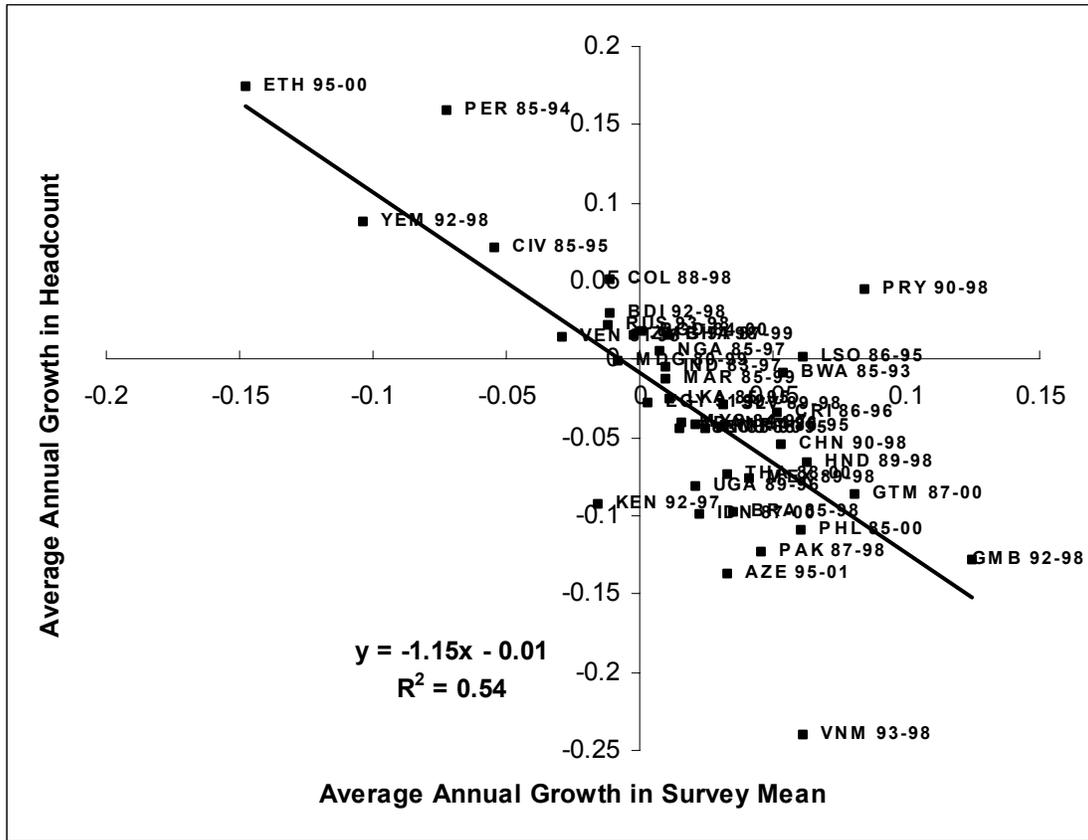


Figure 4: Decomposing Changes in the Headcount (Long Spells)

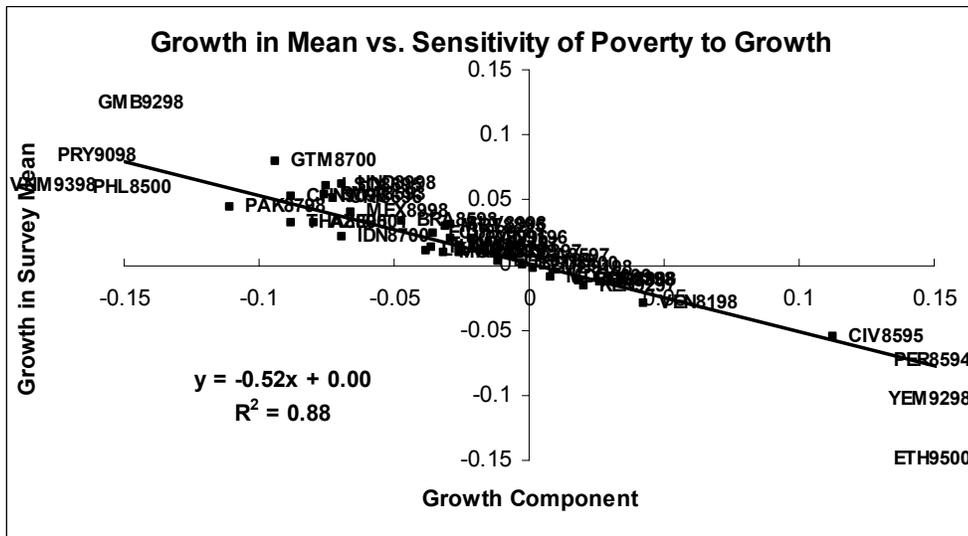
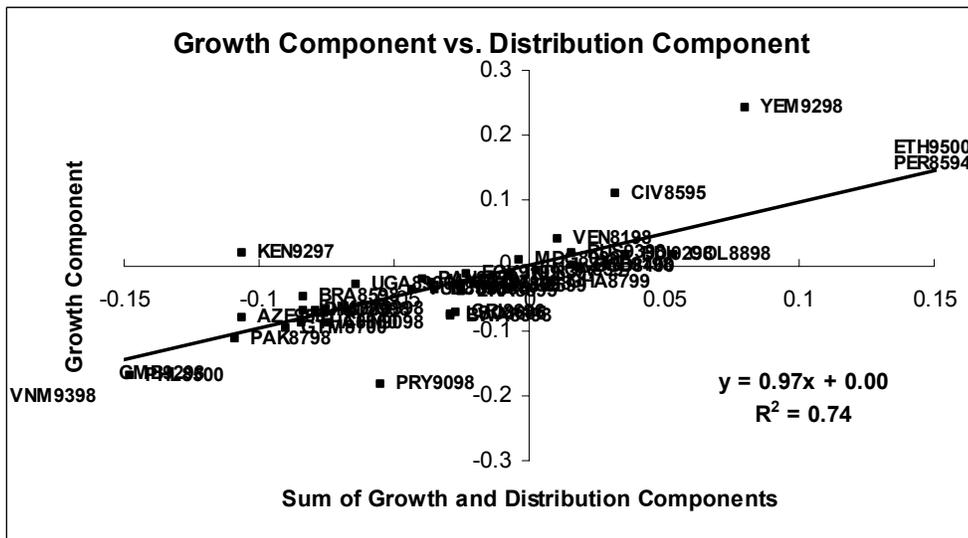
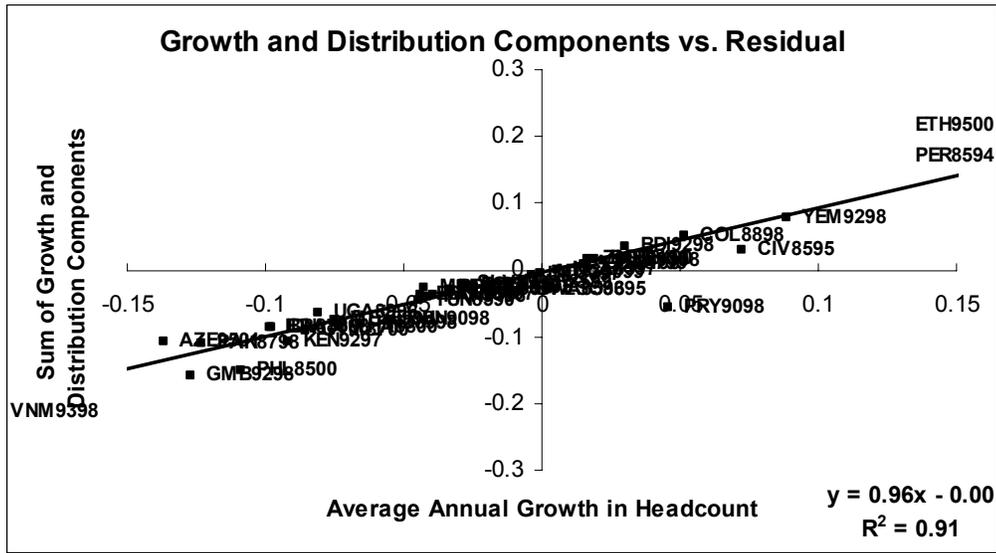
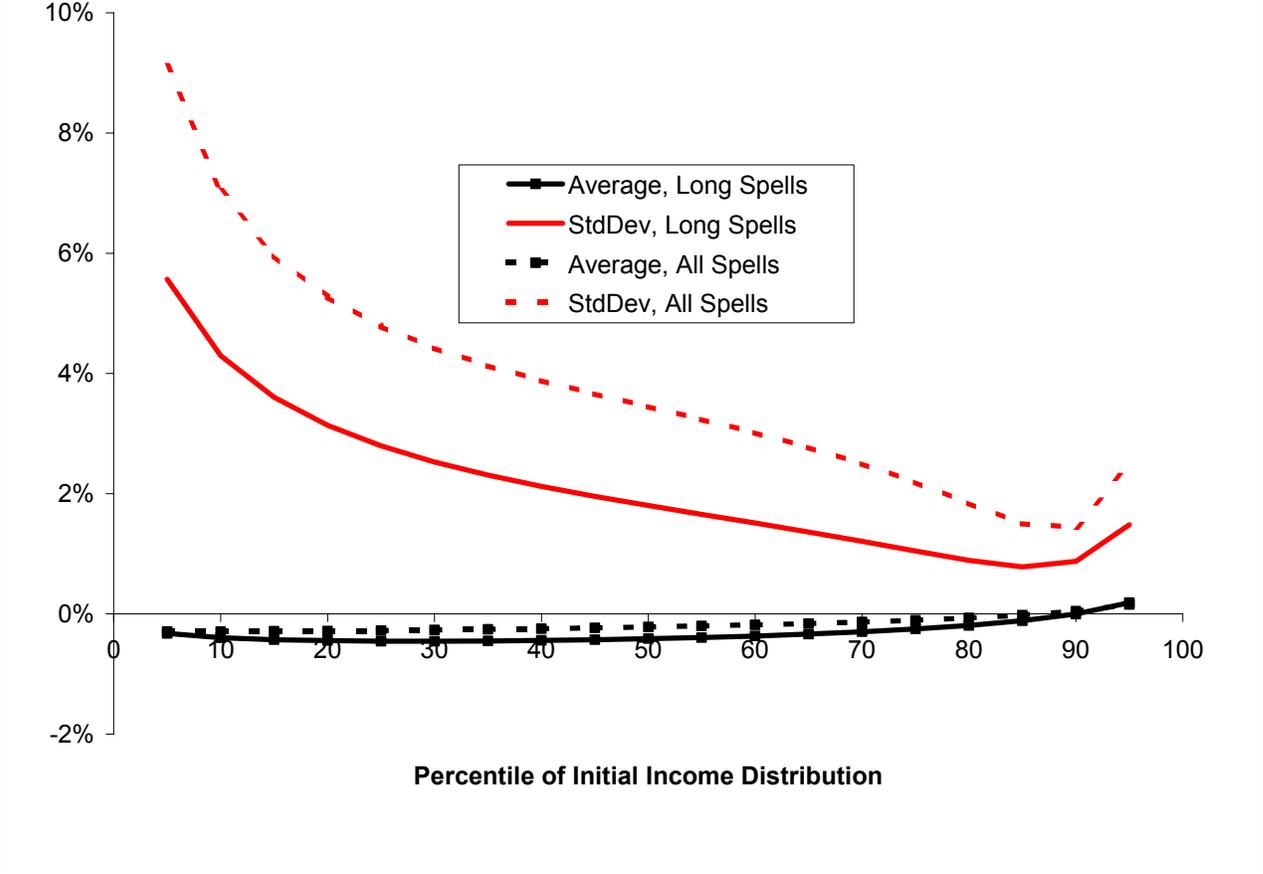


Figure 5: Relative Income Growth Rates, Long and Short Spells
(Average and Standard Deviation Across Spells for Indicated Percentiles)



Appendix 1: Sample of Spells

Table A1.1 lists the sample of spells used in the variance decompositions in Tables 2 and 3. Table A1.2 lists the observations excluded from the sample due to extreme values of growth in the survey mean or growth in the headcount measure of poverty. My assumption is that these extreme observations have a higher noise-to-signal ratio than the rest of the dataset. However, it should be noted that some of the excluded observations in Table A1.2 likely reflect actual changes in living standards, and the convenient shortcut of simply eliminating extreme observations will end up discarding both legitimate as well as problematic data points.

Although the main conclusions of this paper are not sensitive to the exclusion of these datapoints, they do matter for the precise magnitudes of the results. To illustrate this Table A1.3 reports the variance decompositions for the headcount measure, for the unrestricted sample and for the sample used in the main body of the paper dropping extreme observations. In the sample of all spells, the main consequence of dropping outliers is to slightly lower the share of the variance of changes in poverty due to the growth component, from 81 percent to 70 percent. In the sample of long spells, dropping just two outliers (for Zimbabwe and Mali) raises the share of the variance of changes in poverty due to growth from 86 percent to 97 percent. In this sample the share of the variance in the growth component due to growth in the survey mean also falls slightly from 93 percent to 89 percent.

Table A1.1 – Sample of Spells

	<u>All</u>	<u>Long</u>		<u>All</u>	<u>Long</u>		<u>All</u>	<u>Long</u>
Azerbaijan	1995-2001	1995-2001	Ghana	1987-1989	1987-1999	Niger	1992-1995	
Burundi	1992-1998	1992-1998	Gambia	1992-1998	1992-1998	Nigeria	1985-1992	1985-1997
Burkina Faso	1994-1998		Guatemala	1987-1989	1987-2000		1992-1997	
Bangladesh	1984-1985	1984-2000		1989-2000		Pakistan	1987-1990	1987-1998
	1985-1988		Honduras	1990-1992	1989-1998		1990-1996	
	1988-1992			1992-1994		Panama	1991-1995	1991-1996
	1992-2000			1994-1996			1995-1996	
Brazil	1985-1988	1985-1998	Indonesia	1987-1993	1987-2000	Peru	1985-1994	1985-1994
	1988-1989			1993-1996		Philippines	1985-1988	1985-2000
	1989-1993			1996-1999			1988-1991	
	1993-1995		India	1983-1986	1983-1997		1991-1994	
	1995-1996			1986-1987			1994-1997	
	1997-1998			1987-1989		Paraguay	1990-1995	1990-1998
Botswana	1985-1993	1985-1993		1989-1990			1995-1998	
Chile	1987-1990			1990-1992		Russian Fed.	1993-1996	1993-1998
	1990-1992			1992-1994			1996-1998	
China	1990-1992	1990-1998		1994-1995		Senegal	1991-1994	
	1992-1993			1995-1996		El Salvador	1989-1995	1989-1998
	1993-1994			1996-1997			1995-1996	
	1994-1995		Jamaica	1988-1989			1996-1998	
	1995-1996			1990-1993		Thailand	1988-1992	1988-2000
	1996-1997			1993-1996			1992-1996	
	1997-1998		Kenya	1992-1994	1992-1997		1996-1998	
Cote d'Ivoire	1986-1987	1985-1995		1994-1997			1998-2000	
	1988-1993		Sri Lanka	1985-1990	1985-1995	Trinidad & Tob.	1988-1992	
	1993-1995			1990-1995		Tunisia	1985-1990	1985-1990
Colombia	1988-1991	1988-1998	Lesotho	1986-1993	1986-1995	Turkey	1987-1994	
	1991-1995		Madagascar	1980-1993	1980-1999	Uganda	1989-1992	1989-1996
	1995-1996			1993-1999			1992-1996	
Costa Rica	1986-1990	1986-1996	Mexico	1989-1995	1989-1998	Venezuela	1981-1987	1981-1998
	1990-1993		Mauritania	1988-1993	1988-1995		1987-1989	
	1993-1996			1993-1995			1989-1993	
Dominican Rep.	1989-1996		Malaysia	1984-1987	1984-1997		1996-1998	
Ecuador	1988-1994	1988-1995		1987-1989		Vietnam	1993-1998	1993-1998
Egypt	1991-1999	1991-1999		1989-1992		Yemen	1992-1998	1992-1998
Estonia	1993-1995			1992-1995		Zambia	1993-1998	1991-1998
Ethiopia	1995-2000	1995-2000		1995-1997				

Table A1.2 – Discarded Spells

		<i>Average Annual Growth in:</i>	
		<u>Survey Mean</u>	<u>Headcount</u>
All Spells			
Armenia	1996-1998	-0.170	0.150
Brazil	1996-1997	0.238	-0.965
Chile	1992-1994	0.161	-0.235
Cote d'Ivoire	1985-1986	-0.094	-0.432
Cote d'Ivoire	1987-1988	-0.229	0.563
Colombia	1996-1998	0.163	-0.152
Ecuador	1994-1995	0.164	-0.384
Ghana	1989-1992	0.141	-0.562
Ghana	1992-1993	-0.198	0.693
Ghana	1993-1997	-0.340	0.691
Ghana	1997-1999	0.611	-0.748
Honduras	1989-1990	-0.169	0.092
Honduras	1996-1998	0.311	-0.242
Indonesia	1999-2000	0.115	-0.607
Jamaica	1989-1990	0.018	-0.505
Kyrgyz Republic	1993-1997	0.078	-0.508
Lesotho	1993-1995	0.395	-0.165
Lithuania	1993-1994	0.643	-1.552
Mexico	1995-1998	0.160	-0.252
Mali	1989-1994	-0.172	0.313
Mongolia	1995-1998	-0.134	0.394
Pakistan	1996-1998	0.149	-0.440
Philippines	1997-2000	0.173	-0.390
Tanzania	1991-1993	0.051	-0.478
Venezuela	1993-1995	-0.070	0.455
Venezuela	1995-1996	-0.152	0.448
South Africa	1993-1995	0.301	-0.007
Zambia	1991-1993	-0.154	0.095
Zimbabwe	1990-1995	-0.416	0.210
Long Spells			
Mali	1989-1994	-0.172	0.314
Zimbabwe	1990-1995	-0.416	0.210

Table A1.3 – Consequences of Discarding Spells for Variance Decompositions

	<u>Share of Variance of $d\ln P$ due to G+D</u>	<u>Share of Variance of G+D Due to G</u>	<u>Share of Variance of $\ln G$ due to $d\ln m$</u>	<u>Number of Observations</u>
All Spells				
Full Sample	1.20	0.81	0.90	139
Restricted Sample	0.97	0.70	0.91	110
Long Spells				
Full Sample	1.05	0.86	0.93	43
Restricted Sample	0.96	0.97	0.89	41

Appendix 2: Quality of Parametric Approximation to Lorenz Curve

This appendix discusses the quality of the parametric approximation to the Lorenz curve used in the paper, and the consequences for the results, using record-level household survey data from Ghana as an illustration. It is important to note at the outset that using parametric Lorenz curves fitted to grouped data is unavoidable in a large cross-country exercise such as this one. This is because it is difficult if not impossible to obtain access to record-level household survey data for a large set of countries.¹⁶ Nevertheless, it is useful to investigate the extent to which errors introduced by low-dimensional parametrizations of the income distribution may influence the results presented here.

There is a large literature on estimating Lorenz curves, and checking the quality of various parametric approximations to empirical Lorenz curves. The Sarabia et. al. (1999) paper is typical of this literature in that it finds that one- to three-parameter Lorenz curves fit actual distributions quite well on average, although the quality of the approximation is generally worse at the upper and lower ends of the income distribution. There is less systematic evidence on the consequences of these approximations for poverty measurement. Since poverty measures depend on the quantile function, the key issue is the extent to which first derivative of the Lorenz curve is well approximated, rather than the Lorenz curve itself. Ravallion et. al. (1991) use record-level data from Indonesia in 1984 to document that poverty estimates based on three-parameter Lorenz curves fitted to grouped data provide fairly good approximations to the headcount and poverty gap for that country and year, and that the quality of the approximation varies little with the number of groups.

Here I use data from the Ghana 1998/99 Living Standards Measurement Survey to perform a similar exercise. I extract from this survey the household-level consumption aggregate, household size, and the sampling weights, for the 5998 households covered by the survey. To capture intra-household scale economies I construct per capita consumption as household consumption divided by the square root of household size. I then apply the sampling weights to each per capita consumption observation to arrive at

¹⁶ The largest cross-country dataset on poverty and inequality is the Global Poverty Monitoring dataset maintained by the World Bank. The compilers of this dataset have obtained access to record-level data for only about half of the surveys covered.

the distribution of consumption across individuals. To assess the quality of parametric approximations to the Lorenz curve based on a small number of observations on grouped data, I extract just the decile shares from the true Lorenz curve, and then fit the parametric Lorenz curve to these nine data points as described in the text.

The top-left panel of Figure A2 reports the actual and estimated Lorenz curves obtained in this way, and the top-right panel plots the difference between the actual and estimated Lorenz curves. These figures are consistent with existing evidence: the parametric Lorenz curve traces the actual one quite well. The discrepancy between the estimated and actual Lorenz curve is less than one-half of one percent over most of the range, with a maximum error of one percent near the upper end of the distribution. The mean absolute error is 0.3 percent, which is comparable to the errors reported in Sarabia et. al. (1999) for Sweden, Brazil, and the United States.

The bottom panels of Figure A2 show report the logarithm of the actual and estimated quantile functions, and the difference between the two. The quantile function based on the parametric Lorenz curve tracks the actual quantile function quite well around the middle of the income distribution, but there are substantial discrepancies at the upper and lower tails of the distribution. The gap between estimated and actual consumption is less than two percent for the middle three quintiles of the distribution, but the parametric quantile function substantially underestimates (overestimates) consumption in the bottom (top) quintile. It is not hard to see this pattern is not specific to the case of Ghana: the parametric quantile function goes to zero (infinity) as the population share goes to zero (one), while actual consumption is never zero or infinite. As a result, this parametric approximation will always tend to understate consumption at the bottom end of the distribution, and overstate it at the upper end.

This pattern of approximation errors in the quantile function has straightforward implications for poverty estimates. The top two panels of Table A2 report the four poverty indices considered in this paper, for two poverty lines, based on the actual data and on the estimated quantile function. In the top panel, the poverty line is set at the mean, resulting in a poverty headcount 59.7 percent. For this high poverty line, all four poverty measures are quite well approximated using the estimated quantile function: the difference between the actual and estimated poverty indices varies between 0.4 percent

and 4.1 percent. However, when the poverty line is set at half the mean, in the middle panel of Table A2, the approximation errors are significantly larger, with the estimated poverty measures consistently substantially larger than the actual ones. This is not surprising, since with a lower poverty line a greater proportion of poor individuals have their consumption understated by the parametric quantile function. For the same reason, in both panels the poverty estimates based on the parametric quantile function overstate poverty more the more bottom-sensitive is the poverty index.

A further interesting observation is that the poverty measures based on estimated Lorenz curves underestimate the sensitivity of poverty to growth. This can be seen in the bottom panel, which reports the log change between the first and second panels of the estimated and actual poverty measures. Since shifting the poverty line down is equivalent to shifting the mean of the distribution up, this can be interpreted as the proportionate change in poverty that would occur if incomes doubled, holding constant the Lorenz curve. The third column shows that the estimated poverty measures fall less than the actual ones (except for the headcount), and that this discrepancy is larger for increasingly bottom-sensitive measures. To see why, recall from the main text of the paper that the sensitivity of poverty to growth in average incomes is

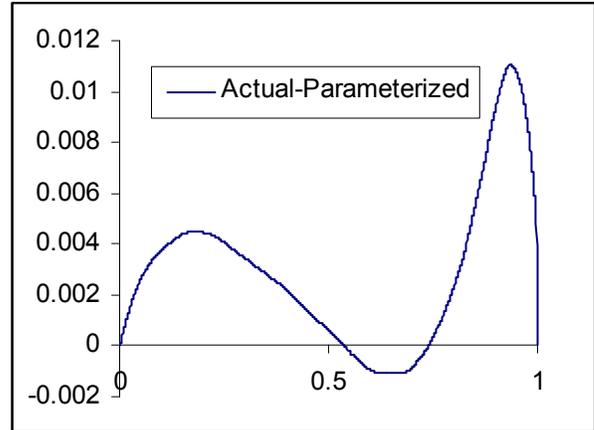
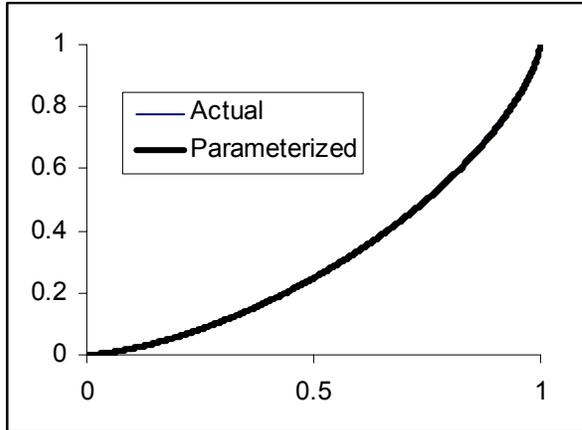
$$\theta \cdot \left(1 - \frac{P_t(\theta - 1)}{P_t(\theta)} \right).$$

Since $P(\theta)$ based on the parametric approximation is biased down by

more than $P(\theta - 1)$, the sensitivity of poverty to growth will also be biased down (in absolute value).

Figure A2: Actual and Parameterized Lorenz Curves for Ghana 1998/99

Lorenz Curve



Logarithm of Quantile Function

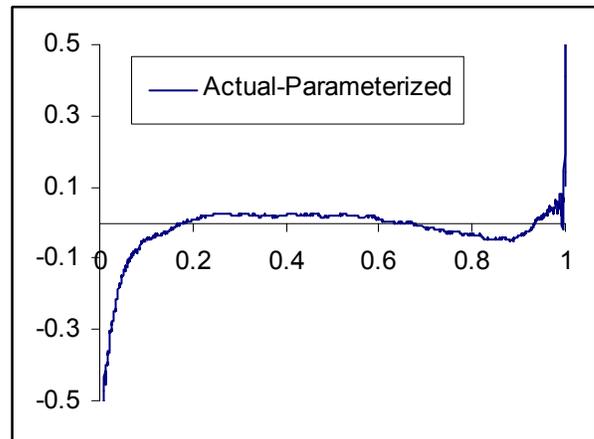
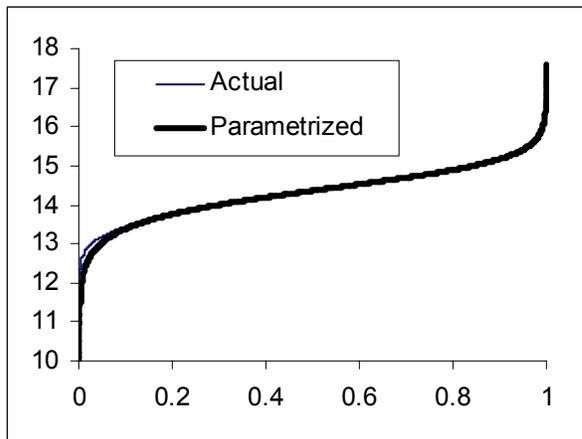


Table A2: Actual and Estimated Poverty Measures for Ghana 1998/99

	Actual	Estimated	(Actual- Estimated)/Actual
<i>Poverty Line = 100% of Mean</i>			
Headcount	0.597	0.591	0.011
Poverty Gap	0.246	0.245	0.004
Squared Poverty Gap	0.132	0.136	-0.033
Watts Index	0.371	0.387	-0.044
<i>Poverty Line = 50% of Mean</i>			
Headcount	0.228	0.221	0.029
Poverty Gap	0.068	0.077	-0.133
Squared Poverty Gap	0.029	0.039	-0.342
Watts Index	0.092	0.116	-0.262
<i>Log Change in Poverty/Log Change in Mean</i>			
Headcount	-0.964	-0.982	0.019
Poverty Gap	-1.288	-1.158	-0.100
Squared Poverty Gap	-1.509	-1.248	-0.174
Watts Index	-1.397	-1.208	-0.136