

A Normal Relationship?

Poverty, Growth, and Inequality

J. Humberto Lopez and Luis Servén*
The World Bank

Abstract

Using a large cross-country income distribution dataset spanning close to 800 country-year observations from industrial and developing countries, this paper shows that the size distribution of per capita income is very well approximated empirically by a lognormal density. Indeed, the null hypothesis that per capita income follows a lognormal distribution cannot be rejected -- although the same hypothesis is unambiguously rejected when applied to per capita consumption. The paper shows that lognormality of per capita income has important implications for the relative roles of income growth and inequality changes in poverty reduction. When poverty reduction is the overriding policy objective, poorer and relatively equal countries may be willing to tolerate modest increases in income inequality in exchange for faster growth -- more so than richer and highly unequal countries.

JEL classification codes: C12, C23, D31, I32.

World Bank Policy Research Working Paper 3814, January 2006

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the view of the World Bank, its Executive Directors, or the countries they represent. Policy Research Working Papers are available online at <http://econ.worldbank.org>.

*This paper was prepared for the World Bank's Latin America Regional Studies program. We thank Lant Pritchett for comments and Patricia Macchi for excellent research assistance.

I. Introduction

In recent years, poverty reduction has been formally enshrined as *the* goal of development policy worldwide, and a rapidly expanding analytical and empirical literature has sought to clarify whether poverty changes are driven mainly by growth in aggregate income or by growth in the relative incomes of the poor. The question is of more than scholarly interest, because it has major implications for the roles of growth-promoting and inequality-reducing policies in the poverty reduction process. For example, if trends in relative incomes were found to account for the lion's share of poverty changes, policy makers might face a tradeoff between fast growth and rapid poverty reduction.¹

In a recent contribution, Kraay (2005) decomposes poverty changes into three ingredients: (i) growth in average incomes; (ii) the sensitivity of poverty to growth; and (iii) changes in the distribution of income. In a large cross-country sample, he finds that growth in average incomes accounts for some 70 percent of the variation in (headcount) poverty changes in the short run, and over 95 percent in the medium to long run. In contrast, cross-country differences in the sensitivity of poverty to growth play a minimal role. Together, these results suggest that growth-oriented policies hold the key to poverty reduction.

In turn, Ravallion (1997, 2004) presents an empirical model relating poverty changes to the distribution-corrected rate of growth.² His estimates underscore the key role of initial inequality: depending on its level, a one-percent increase in income levels reduces poverty by as much as 4.3 percent (in very low inequality countries) or as little as 0.6 percent (in high inequality countries). This suggests that fast poverty reduction will be hard to achieve without declines in inequality, especially in very unequal countries.³

This paper reassesses the roles of growth and inequality for poverty reduction from a different perspective, based on the use of a parametric approach to model the size distribution of income. Thus, the paper follows an abundant literature spanning over a century -- from Pareto (1897) to Gibrat (1931), Kalecki (1945), Rutherford (1955), Metcalf (1969), Singh and Maddala (1976) and Bourguignon (2003) -- that has attempted to approximate the distribution of income using a variety of functional forms. Specifically, the paper uses a large cross-country database including both industrial and developing countries and spanning almost 40 years to test the null hypothesis that the size distribution of per capita income can be described by a lognormal density.

¹ This, of course, need not always be the case, since many policies are likely to be both growth-promoting and equality-enhancing. But some empirical evidence suggests that not all policies have this feature (Barro 2000, Lundberg and Squire 2003, Lopez 2004), and some may force policy makers to face a trade off between faster growth and increasing inequality.

² Specifically, Ravallion (1997) interacts the growth rate with one minus the initial Gini coefficient, whereas Ravallion (2004) considers a distributional term of the form $(1-Gini)^\theta$ with $\theta > 1$, to incorporate possible nonlinear interaction effects between the growth elasticity of poverty and initial inequality.

³ A similar point is made by Bourguignon (2004), who presents a simulation model calibrated on Mexican data. When no change takes place in the distribution of income, the model shows that a per capita growth rate of 3 percent per year over a 10-year period lowers the poverty rate by about 7 percentage points. When the same growth rate is accompanied by declining inequality (a reduction in the Gini coefficient by 10 percentage points), poverty falls twice as much -- over 15 percentage points.

A parametric approach offers a number of advantages over empirical variance decompositions and reduced-form regressions. First, it allows a systematic assessment of the role of country-specific initial conditions for the poverty-reducing effects of growth and distributional change. While the literature has stressed how initial inequality shapes the growth elasticity of poverty, little if any attention has been paid to the roles of the initial level of development or the poverty line itself for the choice of poverty-reducing policies. For example, should the balance between pro-growth and pro-distributional policies be the same in Zambia and El Salvador, which have similar Gini coefficients, despite their wide disparity in terms of per capita income (US\$361 and US\$2,200 respectively in 2002)? The approach taken in this paper allows a rigorous answer to this kind of question.

Second, the parametric approach permits removing the straightjacket of a common poverty line across countries, which otherwise is virtually unavoidable in cross-country empirical work. Most of the available studies use poverty statistics based on the international PPP US\$1 a day poverty line. However, such definition is of little interest for many middle income developing countries (not to mention industrial economies). As an example, in Argentina the headcount poverty was in 2002 close to 60 percent when calculated on the basis of the nationally-defined poverty line, while internationally-comparable poverty indicators based on a dollar-a-day poverty line would place the poverty rate around 3 percent. One implication of these diverging assessments is that the findings from empirical work using cross-country poverty databases, based on a common poverty line, tend to give more weight to countries where the US\$1 a day poverty measure makes sense -- i.e., low-income countries where our results below suggest that growth should be expected to dominate distributional change from the viewpoint of poverty reduction.

Of course, the advantages of the parametric approach matter only if the chosen parameterization fits the data well. We find that the null hypothesis of lognormality cannot be rejected when applied to the distribution of per capita income, regardless of whether income is measured in gross terms (i.e. before taxes and transfers) or net terms (after taxes and transfers). However, the same null hypothesis is unambiguously rejected when applied to per capita consumption data. We conjecture that this rejection may be due to consumption smoothing, under which the log of consumption may not be normally distributed even if the log of income is.

The paper derives some implications of this result for the relative roles of growth and inequality in poverty reduction under alternative initial conditions. We highlight four main points: (i) inequality hampers poverty reduction, both because of its negative impact on the growth elasticity of poverty (as stressed in the literature) but, in most scenarios, also because of its negative impact on the inequality elasticity of poverty; (ii) for a given poverty line, the impact of growth on poverty is stronger in richer than in poorer countries, and hence the latter will find it harder than the former to achieve fast poverty reduction; (iii) the share of the variance of poverty changes attributable to growth should be generally lower in richer and more unequal countries; and (iv) given the initial levels of development and inequality, the relative poverty-reduction effectiveness of growth and inequality changes depends on the poverty line -- the higher the poverty line, the bigger the role of growth and the smaller the role of distributional change.

The rest of the paper is organized as follows. In Section II we describe the test for lognormality and the dataset we employ. Section III reports the empirical results. Section IV derives the implications of lognormality for the poverty reduction roles of growth and inequality changes. Section V offers some concluding remarks.

II. Testing for lognormality of the size distribution of per capita income

Attempts to model the size distribution of per capita income have a long tradition in the economics literature, dating back to Pareto (1897). The use of the lognormal function in this context was pioneered by Gibrat (1931), who found that it offered a good empirical fit to the observed distribution of income, and provided a theoretical justification based on a model in which individual incomes are subject to random proportionate changes.⁴

Gibrat's work was followed by a large literature extending his basic framework and offering additional empirical evidence. Kalecki (1945) modified Gibrat's original setup making negative income changes less likely at low income levels than at high ones, to account for the fact that the variance of log income remained relatively constant over time. Rutherford (1955) expanded Gibrat's model introducing birth and death considerations. He also presented empirical experiments based on the comparison of theoretical and observed quantiles of the distribution of income, searching for a functional form that would improve upon the lognormal.⁵

On the theoretical front, other subsequent papers developed rigorous models that under fairly general conditions yield lognormal distributions of earnings and/or wealth (Sargan 1957, Pestieau and Posen 1979). In turn, on the empirical front, the fit of the lognormal function to the observed distribution of income was found to be somewhat less satisfactory at the upper end of the distribution (Hill 1959, Cowell 1977), specifically the top 3-4 percentiles (Airth 1985). This prompted attempts to fit more complex functional forms -- displaced and/or truncated versions of the lognormal density (Metcalf 1969, Salem and Mount 1974) or alternative functional specifications (Fisk 1961, Salem and Mount 1974, Singh and Maddala 1976, McDonald 1984). Both strategies pose a tradeoff between goodness-of-fit and analytical tractability (Metcalf 1969), as well as interpretability of the parameters (Lawrence 1988), which explains the continuing popularity of the lognormal specification.⁶

More recently, Bourguignon (2003) has offered an indirect reassessment of the empirical validity of the lognormal approximation. He reports OLS regression estimates in a framework explaining the observed change in a selected poverty measure on the

⁴ Specifically, Gibrat (1931) argued that the good empirical performance of the lognormal density could be rationalized under the following three conditions: (i) in each period the distribution of income is derived from that of the previous period by assuming that the variable corresponding to each member of the distribution is affected by a small proportionate change; (ii) such proportions differ for different members of the distribution; and (iii) these differences are determined randomly according to a given frequency distribution.

⁵ Rutherford performed single-country estimations using data for the UK in 1949; USA in 1947 and 1948; Canada in 1947, Australia in 1951, Sri Lanka (Ceylon at the time the paper was written) in 1950 and Bohemia (modern Slovakia and the Czech Republic) in 1932.

⁶ For example, Dollar and Kraay (2002) resort to the assumption of lognormality in order to complete their data sample by generating quintile shares from Gini coefficients for those observations for which only the latter is available.

basis of two regressors: a “growth effect”, given by average income growth times the theoretical growth elasticity of poverty calculated under the lognormality assumption, and an “inequality effect”, given by the change in inequality (as measured by the standard deviation of log income) times the theoretical inequality elasticity of poverty as derived under the lognormality assumption as well. Under the null of lognormality, both regressors should carry coefficients equal to unity.

When this empirical approach is implemented on a sample of developing-country growth spells, with the change in headcount poverty as dependent variable, Bourguignon (2003) rejects the null of lognormality but still finds that the lognormal specification provides a good empirical approximation to actual poverty changes. When the dependent variable is instead the change in the poverty gap, the null hypothesis is still rejected, and in addition the fit of the regression is quite poor.

There are, however, two problems with this approach. The first one, noted by Bourguignon, is that the elasticity-based approach is valid only for infinitesimal changes in poverty and its determinants. Applying it to discrete changes can result in large approximation errors, especially given the long duration (ten years and over) of some of the spells in Bourguignon’s sample. The second problem is the implementation of the approach using poverty databases, which tend to be relatively small, typically include a considerable number of outliers,⁷ and often involve substantial measurement error. The latter is further exacerbated by first-differencing the data for the regressions, which raises the noise-to-signal ratio. Below we present a new test of lognormality of the distribution of income that is not subject to these concerns.

II.1 Empirical approach

In spirit, our approach is closest to that employed by Rutherford (1955). As noted above, for several countries (one at a time) he compared the observed quintiles of the distribution of income with their theoretical counterparts derived under the null hypothesis of lognormality. Formally, we exploit the one-to-one mapping that arises under lognormality between the Gini coefficient and the Lorenz curve $L(p)$ that describes the relative income distribution.⁸ Letting G and σ respectively denote the Gini coefficient, and the standard deviation of log income, Aitchison and Brown (1966) show that lognormality implies

$$\sigma = \sqrt{2} \Phi^{-1}\left(\frac{1+G}{2}\right), \quad (1)$$

and

$$L(p) = \Phi\left(\Phi^{-1}(p) - \sigma\right), \quad (2)$$

where $\Phi(\cdot)$ denotes the cumulative normal distribution. Hence a change in the Gini coefficient, and thus in σ , must be reflected in a matching change in the Lorenz curve

⁷ Kraay (2005) uses a filter to eliminate extreme observations from his poverty dataset. This results in the loss of over one-third of his original sample.

⁸ Recall that $L(p)$ is the aggregate income share of the bottom $100p$ percent of the population. Thus, $L(0)=0$ and $L(1)=1$.

(Aitchison and Brown 1966, chapter 11). Likewise, changes in the Lorenz curve itself can be mapped into changes in the Gini coefficient.

On a cross-country basis, what is usually available to the researcher is some summary information on the shape of the Lorenz curve. One such summary is provided by the income shares of the different quintiles of the population:

$$Q_{20j} = L(.2j) - L(.2(j-1)) \quad \text{for } j = 1, 2, 3, 4. \quad (3)$$

Given the one-to-one mapping between the Gini coefficient and the Lorenz curve that follows from (1) and (2), under lognormality there must be also a one-to-one mapping between the Gini coefficient and the quintile shares (3). Thus, a test of the null hypothesis of lognormality can be based on the comparison of the empirical quintiles, say E_{20j} , with their Gini-based theoretical counterparts Q_{20j} . Following this approach, a formal lognormality test can be performed on the basis of the regression model:

$$E_{20j}^{it} = \alpha + \beta Q_{20j}^{it} + v_j^{it}, \quad (4)$$

where $j=1,2,3,4$ denotes the income quintile; $i=1,2,\dots,N$ is a country index, and $t=1,2,\dots,T_i$ denotes the date of each income (or expenditure) survey available for country i . In general T_i will differ across countries, resulting in an unbalanced sample. In (4), the theoretical quintiles Q_{20j}^{it} are constructed on the basis of the observed Gini coefficients G^{it} , as implied by (1)-(3):

$$Q_{20j}^{it} = \Phi \left(\Phi^{-1}(.2j) - \sqrt{2} \Phi^{-1} \left(\frac{1+G^{it}}{2} \right) \right) - \Phi \left(\Phi^{-1}(.2(j-1)) - \sqrt{2} \Phi^{-1} \left(\frac{1+G^{it}}{2} \right) \right). \quad (5)$$

Testing for lognormality in (4) is equivalent to testing the joint null hypothesis:

$$\alpha=0; \beta=1. \quad (6)$$

We should note that the precise null hypotheses entertained here is that the size distribution of income is described by a *two-parameter* lognormal function. Strictly speaking, rejection of (6) does not quite amount to rejection of lognormality more generally, since the distribution of income might still be characterized by a *three-parameter* lognormal density. This could happen, for example, if per capita income follows a displaced lognormal distribution -- i.e., it is lognormal over the range above some unknown minimum level τ (where τ is expressed as a ratio to average per capita income). In such case,⁹ the variance of log income, and hence the Gini coefficient, remain unaffected, but from Aitchison and Brown (1969, p.15) it follows that

$$L(p) = p\tau + (1-\tau)\Phi \left(\Phi^{-1}(p) - \sigma \right). \quad (7)$$

⁹ Explored, for example, by Metcalf (1969), who fits a displaced lognormal to U.S. personal income.

Expression (3) can still be used to compute the theoretical quintiles that correspond to any given Gini coefficient and shift τ . However, a regression like (4) projecting observed income shares on their counterparts constructed under the null of lognormality, ignoring the shift of the distribution (that is, assuming $\tau = 0$ when in reality $\tau > 0$), will result in a positive intercept and a slope less than 1 under the null. The bigger the shift τ , the larger the constant and the smaller the slope.

Finally, even if the true income distribution is characterized by the conventional two-parameter lognormal function, the observed distribution may follow a more complex form if the availability of data is limited. Data might be completely unavailable outside some income range, like in the textbook truncation case, or availability could vary in some systematic fashion with the level of income, like in the model of survey non-compliance examined by Deaton (2004). Depending on the particulars, a number of possibilities arise regarding the distribution of observed income. Under some special assumptions (illustrated by Deaton 2004), it might still be described by a two-parameter lognormal, although both its mean and variance could differ from those of the true distribution, but under more general conditions it may be characterized by more complex truncated lognormal distributions, in which case the simple relations (1)-(2) break down and inference based on (4) will reject the null hypotheses (6).¹⁰

II.2 Estimation Issues

The choice of estimation technique for (4) is dictated by the properties of the residual term v_j^{it} . If the residuals are i.i.d., OLS suffices to test the null of lognormality. However, there are two reasons why the assumption of independence may not hold. First, the residuals for a given country may be correlated across different surveys. In this regard, Lopez (2004) finds that the Gini coefficient shows significant persistence over time, and this suggests that the discrepancy between observed quintile shares and nonlinear transformations of the Gini as in (5) may also show persistence. Second, for any given survey all four theoretical quintiles are derived from the same Gini coefficient, and hence the residuals of the four regression observations that result may be mutually correlated.

In these circumstances, OLS estimates of α and β will still be consistent but inference based on the usual OLS covariance matrix will be inappropriate. Of course, valid inference can be performed using a robust estimator of the covariance matrix of the OLS coefficients that takes into account the lack of independence of the residuals, as well as their potential heteroskedasticity. However, under appropriate assumptions about the structure of the residual covariance matrix, more efficient inference may be possible. Specifically, assume that the disturbance term follows an error-components model:

$$v_j^{it} = \mu_i + \varepsilon_j^{it}, \quad (8)$$

or the more general

¹⁰ For example, if the sample is truncated from below (i.e., low-income observations are lost) then it can be shown that linear regressions like (4) will yield negative intercepts and slopes above unity.

$$v_j^{it} = \mu_i + \eta_i^t + \varepsilon_j^{it}, \quad (9)$$

where μ_i is an unobservable country-specific effect, assumed to be i.i.d. with zero mean and variance σ_μ^2 ; η_i^t denotes an effect specific to the t th survey for the i th country, also assumed i.i.d. with zero mean and variance σ_η^2 , and ε_j^{it} denotes the residual disturbance, assumed i.i.d. with zero mean and variance σ_ε^2 .¹¹ The μ_i 's, η_i^t 's, and the ε_j^{it} 's are assumed mutually independent. Under these assumptions, the covariance structure of the error term is:

$$E(v_j^{it}, v_k^{ls}) = \begin{cases} \sigma_\mu^2 + \sigma_\eta^2 + \sigma_\varepsilon^2 & \text{if } i = l, s = t, j = k \\ \sigma_\mu^2 + \sigma_\eta^2 & \text{if } i = l, s = t, j \neq k \\ \sigma_\mu^2 & \text{if } i = l, s \neq t \\ 0 & \text{if } i \neq l. \end{cases} \quad (10)$$

These expressions define a two-way error-components model in which the survey-specific effect η_i^t is nested in the country-specific effect μ_i . If $\sigma_\eta^2 = 0$, (9) reduces to the standard error-components model (8).

The parameters of the nested error components model given by (4) and (9)-(10) can be estimated in a variety of ways, ranging from ANOVA-type to minimum quadratic-norm and maximum likelihood estimation.¹² On the whole, Monte Carlo evidence reported by Baltagi *et al.* (2001) suggests that the method chosen makes little difference for the estimates of the regression coefficients in (4). However, when -- like in our case -- one is also interested in the standard errors of the parameter estimates, as well as in the variance components themselves, maximum likelihood estimation offers the best performance, especially if the sample is severely unbalanced.¹³

II.3 Data

We implement the empirical approach described above using the Dollar and Kraay (2002) dataset, which comprises 794 country-year observations for which both the Gini coefficient and the quintile shares are available.¹⁴ The dataset combines observations for which both the Gini coefficient and the quintile shares refer to gross (i.e. before taxes and transfers) income (47 percent of the observations); net (i.e. after taxes and transfers) income (29 percent of the observations); and expenditure (24 percent).

Table 1 presents some descriptive statistics. The average Gini coefficient is roughly the same in the income and expenditure subsamples -- 0.37 for income and 0.38 for expenditure. This is somewhat surprising since on the basis of conventional

¹¹ These assumptions can be further relaxed to allow for heteroskedasticity of ε_j^{it} .

¹² See Baltagi *et al.* (2001) and Davis (2002) for discussion.

¹³ See Rabe-Hesketh *et al.* (2004) on the computational aspects of ML estimation in this context.

¹⁴ We discard the 158 extra observations for which Dollar and Kraay construct the quintile shares from the Gini coefficient using equations (1)-(3), and thus assuming that the distribution of income is lognormal -- which is precisely what our regressions aim to test.

smoothing arguments one would expect less dispersion in the expenditure-based surveys.¹⁵ One might wonder if in our sample the result is driven by the fact that the panel is unbalanced. But even if we give equal weight to each country's average Gini coefficient we find a similar picture: the resulting overall means are 0.40 for income-based and 0.42 for expenditure-based observations, respectively. The latter, however, are more concentrated around their mean. In other words, there is a lower frequency of countries with extreme (whether high or low) inequality in the expenditure-based subsample than in the income-based subsample.

On the other hand, there is a noticeable difference between the average Gini coefficients of the subsamples based on gross and net income. Gross income-based Gini coefficients average 0.40, whereas those based on net income average 0.33. Although one should be careful in attaching any particular economic interpretation to these figures, it is tempting to view them as reflecting the effect of government interventions that lead to income redistribution.

III Empirical results

We turn to the empirical implementation of the lognormality tests. We perform them on the full sample as well as different sub-samples defined by type of data -- i.e., according to whether the observations are based on expenditure or income surveys and, in the latter case, whether income is measured on a net or gross basis. Apart from allowing some robustness checks, this differentiation is also of interest for two other reasons. First, if the distribution of income is lognormal, and households engage in consumption smoothing, the distribution of consumption will not be lognormal in general.¹⁶ Hence the common practice of pooling together income- and expenditure-based observations in applied work could yield misleading test results in our case. Second, gross income could be lognormally distributed while net income is not -- for example if taxes and transfers are lump-sum rather than proportional to income.

To be specific, in addition to (i) the full sample we also consider sub-samples of (ii) income-based observations only (labeled "Income" for short); (iii) expenditure-based observations only ("Expenditure"); (iv) gross income-based observations only ("Gross Income"); (v) a combination of expenditure and net income-based observations ("Net"); and finally (vi) net income-based observations only ("Net Income").

For the full sample, Figure 1.(a) shows a scatter plot of the observed quintile shares (vertical axis) against their theoretical counterparts, as computed under the null hypothesis of lognormality (horizontal axis). The data points cluster along the 45-degree line, suggesting that the lognormal distribution provides a fairly close approximation to

¹⁵ For this reason, Forbes (2000) and Deininger and Squire (1996) raise expenditure-based Gini coefficients by 0.066 to make them comparable with income-based ones.

¹⁶ Under consumption smoothing, current consumption will depend on some weighted sum of current and future anticipated income. In the textbook permanent-income model, the consumption level of household h in period t would take the form $C_t^h = \theta \left[A_t^h + \sum_{s=t}^{\infty} (1+r)^{s-t} E_t \left[x_s^h \right] \right]$, where x is income, A denotes financial wealth, r is the real interest rate, θ is a parameter, and E_t denotes the conditional expectation. Even if x_s is lognormally distributed, the infinite sum in the square brackets will not be in general, a simple consequence of the fact that the sum of lognormally-distributed variables is not itself lognormally distributed.

the size distribution of per capita income / expenditure as summarized by the quintile shares. Figures 1.(b)-1.(f) present similar plots for the various subsamples.¹⁷

Table 2 presents the results of OLS-based lognormality tests for the full sample and the various subsamples described above. The standard errors of the estimates are computed using a clustering procedure to allow for residual dependence; they also allow for heteroskedasticity.¹⁸

The first thing that stands out in the table is the excellent fit of the regressions in all the samples considered, with R^2 ranging from 0.95 to 0.98. For a cross-country sample of this magnitude, such fit is remarkable. In turn, the regression slopes and intercepts are very close to their expected values under the null of one and zero, respectively. It is worth noting that in the samples including expenditure observations (the first, third and fifth columns) the estimated slopes are slightly below one, while the opposite happens in the regressions including only income-based observations. Formally, we can reject the null of unit slope in the Expenditure and Net subsamples (third and fifth columns). In turn, the estimated intercepts are positive in the samples including expenditure-based observations and negative in those including only income-based observations. Like with the slopes, in the Expenditure and Net subsamples we can also reject the null of zero intercept.

The bottom of Table 2 reports Wald tests of the null hypothesis of lognormality (6). Under the null, the test statistic follows a chi-square distribution with two degrees of freedom. As would be expected in the light of the point estimates, the null can be rejected at the 5 percent level in the two samples in which expenditure-based observations represent a sizeable share of the total number of data points. In contrast, the samples containing only income-based observations show little evidence against the null -- the p-values range from 0.62 to 0.93. In the full sample, in which expenditure-based observations represent only about 20 percent of the total, we also fail to reject the null, with a p-value of 0.35.

On the whole, the lognormality tests based on the OLS estimates suggest that the size distribution of income and expenditure follow significantly different patterns. In contrast, the distinction between gross and net income seems to be of little consequence. Table 3 repeats the same tests of Table 2 now based on ML estimation of the nested error-component model given by (4) and (9)-(10). In addition to the information contained in the preceding table, Table 3 also reports the estimated standard deviations of the error components in (9), and the results of tests of their individual and joint significance.

Inspection of Table 3 reveals a picture very similar to the one emerging from Table 2, in terms of both point estimates and standard errors. The pattern of signs and magnitudes of the point estimates across samples is the same as before. The middle block of the table reports the estimated standard deviations of the error components. The

¹⁷ Figures 1(e) and 1(f) show an apparent outlying observation (corresponding to Q2 in Norway 1989), which might be viewed as a candidate for removal from the sample. This, however, would be of no consequence for any of the empirical results in the paper.

¹⁸ The clustering is done by country. Doing it instead by survey yields slightly larger standard errors but does not cause any qualitative changes on the results of the hypothesis tests.

standard deviation of the survey-specific effect σ_η is in all cases quite small -- indeed, much smaller much than that of the country-specific effect σ_μ .

The lognormality tests yield the same qualitative conclusion as before: the evidence from the full sample and the three income-based samples is consistent with the null hypothesis of lognormality, as reflected in p-values ranging from 0.41 in the full sample to 0.92 in the Gross Income sub-sample. In contrast, the null is rejected at the 5 percent level in the Expenditure and Net subsamples.

The last three rows of Table 3 report tests of significance of the error components. The null hypothesis that the variances of the two error components are jointly zero ($\sigma_\mu^2 = \sigma_\eta^2 = 0$) is rejected in all cases. As for the nested survey-specific component, its variance σ_η^2 is insignificant in the Income and Expenditure subsamples; in addition, it falls just short of 5 percent significance in the Gross Income and Net Income subsamples. In turn, the variance of the country-specific component σ_μ^2 is significant at the 5 percent level in four of the six samples, and at the 10 percent level in the other two -- the Expenditure and Gross Income subsamples.

On the whole, we may view these tests as generally supportive of the nested error component specification, except in the Income and Expenditure subsamples, where the test results suggest that a one-way model might be sufficient. To investigate this further, Table 4 reports the results of estimating the standard random effects model given by (4) and (8) for the Income and Expenditure subsamples. The results in the table continue to support the same basic message as before. The parameter estimates are almost identical to those in the preceding table, as are the results of the lognormality tests: strong rejection of the null in the expenditure-based sample, and failure to reject it in the income-based sample.

In summary, the empirical tests reported in this section show that in a large cross-country sample the observed distribution of per capita income is consistent with the hypothesis of lognormality -- regardless of whether income is measured before or after taxes and transfers. In contrast, the same tests reject lognormality of the distribution of per capita expenditure -- although the lognormal specification can still account for a very high proportion of the observed variation in quintile shares even in the expenditure data.

IV. Growth, inequality and poverty

The finding that per capita income follows a lognormal distribution has important practical implications for assessing the respective contributions of growth and inequality to poverty changes. The reason is that under lognormality we can derive simple closed-form expressions for these contributions, which depend only on the prevailing degree of inequality, and on the poverty line relative to mean per capita income.

For concreteness, let us focus on the Foster-Greer-Thorbecke (1984) [henceforth FGT] class of poverty measures, which includes those most widely used in applied work. They are given by the general expression:

$$P_\alpha = \int_0^z \left[\frac{z-x}{z} \right]^\alpha f(x) dx, \quad (11)$$

where $\alpha \in \{0,1,2\}$ is a parameter of inequality aversion, z is the poverty line, x is income, and $f(.)$ is the density function of income. When $\alpha = 0$, (11) reduces to the familiar headcount ratio, which measures the share of the population below the poverty line z . For $\alpha = 1$, we get the FGT measure P_1 , known as the poverty gap, which weighs each poor individual by his / her distance to the poverty line -- heuristically, it provides a measure of the depth of poverty. Finally, for $\alpha = 2$, we have the squared poverty gap P_2 , which weights each poor individual by the square of his/her income shortfall; thus larger shortfalls are weighted more than proportionately.

Denoting average per capita income by ν (i.e., $E(x)=\nu$), the appendix shows that under lognormality we can write

$$P_\alpha = P_\alpha(z/\nu, G). \quad (12)$$

Thus, poverty depends only on the Gini coefficient and the poverty line relative to mean income. Equation (12) provides a starting point to analyze the relative contributions of growth and changes in inequality to poverty reduction. For $\alpha = 0$, Figure 2 plots a set of iso-poverty curves (i.e., level sets of equation (12)); each of them depicts combinations of Gini coefficients and mean per capita income / poverty line ratios (ν/z) that yield a constant poverty headcount P_0 .¹⁹ Curves to the Northeast of the graph correspond to higher levels of the poverty rate.

The slope of these curves depicts the changing tradeoff between growth and redistribution. The steeper the slope, the bigger the decline in the Gini coefficient required to keep poverty constant in the face of a given decline in the ratio of mean income to the poverty line. The curves become increasingly steep, and closer to one another, as we move downward along them. In other words, the more equal and the poorer the economy (as reflected, respectively, by a lower Gini coefficient and a lower mean income / poverty line ratio), the bigger the change in the Gini coefficient required to offset a given change in mean income relative to the poverty line -- i.e., the more effective growth will be relative to redistribution in attacking poverty. As the economy becomes richer and more unequal -- i.e., as we move to the Northwest of the graph -- the curves become less steep, and therefore a smaller change in the Gini coefficient is now needed to offset a given change in mean income relative to the poverty line -- i.e., distributional change now plays a relatively larger role in poverty changes.

We can gauge better the relative roles of growth and distributional change by evaluating numerically the elasticity of poverty with respect to each of them. From (12), for a given poverty line we can write

$$\frac{dP_\alpha}{P_\alpha} = \eta_\nu^\alpha \frac{d\nu}{\nu} + \eta_G^\alpha \frac{dG}{G}. \quad (13)$$

¹⁹ Of course, similar curves can be drawn for $\alpha \in \{1,2\}$.

Here η_v^α and η_G^α respectively are the elasticities of P_α with respect to growth²⁰ and inequality. The appendix derives their exact expressions under the assumption of lognormality, and shows that they depend only on z/v and G .

Tables 5 to 7 report, for the three FGT poverty measures, the values of η_v^α and η_G^α that result from various combinations of the Gini coefficient G and the ratio of per capita income to the poverty line (v/z). In the tables, G runs from 0.3 to 0.6 and v/z from 1 to 6.²¹

Inspection of these tables confirms the well-known result (e.g., Ravallion 1997, 2004; Bourguignon 2003) that the growth elasticity of the various FGT measures is smaller (in absolute value) the higher the level of inequality. Thus, inequality hampers the poverty-reducing effect of growth, as stressed in the literature. In addition, however, the tables show that poverty itself (as measured by low per capita income) is another barrier to poverty reduction: in all three tables, for a given Gini coefficient, the growth elasticity of poverty declines rapidly (in absolute value) as average income declines in relation to the poverty line. This suggests a triple poverty-reducing effect of growth: first, the direct effect of income growth on the average level of income; second, the indirect impact that arises from higher average income via the correspondingly higher growth elasticity of poverty; and third, the indirect impact that arises from the higher average income via the correspondingly higher growth elasticity of poverty.

Under most scenarios, inequality itself also has a doubly deterrent effect on poverty reduction. In addition to lessening the growth elasticity of poverty, as just noted, higher inequality also lessens the impact of progressive distributional change itself on poverty -- i.e., in tables 5-7 the inequality elasticity falls as inequality rises, for a given value of average income relative to the poverty line. However, the relationship is highly nonlinear, and at very low levels of development (captured in the tables by values of (v/z) close to one) its sign is reversed, so that a higher Gini coefficient is associated with a higher inequality elasticity, as shown in the last line of Tables 5 -7.

The implication is that, given a common poverty line, poorer and more equal countries may be in a position to afford some growth-inequality tradeoffs. For low values of (v/z) , the poverty-reducing effects of growth outweigh the poverty-raising effects of a worsening distribution of income. In other words, very poor countries may be willing to tolerate modest deteriorations in income equality in exchange for faster growth. Such tradeoff is much more problematic in richer and highly unequal countries, where small

²⁰ Strictly speaking, η_v^α is the elasticity of poverty with respect to mean income, rather than growth, but we follow the standard practice in the literature and use the term “growth elasticity” to refer to the change in poverty that results when income increases by one percent, at given inequality.

²¹ The range from 0.30 to 0.50 amounts roughly to the mean Gini plus/minus one standard deviation in the overall sample (see Table 1). For the simulations, we raise the upper end to 0.60, a value reminiscent of the inequality encountered in some Latin American countries. As for the ratio of mean income to the poverty line, with poverty defined by the standard US\$1 per person per day, the range from 1 to 6 is equivalent to annual per capita income levels between US\$365 and US\$2140 -- which correspond roughly to those of Zambia and El Salvador respectively. Alternatively, with poverty defined by a US\$2 per person per day poverty line, the range for (v/z) amounts to per capita income levels from US\$730 to US\$4280 -- approximately the income levels of Indonesia and Uruguay, respectively.

inequality increases have a much larger poverty-raising effect, and hence policy makers may be more willing to accept a modest growth decline in exchange for a reduction in inequality.

Another way to gauge the relative importance of growth and redistribution from the poverty reduction viewpoint is to measure the respective contributions of growth and inequality shocks to the observed variation in poverty. This is the approach followed by Kraay (2005). Under lognormality, it follows from (13) that the variance of poverty changes can be approximated as

$$\text{var}\left(\frac{dP_\alpha}{P_\alpha}\right) = (\eta_v^\alpha)^2 \text{var}\left(\frac{dv}{v}\right) + (\eta_G^\alpha)^2 \text{var}\left(\frac{dG}{G}\right) + 2\eta_v^\alpha \eta_G^\alpha \text{cov}\left(\frac{dv}{v}, \frac{dG}{G}\right). \quad (14)$$

In the general case of nonzero covariance between growth and inequality changes, absent any information on their mutual causal precedence, we can use the simplifying assumption that half the covariance can be attributed to growth, and the other half to inequality changes. In such case the share of the total variance attributable to the growth component is:

$$\text{var}\left(\frac{dP_\alpha}{P_\alpha}\right)_v = \frac{(\eta_v^\alpha)^2 \text{var}\left(\frac{dv}{v}\right) + \eta_v^\alpha \eta_G^\alpha \text{cov}\left(\frac{dv}{v}, \frac{dG}{G}\right)}{\text{var}\left(\frac{dP_\alpha}{P_\alpha}\right)}. \quad (15)$$

However, to implement (15) numerically we need three more ingredients, namely the variance of growth, the variance of inequality changes, and the covariance of growth with inequality changes. Table 8 reports these statistics computed on the basis of two alternative datasets: the one used thus far in this paper (i.e., Dollar and Kraay 2002), and the Povmonitor²² database. Figure 3 presents the respective scatter plots. Apart from coverage (much more limited in Povmonitor), the main difference between both databases is that the Dollar and Kraay income data is based on National Accounts, whereas that in Povmonitor is based on household surveys.

Inspection of Table 8 suggests that the different coverage of the two databases is of little consequence for the volatility of changes in the (log) Gini coefficient: in both cases the standard deviation is about 0.05. However, the volatility of growth differs more markedly across the two databases: in the survey-based Povmonitor data, the standard deviation of growth (0.06) is slightly higher than that of inequality changes, while the opposite happens in the National Accounts-based Dollar and Kraay data, in which the standard deviation of growth (0.04) is lower than that of inequality changes. On the other hand, Table 8 also shows that growth and changes in inequality are nearly uncorrelated in both databases -- a conclusion confirmed by the scatter plots in Figure 3. Although the correlation coefficients in the table are of opposite signs (negative in the Dollar and Kraay data, and positive in the Povmonitor data), both are insignificantly different from zero.²³

²² <http://www.worldbank.org/povmonitor>.

²³ Similar evidence is provided by Deininger and Squire (1996), Chen and Ravallion (1997) and Easterly (1999).

Using these results, Tables 9 and 10 report the simulated values of expression (15) for alternative values of the Gini coefficient and the mean income / poverty line ratio, using the variance and covariance patterns of growth and inequality changes shown in the preceding table. Table 9 uses the values from the Dollar and Kraay sample, while Table 10 uses the values from Povmonitor. In both tables, a number close to 1 means that in the scenario in question changes in the poverty measure of interest are mainly driven by growth, whereas a number close to zero means that they are mainly driven by changes in inequality.²⁴

The simulations in Tables 9 and 10 suggest that, consistent with the previous discussion, in poorer and more equal countries growth accounts for a larger share of the variance of poverty changes. Conversely, inequality changes tend to play a more prominent role in richer and/or more unequal countries. This again suggests that in the former countries growth should be expected to be the main driver of poverty reduction, while the opposite would happen in the latter countries.

Notice also that, for any given configuration of Gini coefficient and per capita income / poverty line ratio, the relative contribution of growth to the overall variance of poverty changes declines as the poverty measure of interest varies from headcount poverty to the poverty gap and then its square. As noted by Kraay (2005), this is a natural consequence of the fact that more bottom-sensitive poverty measures place more weight on changes in the distribution of income than on changes in average income.

So far we have implicitly viewed alternative values of (ν/z) as reflecting different levels of average per capita income with a given poverty line. But they could also be interpreted the other way around, namely reflecting alternative poverty lines with a given level of average per capita income. In this view, the numerical results above imply that as the relevant poverty line z becomes more generous -- i.e., as (ν/z) declines -- the relative role of growth in the overall variation of poverty changes must go up as well, and other things equal this offers a rationale for shifting poverty reduction priorities in favor of growth-oriented policies. In the limit, as (ν/z) falls to zero -- so that the poverty rate approaches 1 -- distributional change becomes completely ineffective for poverty reduction. In other words, given the choice of poverty measure P_α and the initial conditions in terms of average income and inequality, the location of the poverty line is a key determinant of the relative effectiveness of growth and redistribution for poverty reduction.

V. Conclusions

The focus on poverty reduction as the key objective of development policy has opened a debate on the relative merits of aggregate growth and distributional change as anti-poverty strategies, and the conditions under which one may be more effective than the other. In this paper we have reexamined that question using a parametric approach to model the distribution of per capita income. The parametric approach has a long tradition in the literature, going back over a century. One of its key advantages is that it allows a

²⁴Note that in this calculation we are implicitly assuming that the same variance / covariance pattern applies regardless of the particular configuration of Gini coefficient and mean income/ poverty line ratio.

systematic assessment of how country-specific initial conditions -- inequality, level of development and poverty line definition -- affect the poverty-reduction effectiveness of growth and income redistribution.

The paper's approach is based on the use of a lognormal approximation to the size distribution of per capita income. We implement this approach in two stages. First, we perform empirical tests of lognormality of the observed distribution of income, using a large cross-country income and expenditure distribution dataset covering over 100 countries and 40 years. Our testing strategy is based on the comparison of the observed quintile shares of the distribution with their theoretical counterparts under the null hypothesis of lognormality.

The empirical tests, performed on the full data sample as well as a variety of subsamples, are very supportive of the lognormal approximation to the distribution of per capita income, but less so for per capita expenditure. In the former case, the null of lognormality cannot be rejected in any of the samples considered; in the latter case, it is consistently rejected. In both cases, however, the lognormal specification yields a very close approximation to the observed distribution.

Lognormality of the distribution of income allows us to derive, at the second stage of the analysis, some qualitative and quantitative implications for the relative roles of growth and inequality in poverty reduction under alternative initial conditions, using a variety of poverty measures.

Our conclusions can be summarized in four main points. First, inequality hampers poverty reduction, not only because of its negative impact on the growth elasticity of poverty (as stressed in the literature) but, in most scenarios, also because of its negative impact on the inequality elasticity of poverty. Second, for a given poverty line, the impact of growth on poverty is stronger in richer than in poorer countries, and hence the latter will find it harder than the former to achieve fast poverty reduction. Third, the share of the overall variance of poverty changes attributable to growth should be generally lower in richer and more unequal countries. And fourth, for given initial levels of development and inequality, the relative poverty-reduction effectiveness of growth and inequality changes depends on the poverty line -- the higher the poverty line, the bigger the role of growth and the smaller the role of distributional change.

References

- Aitchison, J. and J. Brown (1966): *The Lognormal Distribution*, Cambridge, Cambridge University Press.
- Airth, A. (1985): "The Progression of Wage Distributions", *Eurostat News* (special issue), 139-161.
- Baltagi, B., S. Heun Song and B. Cheol Jung (2001): "The Unbalanced Nested Error Component Model Regression Model", *Journal of Econometrics*, 101, 357-381.
- Barro, R. (2000): "Inequality and Growth in a Panel of Countries", *Journal of Economic Growth*, 5, 5-32.
- Bourguignon, F. (2003): "The Growth Elasticity of Poverty Reduction; Explaining Heterogeneity across Countries and Time Periods", in T. Eicher and S. Turnovsky, eds. *Inequality and growth: Theory and Policy Implications*. Cambridge: The MIT Press.
- Bourguignon, F. (2004): "The Poverty-Growth-Inequality Triangle", mimeo, The World Bank.
- Chen, S. and M. 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.
- Cowell, F. (1977): *Measuring Inequality*. Oxford: Phillip Allen.
- Davis, P. (2002): "Estimating Multi-way Error Components Models with Unbalanced Data Structures", *Journal of Econometrics* 106, 67-95.
- Deaton, A. (2004): "Measuring Poverty in a Growing World (Or Measuring Growth in a Poor World)", mimeo, Woodrow Wilson School.
- 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 (2002): "Growth is Good for the Poor", *Journal of Economic Growth*, 7, 195-225.
- Easterly, W. (1999): "Life during Growth", *Journal of Economic Growth*, 4, 239-276.
- Fisk, P. (1961): "The Graduation of Income Distribution", *Econometrica*, 29, 171-185.
- Forbes, K. (2000): "A Reassessment of the Relationship between Inequality and Growth", *American Economic Review*, 90, 869-897.
- Foster, J., J. Greer, and E. Thorbecke (1984): "A Class of Decomposable Poverty Measures", *Econometrica*, 52, 761-766.
- Gibrat, R. (1931): *Les inégalités économiques*, Paris, Sirey.

- Hill, T. (1959): "An Analysis of the Distribution of Wages and Salaries in Great Britain", *Econometrica*, 27, 355-381.
- Kakwani, N. (1990): "Poverty and Economic Growth: With Applications to Cote d'Ivoire". The World Bank, LSMS Working Paper No. 63.
- Kalecki, M. (1945): "On the Gibrat Distribution", *Econometrica*, 13, 161-170.
- Kraay, A. (2005): "When is Growth Pro-Poor? Evidence from a Panel of Countries", *Journal of Development Economics*. Forthcoming.
- Lawrence, R. J. (1988): "Applications in Economics and Business", in Crow, E. L. and Shimizu, K. (Eds.), *Lognormal Distributions: Theory and Applications*. New York: Dekker.
- Lopez, H. (2004): "Pro-Poor-Pro-Growth: Is there a Trade Off?", The World Bank, Policy Research Working Paper No. 3378.
- Lundberg M. and L. Squire (2003): "The Simultaneous Evolution of Growth and Inequality", *The Economic Journal*, 113, 326-344.
- Metcalf, C. (1969): "The Size Distribution of Personal Income during the Business Cycle", *The American Economic Review*, 59, 657-668.
- McDonald, J. (1984): "Some Generalized Functions for the Size Distribution of Income", *Econometrica*, 52, 647-663.
- Pareto, W. (1897): "Cours d'Economie Politique 2", F. Rouge, Lausanne.
- Pestieau, P. and U. Possen (1979): "A model of wealth distribution", *Econometrica* 47, 761-772.
- Rabe-Hesketh S., A. Skrondal and A. Pickels (2004): "GLLAMM Manual", U. C. Berkeley Division of Biostatistics WPS, 160.
- Ravallion, M. (1997): "Can High Inequality Development Countries Escape Absolute Poverty?", *Economics Letters*, 56, 51-57.
- Ravallion, M. (2004): "Pro-Poor Growth: A Primer", The World Bank, Policy Research Working Paper No. 3242.
- Rutherford, R.S.G. (1955): "Income Distributions: A New Model". *Econometrica*, 23, 277-294.
- Salem, A. and T. Mount (1974): "A Convenient Descriptive Model of Income Distribution: The Gamma Density", *Econometrica*, 42, 1115-1128.
- Sargan, D. (1957): "The distribution of wealth", *Econometrica* 25, 568-590.

Singh, S. and G. Maddala (1976): "A Function for Size Distribution of Incomes", *Econometrica*, 44, 963-970.

Table 1. Descriptive statistics, by sample

	Gini Coefficient	Q1	Q2	Q3	Q4
<i>All (3,176 observations)</i>					
Mean	0.37	0.07	0.11	0.16	0.22
Median	0.35	0.07	0.12	0.16	0.22
Standard Deviation	0.10	0.02	0.03	0.03	0.02
<i>Income (2,420 observations)</i>					
Mean	0.37	0.06	0.11	0.16	0.22
Median	0.35	0.06	0.12	0.17	0.23
Standard Deviation	0.10	0.02	0.03	0.03	0.03
<i>Gross Income (1,472 observations) ^{a/}</i>					
Mean	0.40	0.06	0.11	0.15	0.22
Median	0.37	0.05	0.11	0.16	0.23
Standard Deviation	0.11	0.02	0.03	0.03	0.03
<i>Net Income (892 observations) ^{a/}</i>					
Mean	0.33	0.07	0.12	0.17	0.23
Median	0.31	0.08	0.13	0.17	0.23
Standard Deviation	0.92	0.02	0.02	0.02	0.02
<i>Expenditure (756 observations)</i>					
Mean	0.38	0.07	0.11	0.15	0.21
Median	0.36	0.07	0.12	0.16	0.22
Standard Deviation	0.87	0.02	0.02	0.02	0.01

Note: (^{a/}) The total number of observations of the Gross Income and Net Income subsamples does not equal to the number of observations of the Income subsample because 56 observations are not classified.

Table 2. Lognormality tests
Pooled OLS

	Sample					
	All	Income	Expenditure	Gross Income	Net	Net Income
β	0.983	1.009	0.897 *	1.014	0.961 *	1.006
<i>s.e.</i>	0.014	0.015	0.012	0.021	0.016	0.017
α	0.002	-0.001	0.013 **	-0.001	0.005 **	-0.001
<i>s.e.</i>	0.002	0.002	0.002	0.002	0.002	0.002
R^2	0.96	0.96	0.98	0.95	0.98	0.98
# Observations	3,176	2,420	756	1,472	1,484	892
# Countries	130	98	65	75	97	55
Test of the joint hypothesis						
Ho: $\alpha=0$; $\beta=1$						
(p-value)	0.35	0.76	0.00	0.62	0.05	0.93

Notes: The table reports regression results with the observed quintile as dependent variable and the theoretical quintile as explanatory variable. All regressions include a constant. Robust standard errors using a clustering procedure are reported below the coefficients.

(*) Ho: $\beta=1$ rejected at the 5%.

(**) Ho: $\alpha=0$ rejected at the 5%.

Table 3. Lognormality tests
Nested Error Component Model

	Sample					
	All	Income	Expenditure	Gross Income	Net	Net Income
β	0.980	1.007	0.894 *	1.009	0.960 *	1.009
<i>s.e.</i>	0.015	0.016	0.012	0.023	0.016	0.017
α	0.002	-0.001	0.013 **	-0.001	0.005 **	-0.001
<i>s.e.</i>	0.002	0.002	0.002	0.003	0.002	0.002
# Observations	3,176	2,420	756	1,472	1,484	892
# Countries	130	98	65	75	97	55
σ_ε	1.00E-02	1.24E-02	7.35E-03	1.41E-02	2.59E-02	8.62E-03
σ_η	3.24E-07	3.67E-12	1.24E-08	3.51E-07	6.55E-07	4.69E-08
σ_μ	2.68E-03	3.42E-03	2.10E-03	5.20E-03	1.90E-03	1.92E-03
Hypothesis tests						
(p-values)						
Ho: $\alpha=0$; $\beta=1$	0.41	0.90	0.00	0.92	0.05	0.80
Ho: $\sigma_\eta=\sigma_\mu=0$	0.00	0.00	0.02	0.00	0.00	0.00
Ho: $\sigma_\eta=0$	0.04	0.50	0.50	0.08	0.03	0.07
Ho: $\sigma_\mu=0$	0.00	0.03	0.08	0.08	0.00	0.01

Notes: The table reports regression results with the observed quintile as dependent variable and the theoretical quintile as explanatory variable. All regressions include a constant. Robust standard errors are reported below the coefficients.

(*) Ho: $\beta=1$ rejected at the 5%.

(**) Ho: $\alpha=0$ rejected at the 5%.

Table 4. Lognormality tests
Random Effects Model

	Sample	
	Income	Expenditure
β	1.007	0.894 *
<i>s.e.</i>	0.016	0.012
α	-0.001	0.013 **
<i>s.e.</i>	0.002	0.002
# Observations	2,420	756
# Countries	98	65
Test of the joint hypothesis		
Ho: $\alpha=0$; $\beta=1$		
(p-value)	0.90	0.00

Notes: The table reports regression results with the observed quintile as dependent variable and the theoretical quintile as explanatory variable. All regressions include a constant. Robust standard errors are reported below the coefficients.

(*) Ho: $\beta=1$ rejected at the 5%.

(**) Ho: $\alpha=0$ rejected at the 5%.

Table 5. Headcount: Theoretical elasticities under lognormality

		<u>Growth Elasticity</u>			
		Gini Coefficient			
Mean Income /	Poverty Line	0.30	0.40	0.50	0.60
6		-6.05	-3.25	-1.95	-1.22
3		-3.94	-2.18	-1.33	-0.86
2		-2.80	-1.60	-1.01	-0.66
1.5		-2.06	-1.23	-0.80	-0.54
1		-1.16	-0.78	-0.55	-0.39

		<u>Inequality Elasticity</u>			
		Gini Coefficient			
Mean Income /	Poverty Line	0.30	0.40	0.50	0.60
6		12.34	7.38	5.10	3.89
3		5.17	3.28	2.42	1.97
2		2.48	1.70	1.35	1.18
1.5		1.20	0.92	0.81	0.77
1		0.18	0.24	0.29	0.35

Note: The table reports the theoretical growth and inequality elasticities, computed under the assumption of log normality, as a function of the ratio of mean income/poverty line and the Gini coefficient.

Table 6. Poverty gap: Theoretical elasticities under lognormality

		<u>Growth Elasticity</u>			
		Gini Coefficient			
Mean Income /	Poverty Line	0.30	0.40	0.50	0.60
6		-6.45	-3.59	-2.22	-1.44
3		-4.45	-2.57	-1.64	-1.09
2		-3.37	-2.02	-1.32	-0.90
1.5		-2.68	-1.67	-1.12	-0.77
1		-1.83	-1.23	-0.86	-0.62

		<u>Inequality Elasticity</u>			
		Gini Coefficient			
Mean Income /	Poverty Line	0.30	0.40	0.50	0.60
6		14.08	9.01	6.64	5.38
3		6.70	4.69	3.73	3.22
2		3.82	2.92	2.49	2.27
1.5		2.36	1.98	1.81	1.73
1		1.03	1.05	1.08	1.13

Note: The table reports the theoretical growth and inequality elasticities, computed under the assumption of log normality, as a function of the ratio of mean income/poverty line and the Gini coefficient.

Table 7. Squared poverty gap: Theoretical elasticities under lognormality

		<u>Growth Elasticity</u>			
		Gini Coefficient			
Mean Income /	Poverty Line	0.30	0.40	0.50	0.60
	6	-6.79	-3.85	-2.41	-1.59
	3	-4.84	-2.85	-1.84	-1.24
	2	-3.80	-2.32	-1.53	-1.05
	1.5	-3.12	-1.97	-1.33	-0.92
	1	-2.27	-1.52	-1.06	-0.76

		<u>Inequality Elasticity</u>			
		Gini Coefficient			
Mean Income /	Poverty Line	0.30	0.40	0.50	0.60
	6	15.58	10.34	7.83	6.46
	3	7.98	5.80	4.71	4.11
	2	4.92	3.88	3.34	3.05
	1.5	3.31	2.81	2.55	2.42
	1	1.73	1.69	1.68	1.70

Note: The table reports the theoretical growth and inequality elasticities, computed under the assumption of log normality, as a function of the ratio of mean income/poverty line and the Gini coefficient.

Table 8. GDP growth and changes in inequality: Descriptive statistics

Source	<u>Standard Deviation</u>		Correlation between Change in Inequality and GDP Growth
	Change in Inequality	GDP Growth	
Dollar and Kraay (2002) ^{a/}	0.054	0.037	-0.020
Povmonitor ^{b/}	0.049	0.063	0.070

Notes:

(^{a/}) Based on PWT national accounts.

(^{b/}) Based on survey data (<http://www.worldbank.org/povmonitor>).

Table 9. Share of variance in poverty changes due to growth ^{a/}
(Based on Dollar and Kraay (2002) database)

		<u>Headcount</u>			
		<u>Gini Coefficient</u>			
Mean Income /	Poverty Line	0.30	0.40	0.50	0.60
6		0.28	0.24	0.19	0.14
3		0.48	0.41	0.32	0.23
2		0.66	0.58	0.47	0.33
1.5		0.82	0.73	0.60	0.44
1		0.98	0.94	0.84	0.66

		<u>Poverty Gap</u>			
		<u>Gini Coefficient</u>			
Mean Income /	Poverty Line	0.30	0.40	0.50	0.60
6		0.25	0.20	0.15	0.11
3		0.41	0.32	0.24	0.16
2		0.55	0.43	0.31	0.20
1.5		0.67	0.53	0.38	0.24
1		0.83	0.68	0.50	0.32

		<u>Squared Poverty Gap</u>			
		<u>Gini Coefficient</u>			
Mean Income /	Poverty Line	0.30	0.40	0.50	0.60
6		0.23	0.18	0.13	0.09
3		0.37	0.28	0.20	0.13
2		0.48	0.36	0.25	0.16
1.5		0.58	0.44	0.30	0.19
1		0.73	0.56	0.39	0.24

Notes: The table reports the share of the overall variance of the poverty measures, computed under the assumption of lognormality, attributable to income growth as a function of the ratio of mean income/poverty line and the Gini coefficient.

(^{a/}) Calculated using the variances and covariance of growth and changes in inequality from the Dollar and Kraay (2002) database.

Table 10. Share of variance in poverty changes due to growth ^{a/}
(Based on Povmonitor database)

		<u>Headcount</u>			
		<u>Gini Coefficient</u>			
Mean Income /	Poverty Line	0.30	0.40	0.50	0.60
6		0.12	0.10	0.07	0.05
3		0.27	0.22	0.16	0.10
2		0.47	0.37	0.27	0.16
1.5		0.68	0.56	0.40	0.24
1		0.98	0.90	0.73	0.46

		<u>Poverty Gap</u>			
		<u>Gini Coefficient</u>			
Mean Income /	Poverty Line	0.30	0.40	0.50	0.60
6		0.11	0.08	0.06	0.03
3		0.22	0.15	0.10	0.06
2		0.34	0.23	0.14	0.08
1.5		0.47	0.32	0.19	0.10
1		0.70	0.49	0.29	0.15

		<u>Squared Poverty Gap</u>			
		<u>Gini Coefficient</u>			
Mean Income /	Poverty Line	0.30	0.40	0.50	0.60
6		0.10	0.07	0.05	0.03
3		0.19	0.13	0.08	0.04
2		0.28	0.18	0.11	0.06
1.5		0.37	0.24	0.14	0.07
1		0.55	0.35	0.20	0.10

Notes: The table reports the share of the overall variance of the poverty measures, computed under the assumption of lognormality, attributable to income growth as a function of the ratio of mean income/poverty line and the Gini coefficient.

(^{a/}) Calculated using the variances and covariance of growth and changes in inequality from the Povmonitor database (<http://www.worldbank.org/povmonitor>).

Figure 1. Empirical and Theoretical Quintiles

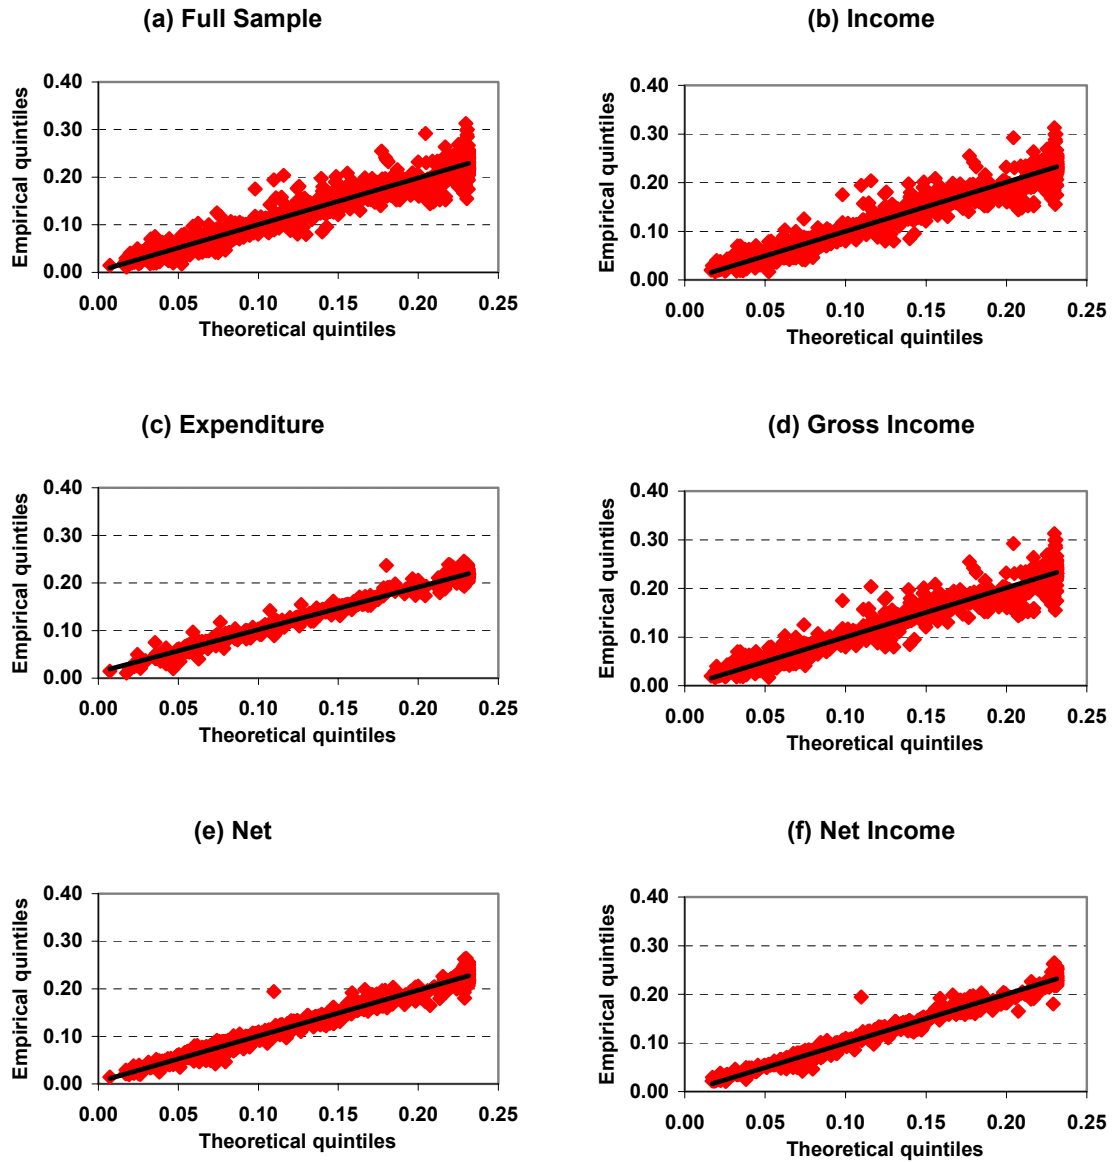


Figure 2. Iso-Poverty Curves for the Headcount Ratio

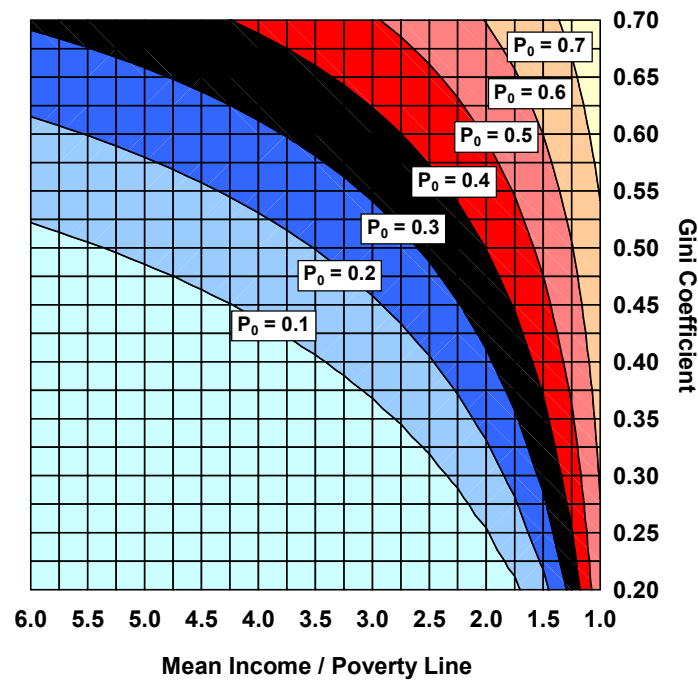
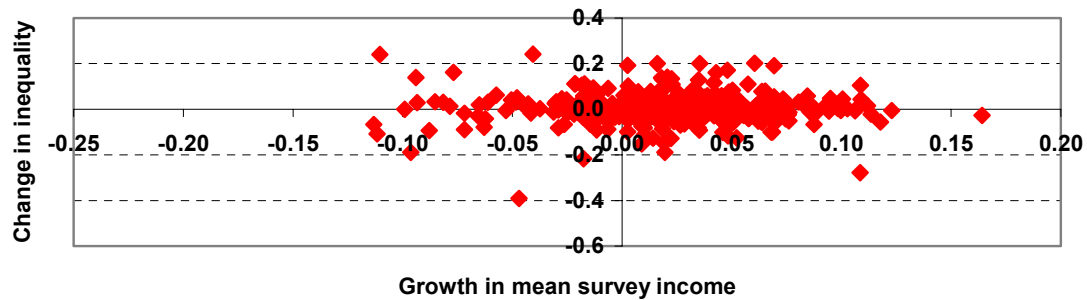
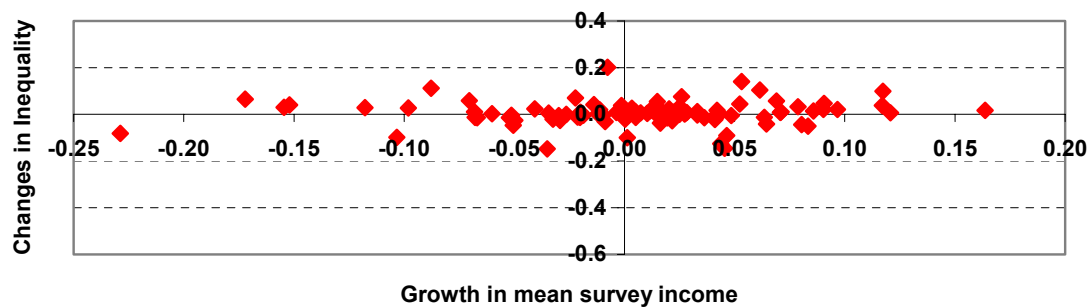


Figure 3. Income Growth and Changes in Inequality

(a) Dollar and Kraay (2002) Database



(b) Povmonitor^{a/} Database



Note: ^(a/) <http://www.worldbank.org/povmonitor>.

Appendix

In this appendix we derive the growth and inequality elasticities of the FGT family of poverty measures, given by

$$P_\alpha = \int_0^z \left[\frac{z-x}{z} \right]^\alpha f(x) dx \quad (\text{A1})$$

where $\alpha \in \{0,1,2\}$. When $\log x$ is normally distributed with mean μ and variance σ^2 we can denote $f(x) dx$ by $d\Lambda(x/\mu, \sigma^2)$, so that we can express the different FGT measures as:

$$P_0 = \int_0^z d\Lambda(x/\mu, \sigma^2) \quad (\text{A2})$$

$$P_1 = \int_0^z \left[\frac{z-x}{z} \right] d\Lambda(x/\mu, \sigma^2) = \int_0^z d\Lambda(x/\mu, \sigma^2) - \int_0^z \frac{x}{z} d\Lambda(x/\mu, \sigma^2) = \quad (\text{A3})$$

$$P_2 = \int_0^z \left[\frac{z-x}{z} \right]^2 d\Lambda(x/\mu, \sigma^2) = \quad (\text{A4})$$

$$\int_0^z d\Lambda(x/\mu, \sigma^2) - 2 \int_0^z \frac{x}{z} d\Lambda(x/\mu, \sigma^2) + \int_0^z \left[\frac{x}{z} \right]^2 d\Lambda(x/\mu, \sigma^2).$$

In order to derive the growth and inequality elasticities of the FGT family under lognormality we make use of the following result:

$$\int_0^z x^j d\Lambda(x/\mu, \sigma) = e^{j\mu + \frac{1}{2}j^2\sigma^2} \int_0^z d\Lambda(x/\mu + j\sigma^2, \sigma) dx, \quad (\text{A5})$$

which follows from Theorem 2.6 in Aitchison and Brown (1966). Using (A5) we can express (A2-A4) respectively as:

$$P_0 = \int_0^z d\Lambda(x/\mu, \sigma^2), \quad (\text{A6})$$

$$P_1 = \int_0^z d\Lambda(x/\mu, \sigma^2) - \frac{V}{z} \int_0^z d\Lambda(x/\mu + \sigma^2, \sigma^2), \quad (\text{A7})$$

$$P_2 = \int_0^z d\Lambda(x/\mu, \sigma^2) - \frac{2V}{z} \int_0^z d\Lambda(x/\mu + \sigma^2, \sigma^2) + \left(\frac{V}{z} \right)^2 e^{\sigma^2} \int_0^z d\Lambda(x/\mu + 2\sigma^2, \sigma^2). \quad (\text{A8})$$

Combining these expressions with the relationship linking the normal and lognormal distributions

$$\int_0^z d\Lambda(x/\mu, \sigma^2) = P(x < z) = P(\log x < \log z) = \Phi((\log z - \mu)/\sigma),$$

where $\Phi(\cdot)$ is the standard normal cumulative density function; and using also the identity linking average per capita income v to the mean and variance of log income, $\log v = \mu + \sigma^2/2$, (A6) -- (A8) can be rewritten as

$$P_0 = \Phi\left(\frac{\log(z/v)}{\sigma} + \frac{\sigma}{2}\right), \quad (\text{A9})$$

$$P_1 = \Phi\left(\frac{\log(z/v)}{\sigma} + \frac{\sigma}{2}\right) - \frac{v}{z} \Phi\left(\frac{\log(z/v)}{\sigma} - \frac{\sigma}{2}\right), \quad (\text{A10})$$

$$P_2 = \Phi\left(\frac{\log(z/v)}{\sigma} + \frac{\sigma}{2}\right) - 2\frac{v}{z} \Phi\left(\frac{\log(z/v)}{\sigma} - \frac{\sigma}{2}\right) + \left(\frac{v}{z}\right)^2 e^{\sigma^2} \Phi\left(\frac{\log(z/v)}{\sigma} - \frac{3\sigma}{2}\right). \quad (\text{A11})$$

which shows that $P_\alpha = P_\alpha(z/v, \sigma)$, and inverting equation (1) in the text we can further write $P_\alpha = P_\alpha(z/v, G)$.

Growth elasticities

For P_0 , the growth elasticity can be found in Bourguignon (2003). It is given by

$$\eta_v^0 = \frac{1}{d \log(v)} \frac{dP_0}{P_0} = -\frac{1}{\sigma} \lambda\left[\frac{\log(z/v)}{\sigma} + \frac{1}{2}\sigma\right] \quad (\text{A12})$$

where $\lambda(\cdot) = \phi(\cdot)/\Phi(\cdot)$, and $\phi(\cdot)$ denotes the standard normal density. In turn, for $\alpha \in \{1, 2\}$ we can use a result by Kakwani (1990):

$$\eta_v^\alpha = \frac{\alpha(P_\alpha - P_{\alpha-1})}{P_\alpha}. \quad (\text{A13})$$

Inequality elasticities.

For P_0 , the elasticity of poverty with respect to the standard deviation of log per capita income is given by:

$$\frac{\sigma}{P_0} \frac{\partial P_0}{\partial \sigma} = \lambda\left[\frac{\log(z/v)}{\sigma} + \frac{1}{2}\sigma\right] \left[\left(\frac{1}{2}\sigma - \frac{\log(z/v)}{\sigma}\right)\right] \quad (\text{A14})$$

In turn, from (1) in the text

$$dG = \sqrt{2} \phi\left(\frac{\sigma}{\sqrt{2}}\right) d\sigma \quad (\text{A15})$$

and hence the Gini elasticity of poverty is given by

$$\eta_G^0 = \frac{G}{P_0} \frac{\partial P_0}{\partial G} = \lambda \left[\frac{\log(z/\nu)}{\sigma} + \frac{1}{2} \sigma \right] \left[\left(\frac{1}{2} \sigma - \frac{\log(z/\nu)}{\sigma} \right) \right] / [\sqrt{2} \sigma \phi(\sigma/\sqrt{2}) / G] \quad (\text{A16})$$

The sign of this expression depends on that of $\frac{\sigma^2}{2} - \log(z/\nu)$; it is positive for $\nu > z$ -- i.e., when mean income is above the poverty line -- but becomes negative as $(\nu/z) \rightarrow 0$.

In order to derive the expressions for the Gini elasticity of P_1 and P_2 we define:

$$a \equiv \frac{\log(z/\nu)}{\sigma} + \frac{\sigma}{2} \quad (\text{A17})$$

$$b \equiv a - \sigma = \frac{\log(z/\nu)}{\sigma} - \frac{\sigma}{2} \quad (\text{A18})$$

$$c \equiv a + \sigma = \frac{\log(z/\nu)}{\sigma} + \frac{3\sigma}{2}. \quad (\text{A19})$$

First, we consider the elasticities of P_1 and P_2 with respect to the standard deviation of log per capita income, which follow from (A10) and (A11):

$$\frac{\sigma}{P_1} \frac{\partial P_1}{\partial \sigma} = \frac{\phi(a)(-b) - \frac{\nu}{z} \phi(b)(-a)}{P_1} \quad (\text{A20})$$

$$\frac{\sigma}{P_2} \frac{\partial P_2}{\partial \sigma} = \frac{\phi(a)(-b) - 2 \frac{\nu}{z} \phi(b)(-a) + \left(\frac{\nu}{z} \right)^2 e^{\sigma^2} (2\sigma^2 \Phi(c) + \phi(c) \left(\frac{-\log(z/\nu)}{\sigma} - \frac{3}{2} \sigma \right))}{P_2}. \quad (\text{A21})$$

Using again (A15) we finally get:

$$\eta_G^1 = \left(\frac{\phi(a)(-b) - \frac{\nu}{z} \phi(b)(-a)}{P_1} \right) / \left(\sqrt{2} \phi\left(\frac{\sigma}{\sqrt{2}}\right) \frac{\sigma}{G} \right), \quad (\text{A22})$$

and

$$\eta_G^2 = \left(\frac{\phi(a)(-b) - 2 \frac{\nu}{z} \phi(b)(-a) + \left(\frac{\nu}{z} \right)^2 e^{\sigma^2} (2\sigma^2 \Phi(c) + \phi(c) \left(\frac{-\log(z/\nu)}{\sigma} - 3\sigma \right))}{P_2} \right) / \left(\sqrt{2} \phi\left(\frac{\sigma}{\sqrt{2}}\right) \frac{\sigma}{G} \right). \quad (\text{A23})$$