Is Economic Growth Good for the Poor?

Tracking Low Incomes Using General Means

JAMES E. FOSTER and MIGUEL SZÉKELY

James E. Foster is an economics professor at Vanderbilt University, Nashville, USA. Miguel Szekely is with the Research Department of the Inter-American Development Bank. This paper is to be delivered at the ***Asia and Pacific Forum on Poverty: Reforming Policies and Institutions for Poverty Reduction***, to be held at the Asian Development Bank, Manila, 5-9 February 2001.

I. Introduction

Several recent papers have revisited the relationship between economic growth and poverty.[1] Does economic growth tend to "raise all boats'' including the conditions of the poor? Or is the main impact of economic expansion felt by the rich, with little if any benefits "trickling down" to the lower income groups? The answer has important implications for economic policy, since if the benefits of economic growth are already being shared across the various strata of an economy, departures from an unmitigated growth oriented policy need not be made in concession to distributional goals. However, if economic growth typically leaves the poor behind, pro-growth policies may have to be tempered by other considerations.

At issue here is whether economic growth, as measured by the rate of increase in per capita income, is associated with marked improvement in the condition of the poor. The latter outcome variable clearly has many dimensions. A proper evaluation would track each of the key attainments and capabilities of the poor and determine how they are altered during the growth process. But given data availability and the time lags with which capabilities are affected by economic conditions, it is not surprising that researchers have instead focused on the intermediate variable, namely, income. The question that is eventually brought to data is thus quite a bit narrower: How does economic growth affect the incomes of the poor?[2]

There are many ways of tracking the incomes of the poor. The traditional methods are based on the twin components of income poverty measurement as elucidated by Sen (1976), namely, the identification question (Who is poor?) and the aggregation question (Which function of incomes is to be used to track the condition of the poor?). Ravallion (2000), for example, employs absolute poverty standards of $1 and $2 a day to identify the poor and then aggregates using the most common measures of poverty, the headcount ratio and the per capita poverty gap. In contrast, Dollar and Kraay (2000) employ a purely relative definition of the poor as persons in the lowest fifth of the distribution, and then aggregates using the most common indicator of relative affluence, the per capita income of this group. Each approach is comprehensible and leads to results that inform the

1. See for example Ravallion and Chen (1997), Roemer and Gugerty (1997), Timmer (1997), Bruno et al. (1998), Gallup et al. (1999), Dollar and Kraay (2000), Ravallion (2000), Morely (2000), and De Janvry and Sadoulet (2000).

2. It is important, therefore, to be careful in the interpretation of results. Increased income is but one goal of development policies. See Sen (1999).

discussion. However, several methodological difficulties can be noted. A purely *absolute* poverty line (say, of $2 per day) marginalizes poverty in richer countries, and lessens the relevance of the findings across a broad spectrum of countries. A thoroughgoing *relative* poverty threshold (say, at the 20th percentile) can hardly be justified as a coherent line of separation between poor and nonpoor. In richer countries, the lowest 20 percent likely include many persons in the middle class, and hence some of the observed growth in poor incomes is actually growth in middle incomes. In poorer countries, the majority of poor persons may well be excluded, resulting in an estimate based on partial data. An alternative is to employ the country's own poverty standard in identifying the poor; but this introduces country-specific, idiosyncratic elements into the choice of the poverty line which in turn can lead to suspicious cross-country results. Clearly, none of these methods of identifying the poor is entirely above reproach in the demanding environment of cross-country evaluations over time.

Even if there were a thoroughly acceptable methodology for setting poverty lines in this context, there would still be significant questions about the use of an abrupt 0-1 cutoff. Suppose that an income is part of the evidence employed in evaluating the effect of growth on poor incomes. Why should an income slightly higher be ignored, just because it is above the arbitrary cutoff that is being employed? There is always a certain arbitrariness in selecting a particular poverty standard.[3] This is true for the $2 per day standard (why not $2.1?) as well as the 20th percentile cutoff (why not 21?), and an analogous argument likely holds for any given methodology. Further questions pertain to the aggregation methods typically used in these studies. For example, why should an income that is just below the poverty standard receive the same weight in the aggregation process as one that is much lower, as is implicit in the use of the headcount ratio, the poverty gap and the per capita income? A more defensible position might be to require progressively more weight to be placed on incomes further down the distribution

In this paper we propose the use of an alternative methodology to track low incomes based on Atkinson's (1970) family of "equally distributed equivalent income" functions, called "general means" here. Each of Atkinson's general means is an income standard that emphasizes the incomes of the

3. See the discussion in Foster and Shorrocks (1988), for example.

poor without ignoring the incomes of the near poor. Progressively less weight is placed on higher incomes. No arbitrary poverty standard is used. Rather, the curvature properties of a general mean ensure that higher incomes contribute very little to its value. In a sense, the presence of low incomes endogenously suppresses the impact of changes in higher incomes. The family of general means is indexed by a parameter that indicates the extent to which poorer incomes are emphasized in the income standard, or its “bottom sensitivity”.

Our method of evaluating the effects of growth on poor incomes is based on a comparison of growth rates for two standards of living: the ordinary mean; and a bottom sensitive general mean. The motivating question becomes: To what extent is growth in the ordinary mean accompanied by growth in the general mean? If growth were distributionally neutral, in that all incomes rise by the same proportion, then both standards would grow at the same rate. However, if the bulk of the increase in the mean takes place at the high end of the distribution, the growth rate in the general mean will lag behind the growth in the ordinary mean. Alternatively, if the general mean grows faster than the ordinary mean, this is a signal that growth disproportionately benefits the poor.

A key indicator in this approach is the growth elasticity of the general mean, or the percentage change in the general mean over the percentage change in the mean. Then proportional growth would lead to an elasticity of one, while pro-poor growth would be associated with an elasticity greater than one. If the elasticity is positive, but less than one, this indicates that although the growth favors the richer incomes, it also includes the poor to some extent. However, a nonpositive elasticity is a strong indicator of growth that does not benefit the poor.

We should mention that this approach yields additional dividends beyond the motivating question. Reinterpreting the general mean as a measure of social welfare as in Atkinson (1970), the question becomes the extent to which growth in per capita income is accompanied by growth in social welfare. So if the above elasticity is greater than one, this indicates growth that favors an even more rapid expansion of social welfare; while an elasticity less than one indicates that economic growth is somewhat less effective in generating an increase in social welfare. Hence, there is a useful normative interpretation of the methodology.

It is also possible to make statements about inequality in this framework, using Atkinson’s definition of inequality as unity minus the ratio of the general mean to the ordinary mean. Clearly if the data show that the elasticity is generally greater than one, the general mean rises faster than the ordinary mean, implying that the associated Atkinson inequality measure is falling. Hence, we have a method of evaluating the impact of growth on welfare and inequality, in addition to our original concern with low incomes.

Our empirical analysis estimates the growth elasticity of the general mean for a data set containing 144 household surveys from 20 countries over the last quarter century. Among other results, we find that the growth elasticity of bottom sensitive general means is positive, but significantly smaller than one. This suggests that the incomes of the poor do *not* grow one-for-one with increases in average income. The conclusion is robust to changing the composition of our sample of countries, to different estimation techniques, and to the inclusion of a set of control variables. Our conclusions differ from those in recent papers that use the per capita income of individuals in the first quintile as indicator of incomes of the poor and which argue that the growth elasticity of the incomes of the poor is equal to one. We confirm that the main reason why we arrive at a different result is because of the use of a different methodology to track the incomes of the poor.

The rest of the paper is organized as follows. Section II reviews previous empirical evidence on poverty and growth. Section III presents the general means and discusses their usefulness as measures of living standards. Section IV extends the discussion to the use of general means for tracking low incomes. Section V presents our empirical evidence, while Section VI concludes.

II. Previous Evidence on Poverty and Growth

Empirical evidence on whether the benefits of economic growth are shared by the poor started to be produced systematically around the 1970s, when compilations of income distribution statistics for several countries started to become available. The first papers on the subject focused on the relationship between growth and inequality since they were mainly concerned with verifying the Kuznets hypothesis that inequality increases during the initial phases of development, and declines after a turning point. The earlier papers were also specifically concerned with the effects of growth over the standard of living of the poor. For instance, Adelman and Morris (1973), Ahluwalia (1976) and Ahluwalia, Carter and Chenery (1979) asked whether there was a systematically relationship between economic growth and the income share of the bottom quintile. They concluded that this share tends to decline in the early stages of development, but increases in the long run. [4]

The growth–inequality relationship took center-stage during the 1980s, and only recently has there been renewed interest on the question of whether the poor specifically—rather than all sectors of society—share the benefits of growth proportionally. Recent papers follow two different approaches for classifying the population into poor and nonpoor. The first uses a relative concept of poverty by estimating the growth elasticity of the per capita income of individuals in the first quintile of the distribution.[5] There are two opposing views on the relation. While Roemer and Gugerty (1997), Gallup et al. (1999), and Dollar and Kraay (2000) argue that the elasticity is practically one, Timmer (1997) obtains an elasticity of around 0.8. Interestingly these four studies use the same data and similar econometric techniques, but disagree on whether growth in average income leads to a one-to-one increase in the incomes of the poor or to considerable smaller gains for this group.

The second approach has been to examine the growth elasticity of poverty defined in absolute terms. Ravallion (2000), Ravallion and Chen (1997), and Bruno et al. (1998) find that the elasticity of the headcount ratio is typically higher than 2, or in other words, that when average income increases by 10 percent, the proportion of poor declines by more than 20 percent. Other authors such as Morley (2000), De Janvry and Sadoulet (2000), and Smolensky et al. (1994) report a smaller elasticity of around one percent, but these are obtained from a smaller sample of countries. Ravallion and Chen (1997) also use poverty lines that combine an absolute and a relative component, but their elasticities are highly sensitive to where the poverty line is located. The elasticity of poverty to growth ranges from -2.59 to -0.69 depending on whether the threshold is

4. With the appearance of better data and the availability of improved econometric techniques, conclusions on the relation between inequality and economic growth have been repeatedly challenged. For instance, Anand and Kanbur (1993) argue that if the specification is improved, the inverted "U" shape relationship between inequality and growth vanishes. Bruno et al. (1998), Deininger and Squire (1998), Li et al. (1998), and Ravallion and Chen (1997) use an improved data set and argue that there is no systematic relation between the Gini inequality index and GDP per capita growth. But according to Barro (1999), inequality and growth do follow the inverted "U" shape relationship suggested by Kuznets. De Janvry and Sadoulet (2000) and Morley (2000) arrive to the same conclusion by using a data set that includes only Latin American countries. A recent paper by Lundberg and Squire (2000) argues that changes in GDP and in income inequality are jointly determined and should therefore be examined in a system of simultaneous equations where the direct relationship between these two variables is no longer of central interest.

5. The paper by Barro (1999) is one of the only to follow the earlier literature by estimating the growth elasticity of the income share of the poorest 20 percent.

established at 50 percent or 100 percent of the average income observed at the initial period of observation.

As compared to the literature on the growth–inequality relationship, the above approaches have the advantage that an intuitive interpretation can be given to the estimated growth elasticity. But as mentioned in the introduction, the advantage of a simple interpretation comes at the cost of having to specify the cut-off point after which changes in income are ignored. The need to define a threshold for dividing the population into poor and nonpoor introduces three issues into the analysis. The first is that poverty measures are highly sensitive to where the poverty line is set. For instance Chen and Ravallion (1997) report that when using a definition of poverty of one dollar-per-day PPP adjusted to 1993 prices, poverty in Latin America & the Caribbean is 15.57 percent, while if a relative poverty line is applied to the same data, the proportion is 51.35 percent. Székely et al. (2000) arrive at a similar conclusion: the proportion of poor in Latin America ranges from 22.8 to 56.8 percent, depending on which of the poverty lines that are commonly used in the region is adopted. Moreover, in countries where there is high income-concentration around the poverty line, even inframarginal variations in the value of the threshold may lead to large differences in poverty rates.[6]

The second issue is that the absolute and relative poverty measures used in the literature give exactly the same weight to all the poor. For instance, household survey data for Argentina 1998 reveal that the highest income among the poorest 20 percent of the population is $90 per month PPP at 1985 prices, while the average income among the poorest 3 percent is less than $7 PPP. Should a marginal increase in income for the second individual have the same value as a marginal gain for the first? If the interest is on whether relatively poorer individuals gain more from growth, then the answer is clearly no.

The third issue refers to the meaning of being poor. If a relative poverty line—such as the lower 20 percent of the distribution—is adopted then all individuals in the first quintile will be classified as poor, regardless of their absolute standard of living level. This implies giving the same weight to an individual in the lower 20 percent of the distribution in Sweden, who has an income of $450 PPP per

6. For one example take the case of El Salvador in 1997. If the poverty line is defined as 1985 PPP 2-dollars-a-day, the highest income among the poor turns out to be $1.99999, while the individual marginally above the poverty line has an income of $2.00073 per day, a difference of less than one thousandth of a cent. In fact, around 1 percent of the total population has an income within 3 cents of the value of the poverty line.

month in 1991, than to an individual in the poorest 20 percent in Kenya, with an income of barely $12.

III. General Means as Income Standards

Before describing our methodology, we present the general analytical framework and several key definitions. An *income distribution* is a vector of the form $x = (x_1,\ldots,x_n)$ where $x_i > 0$ is the income of the ith person. The population size n may vary across all positive integers. The set of all income distributions under consideration is given by the set $D = \cup_{n=1}^{\infty} R_{++}^{n}$. The income standard traditionally used in the evaluation of economic growth is the *per capita* or *mean* income $\mu = \mu(x) = (x_1 + \ldots + x_n)/n$, which is the aggregate income in x divided by the population size n of x. The class of *general means* is given by the formula $\mu_\alpha(x) = [(x_1^\alpha + \ldots + x_n^\alpha)/n]^{1/\alpha}$ for all $\alpha \neq 0$ and $\mu_\alpha(x) = (x_1 \cdots x_n)^{1/n}$ for $\alpha = 0$. Clearly, the general mean reduces to the standard mean when $\alpha = 1$. The case where $\alpha = 0$ is often called the *geometric mean* while $\alpha = -1$ is known as the *harmonic mean*. It is an easy matter to show that for fixed x, the general mean $\mu_\alpha(x)$ is increasing in the parameter α, with the limit as α falls to $-\infty$ being the minimum income in x, while the limit as α rises to ∞ is its maximum income. Each $\mu_\alpha(x)$ provides an alternative *income standard* or *representative income* for x, which places more weight on higher incomes for higher parameter values and more weight on lower incomes at lower parameter values. This paper will focus on general means that emphasize lower incomes, namely, $\mu_\alpha(x)$ for $\alpha < 1$.

One interesting characteristic of the general means is that for a given population size n, $\mu_\alpha(x)$ is strictly S-concave for $\alpha < 1$ and strictly S-convex for $\alpha > 1$.[7] Consequently, if distributions x and y share the same mean and population size, and if x is unambiguously more equal than y (in the Lorenz sense) then distribution x must have a higher general mean than distribution y, for every $\alpha < 1$, while over the range $\alpha > 1$ the inequality will be reversed. More equality "flattens" out the graph of the (increasing) function $\mu_\alpha(x)$ in parameter α, so that in the limit, where all incomes are equalized to $\mu(x)$, the graph becomes horizontal with all general means becoming equal.

7. See Foster and Shneyerov (2000) or Marshall and Olkin (1979, 54) for a discussion of the properties of general means.

The comparison between the value of $\mu_\alpha(x)$ for the United States (US), United Kingdom (UK), and Sweden illustrate this interpretation.[8] Figure 1 plots $\mu_\alpha(x)$ for α= –3, -2, -1, 1, 2 and 3 respectively, for each country.[9] The figure shows that the US has a considerably higher mean income than the other two countries (see μ_1). However, the fact that for α<1 the US ranks lower than Sweden, while for all α>1 it ranks much higher, reveals that the distribution of income is more unequal in the US, which is a well-known fact. Additionally, the figure shows that even though the average income in the US is higher (due to its higher μ_1), the incomes of individuals at the bottom of the US distribution is considerably lower than in the UK distribution and much lower than in Sweden.

As noted in Foster and Shneyerov (1999), there is a close link between the general means and decomposable inequality measures. Indeed, virtually every commonly used inequality measure (apart from the Gini) is a function of a ratio of two general means, or a limit of such functions.[10] Foster and Székely (2001) exploit this observation to derive new ways of evaluating inequality and growth by comparing levels of general means across different values of α, and by comparing the rates at which different general means grow. The approach is most easily illustrated for the Atkinson class of inequality measures, $A_\alpha = (\mu - \mu_\alpha)/\mu = 1 - \mu_\alpha/\mu$, for $\alpha < 1$, which takes inequality to be the gap between standard mean and the (smaller) general mean μ_α, normalized by μ. According to the Atkinson measure, inequality increases over time when the standard growth rate (of μ) exceeds the rate of growth of μ_α, or in words, poorer incomes grow less rapidly than average.

A similar conclusion holds for the generalized entropy measures over the range $\alpha < 1$, which are increasing transformations of Atkinson measures. For $\alpha > 1$, though, the general mean μ_α emphasizes higher incomes and takes higher values than the standard mean; the generalized entropy measures can be represented as positive transformations of $\mu_\alpha/\mu - 1 = (\mu_\alpha - \mu)/\mu$. It is clear, then, that for this parameter range inequality increases over time whenever the growth rate of μ_α is higher than the standard growth rate, i.e., higher incomes tend to grow more rapidly than average income.

8. Strictly speaking, inequality comparisons require proportional shifts in the graphs until all three intersect at the same level of μ_1.

9. The values are computed from household survey data for 1995 accessed through LIS. To make the values comparable across countries, household incomes were adjusted so that they equal PPP adjusted GDP per capita (taken from the World Bank *World Development Indicators* 2000).

10. This includes the generalized entropy measures, the Atkinson measures and the variance of logarithms.

The connection between changes in inequality and in μ_α are illustrated in Figure 2. Calculations from household survey data show that inequality increased substantially in Mexico between 1984 and 1996 (for example, the Gini index rose from .48 to .53). During the same years the mean income reported in the surveys increased by 50 percent in real terms. During a similar period of time (1985-1995), the Gini coefficient for Costa Rica declined from .47 to .45, while per capita incomes increased at practically the same rate as in Mexico. Figure 1 plots the change in the value of μ_α for parameter values of α=-3 to α=+3 for the two countries, and allows to identify the section of the income distribution where the changes in each country took place. As can be seen, while in Mexico μ_{-3}, μ_{-2}, and μ_{-1} decline substantially, in Costa Rica they grow much more than μ_1. The converse is the case for μ_3 and μ_2, while mean income (μ_1) increased at the same rate in both countries. Thus, in Mexico inequality increased because while there were substantial income losses among the poorest of the poor, the richest sectors of society registered large gains. In Costa Rica, the poor gained much more than the rich from the 50 percent increase in average income.

In his welfare-based approach to inequality, Atkinson (1970) introduced $\mu_\alpha(x)$ for $\alpha < 1$ as the *equivalent equally distributed income* associated with a given distribution x. This is the level of income which, if distributed equally, would yield the same level of social welfare as the original income distribution x, given a specific, symmetric, utilitarian social welfare function with decreasing marginal utility. As such, $\mu_\alpha(x)$ is clearly an increasing transformation of the original social welfare function having the additional property of being homogeneous of degree 1 (so that doubling all incomes doubles $\mu_\alpha(x)$). With this interpretation of $\mu_\alpha(x)$, we then see that Atkinson's measure of inequality is the shortfall of actual social welfare $\mu_\alpha(x)$ from the maximum social welfare achievable with the given total income (namely $\mu = \mu_\alpha(\mu_{,\ldots,}\ \mu)$), expressed as a percentage of maximum social welfare.

The welfare interpretation of general means for low values of α makes them valuable tools for normative evaluation of the income distribution. Our focus here, though, is on low incomes and how they are altered in the course of economic growth. In the next section we show how μ_α can be used to track the low incomes in a distribution.

IV. Tracking Low Incomes with General Means

Standard methods of evaluating poverty are difficult to apply in a consistent fashion across many countries at different levels of development. In particular, what is the right mix of relative and absolute standards? Even if a consistent general basis can be found for setting a poverty line for each country, the specific cut-off is still likely to be arbitrary. And the effect of this arbitrariness is amplified by the 0-1 approach to identifying the poor and aggregation methods that weight their incomes equally. The general means, on the other hand, represent an alternative approach to evaluating low incomes that avoids these difficulties. There is no arbitrarily set standard below which incomes are not counted. And incomes receive progressively less weight higher up the distribution.

Note, though, that each general mean is a function of all incomes, and this in turn might raise questions about its suitability in the present use. However, the seriousness of this concern depends entirely on the value of α. For values of α sufficiently close to 1, the general mean approximates the mean itself, which clearly places too much weight on high incomes for the present purpose. However, as α falls below 0, the curvature inherent in the functional form forces high incomes to be substantially muted. So, for example, the harmonic mean (α = -1) of the distribution (1, 2, 10) is 15/8, and no matter what level the highest income rises to, the harmonic mean of the distribution stays below 16/8. This insensitivity arises because the harmonic mean takes the average of the inverses of the incomes, which emphasizes the low incomes and essentially ignores the higher incomes (since their inverses are relatively small). This insensitivity is even more pronounced for lower values of α, with very small changes in low incomes having much larger effect than very large changes in middle and upper incomes. While these general means are a function of all incomes, and no external poverty standard is being imposed, their functional form ensures that they focus only on the lower incomes in the distribution.

To evaluate how growth affects the poor, we interpret μ_α as an income standard for the poor, and then estimate how rapidly μ_α changes when there is a change in the mean income. The resulting growth elasticity of the general mean then provides an answer to the question: To what extent do the poor share in economic growth? A negative elasticity indicates that the poor are negatively impacted by growth. An elasticity between 0 and 1 suggests that the poor are helped by economic growth, but at a rate that is somewhat less than the rest of the population. Elasticities beyond 1 imply that the poor benefit more than proportionately from economic growth. In the next section, we obtain estimate of growth elasticities for a dataset of household surveys from 20 countries over a quarter century.

V. Is Economic Growth Good for the Poor?

We now use the general means approach to ask whether growth has been good for the poor. Before discussing the empirical results, we briefly describe the data and estimation issues.

A. Data Description and Estimation Issues

Practically all the recent papers asking whether growth is good for the poor use the data set by Deininger and Squire (1996), which includes Gini coefficients and quintile shares for a large number of countries and years. This kind of aggregate data is not suitable for our analysis because to compute the general means it is necessary to have access to micro data in order to apply a weight to each individual in the distribution. Therefore, for this paper we construct our own data.

We have direct access to 144 household surveys from 20 countries ranging between 1976 and 1999, which we use to compute the general means. The sample includes 17 Latin American countries; Thailand; Taipei,China; and US, and the number of observations per country ranges from 2 for the Dominican Republic, Ecuador, Nicaragua, and Paraguay; to more than 11 data points for Brazil; Costa Rica; Taipei,China; and US. Appendix Table A1 gives more details of the household surveys included in our sample

Since our interest is on changes in welfare at the bottom of the distribution, we compute general means for each household survey for parameter values of $\alpha=0$, $\alpha=-1$, $\alpha=-2$, $\alpha=-3$, $\alpha=-4$, and then link each one of these measures to the growth in average income from the same survey. Rather than using the original incomes we adjust the data to make per capita incomes equal to PPP-adjusted GDP per capita so that our results are more comparable to other elasticities reported in the literature.[11]

The central equation of interest is therefore:

11. PPP GDP per capita figures between 1975 and 1992 are obtained from the World Penn Tables. To expand the series up to 1998 we use the rate of growth of real GDP in local currency from the *World Development Indicators 2000*. The GDP for 1999 was obtained from the ESDB at the Inter American Development Bank. Deaton (2000) has pointed out the various problems introduced by the use of these PPP conversion factors but we still use this methodology to make our results comparable with estimates published by other authors.

$$\log m_a(x)_{i,t} - \log m_a(x)_{i,t-1} = c + \log m_1(x)_{i,t} - \log m_1(x)_{i,t-1} + e_i$$

where *t* is the year in which a household survey for country *i* is available, and ✳ is an error term.

Having direct access to each of the 144 household surveys allows us to produce a data set with a high degree of comparability across observations, which minimizes measurement error in the dependent variable. [12] On one hand we are able to assure that the income concept is comparable *within* each country over time, while on the other, the lack of comparability *across* countries that inevitably remains becomes irrelevant when regressions are estimated in first differences, as we do below.[13] The high degree of comparability comes at the cost of having a reduced sample covering mostly developing countries from only one region in the world, but as we illustrate later, this characteristic of the sample does not seem to drive our conclusions.

The timing of the available data points to explore the relation between growth and poverty has been an issue in applied work, mainly because the Deininger-Squire data provides an unbalanced panel with observations scattered over several years. It has been standard practice to produce a reduced data set by spacing observations over 5 to 10 years and to use various estimation methods to impute information when an observation for a specific year is missing.[14] As shown in Appendix Table 1, household surveys for several countries in our sample belong to successive years or are only 2 or 3 years apart. Therefore, eliminating observations to produce a balanced panel would entail a significant loss in sample size, so our base estimates will refer to the full 144 observations. However, we show later that our conclusions are robust to eliminating information for successive

12. Although the Deiniger-Squire data is a major improvement over the information available to authors such as Ahluwalia, it still has some limitations. The most important is that the data is not comparable across or even within countries, so it is difficult to distinguish the noise-to-signal ratio it produces. This lack of comparability introduces measurement errors that are very hard to assess. Székely and Hilgert (1999), Atkinson and Brandolini (1999), and Pyatt (1999) provide a more thorough discussion of the problems of these kinds of secondary data sets. Most users of the Deininger-Squire data include some information on the characteristics of household surveys into regression specifications to reduce the comparability problems. However, these controls are not able to correct for fundamental differences such as the types of incomes captured by a survey, or the survey timing. Their importance is clearly illustrated in the recent paper by Panizza (2001), who shows that the relation between inequality and economic growth changes substantially when strictly comparable data is used.

13. When the income definition changes between surveys for a country, we construct a series with the minimum common denominator to ensure consistency. This entails some loss of information, but we believe that there are larger gains from reducing the noise-to-signal ratio of the series.

14. For instance, when an observation for certain year is missing, Ravallion and Chen (1997) use the distribution of the closest year available, and apply the average income of the target year to produce an estimate of poverty for that specific year. Another example is that when a country-year observation in the Deininger-Squire data has a Gini coefficient but no information on the quintile income shares, Dollar and Kraay (2000) estimate the quintile shares.

years and estimating the growth–poverty relationship with a reduced data set with observations every 3 years.

Regarding estimation techniques, the standard practice has been to estimate the poverty-growth elasticity in first differences, which eliminate the effect of time-invariant country characteristics. However, there are some discrepancies in the literature on how standard errors are corrected. Some authors acknowledge that successive spells within countries have one survey in common, and are therefore not independent observations (see especially Ravallion and Chen 1997). To produce our results we estimate our base regression in first differences, but in all cases we report Huber-White robust standard errors to address this issue.

Most of the poverty–growth elasticities reported in the literature do not acknowledge the potential endogeneity problem that arises from the fact that average income and measures of the standard of living of the poor are computed by using basically the same information. Dollar and Kraay (2000) deal with this by using instrumental variables and also address the problem arising with the inclusion of lagged endogenous variables.[15] Although our central results refer to standard first difference estimations, we also test whether our results are robust to the use of these techniques.[16]

B. Empirical Results

Table 1 presents our main results.**[17]** The table reports the value of the elasticity estimated through equation (1), from five separate regressions, corresponding to the use of a different general mean as dependent variable ('*t*' statistics are included under the coefficient). The results are quite striking since the lower the value of α, the smaller the elasticity. In other words, the greater the weight attached to the incomes of the poorest individuals, the smaller the gains from growth. $\mu_{0,}$ for instance, applies a slightly greater weight than μ_1 to lower income individuals, but the difference in weights between the poorest of the poor and individuals close to the mean is not very large. The elasticity of 1.08 suggests those individuals close to the middle of the distribution gain significantly

15. Specifically, Dollar and Kraay (2000) use the Arellano and Bover estimator, which is similar to GMM estimators but does not include fixed effects.

16. Lundberg and Squire (2000) use a variant of the forward differencing method to estimate a set of simultaneous equations that deal with the problem of reverse causality. We do not pursue this estimation here because it does not yield the growth elasticity that is of interest for this paper.

17. To perform the estimation we use the Huber iteration to reduce the effect of outlier observations.

more than one-to-one with growth in the mean. However, once greater weight is given to lower incomes, the elasticity becomes smaller. For μ_{-2}, μ_{-3}, and μ_{-4}, the elasticity is 0.77, 0.36 and 0.33, respectively, and in all cases, the coefficients are statistically insignificant. The conclusion is that living standards at the bottom of the distribution improve with growth, but that the poor gain proportionally much less than the average individual.

The conclusions from Table 1 are at odds with the recent papers by Roemer and Gugerty (1997), Gallup et al. (1999), and Dollar and Kraay (2000), which argue that the poor gain one-for-one from growth in mean income.[18] A straightforward question is whether the differences are due to the fact that those authors use the Deininger-Squire database, while our sample of countries and years is different. To explore this possibility we use our series of household surveys to compute the average income of individuals in the bottom 20, 10, and 30 percent of the distribution and estimate equation (1) by using each of these as dependent variables.[19]

Table 2 presents the results. The first line reports the growth elasticity of the mean income of individuals in the first quintile, which is the dependent variable used by the other authors. We obtain an estimate of 1.03, which is higher than the elasticity of 1.019 and .92 reported by Dollar and Kraay (2000), and Roemer and Gugerty (1997), respectively, but which is lower than the elasticity of 1.16 in Gallup et al. Some of these authors report more than one estimate, but we choose the ones that use the same methodology as in Table 1.

Interestingly, when the cutoff point is moved down, the elasticity declines. The second line in Table 2 reports the growth elasticity of the per capita income of individuals in the first decile. For every 10 percent increase in average income, the mean among the poorest 10 percent grows by 9.2 percent. When the cutoff is moved up to the 30th percentile, the elasticity is 1.06 (third line in the table). So, the lower the section of the distribution under examination, the smaller the gains from growth. The differences among percentiles 10, 20, and 30 are consistent with the results in Table 1 that when greater weight is given to lower incomes, the growth elasticity is smaller.

18. It must be said, however, that Timmer (1997) obtains a growth elasticity for the per capita income of individuals in the bottom 20 percent of the distribution that is significantly smaller than one, and which is similar to the elasticity we obtain for μ_{-2}. Interestingly, Timmer uses the same data and similar econometric techniques than the authors of these papers.

19. Household incomes are adjusted to match PPP GDP per capita.

Table 2 also includes the growth elasticity of the headcount ratio and the poverty gap index. We include these measures to determine whether the use of our database leads to the same conclusion than those obtained by Ravallion and Chen (1997) and Bruno et al. (1998) with respect to the effect of growth on these two poverty measures. [20] Our estimate of the growth elasticity of the headcount ratio and the poverty gap is of -1.49, and –2.09, respectively. Both coefficients are statistically significant. The elasticities are smaller than those in Ravallion and Chen (–3.12 and –3.69 for the headcount ratio and the poverty gap, respectively), but are of very similar magnitude than those obtained by Morley (2000) and De Janvry and Sadoulet (2000), who use a sample restricted to Latin American countries. These comparisons suggest that the head count ratio and the poverty gap are less responsive to growth in Latin America. But in any case, the conclusions derived from Table 2 are still in line with those of Ravallion and Chen.

Therefore, our data confirms previous results: (i) that the growth elasticity of per capita incomes in the first quintile is roughly equal to 1, and (ii) that the proportion of poor and the poverty gap decline significantly with growth. The use of a different data set and the inclusion of data from consecutive years does not explain why the general means lead to a different conclusion. The explanation is the difference in methodology.

C. Robustness Tests

This section performs a series of robustness test to check whether the conclusions from Table 1 change when we modify the country composition in our data set, when we apply different econometric techniques, and when control variables are introduced into equation (1).

Table 3 splits the data in different ways. The first column presents the growth elasticities for general means with $\alpha<1$, obtained by excluding the US from the sample. The elasticity of the general means in this sample of developing countries decline substantially as α becomes smaller, so the general conclusion that growth leads to less than proportional gains for the poorest individuals, stands. However, the coefficients for general means with $\alpha=-3$ and $\alpha=-4$ are substantially smaller

20. To compute the headcount ratio and the poverty gap we use a poverty line of $2-a-day PPP adjusted to 1985 prices, and as before, we blow up survey incomes to make them equal to PPP-adjusted GDP per capita so that they are comparable with the relative poverty measures and the results for the general means.

when the US is excluded. This suggests that the incomes of the poor are relatively less responsive to growth in the sample of developing countries.

The second column in Table 3 restricts the sample to the Latin American countries. In this case also, the general relation between growth and the incomes of the poor is consistent with the results in Table 1. The only difference is that $\mu_{0,}$ and μ_{-1} have a higher growth elasticity in these countries. The third column in the Table presents elasticities estimated with a sample that includes only episodes of positive growth in mean income. The coefficients are very similar to those in Table 1, which suggests that the poor suffer from income contractions in a similar way that they gain with expansion in the mean.

Table 4 tests whether our results hold when changing some elements of the estimation procedure. In the first column we present elasticities from a sample that drops information from consecutive years. The data is modified so that there are at least three years between each observation. Generating a data set with 5-year intervals as is common in the literature would imply dropping more than half of the sample. We chose 3-year episodes because this is the greatest interval by which observations can be separated, while still maintaining at least one half of the total number observations. The growth elasticities derived from the restricted sample lead to exactly the same conclusions that those in Table 1.

One important criticism to which several of the papers in this strand of literature have been subject, is that comparing the growth in incomes from a household survey to changes in GDP from National Accounts is similar to asking whether “growth in the People’s Republic of China has any effect over the incomes of the poor in India” (Deaton 2000). So far we have adjusted household incomes to GDP only to make our estimates more comparable with those in the literature, but to address this concern, we include a set of estimations that use the original household incomes (in real terms) to compute the general means as well as mean incomes. The second column in Table 4 presents the results. All elasticities are considerably smaller than our base estimates in Table 1, but the general conclusion is the same. The main difference is that even for μ_0 the growth elasticity is smaller than 1 (although not significantly smaller). This suggests that elasticities derived from data adjustments to match national account aggregates may overestimate the real relation between changes in average income and changes in the standard of living of the poor.

One of the concerns with specifications such as equation (1) is endogeneity. To address this issue, we follow Dollar and Kraay (2000) and instrument mean incomes by using the growth rate of the mean for the previous five years for each observation. The third column in Table 4 presents the results. Our central conclusions are not modified, and the value of each of the growth elasticities is very similar to those obtained by simply applying OLS to equation (1).

VI. Conclusion

We have argued in this paper that the general means are useful tools for illuminating the relationship between growth and poverty. Their importance for normative inequality analysis has been well-known since they were introduced as "equally distributed equivalent income" functions in Atkinson (1970). This paper shows that, in addition to their usefulness in evaluating social welfare and inequality of distributions, the general means can be employed to track low incomes for the purpose of evaluating how poor incomes are affected by growth.

We illustrate this method in an extensive empirical application involving household surveys from 20 countries over a quarter century. We replicate previous results that find a growth elasticity of about 1 for the income of the lowest 20 percent in these countries. We then find growth elasticities for the general means that are significantly below 1, suggesting that when the lowest incomes receive greater emphasis (as they do with the general means) then the effect of growth on the poor is not quite as strong as previously thought. This suggests a role for policies that take into account the distributional impact of growth.

Bibliography

Adelman, I. and Morris, C., 1973 *"Economic Growth and Social Equity in Developing Countries"*, Stanford University Press.

Ahluwalia, M., 1976 "Inequality, Poverty and Development", *Journal of Development Economics*, 3, pp. 307-42.

Ahluwalia, M., Carter. N. amd Chenery, H. 1979, "Growth and Poverty in Developing Countries", *Journal of Development Economics*, 6, pp. 299-341.

Anand, S. and Kanbur, R. 1985 "Poverty Under the Kuznets Process", *Economic Journal*, 95, pp. s42-s50.

Anand, S. and Kanbur, R. 1993, "The Kuznets Process and the Inequality-Development Relationship", *Journal of Development Economics*, 40, pp. 25-52.

Anand, S. and Kanbur, R. 1993a, "Inequality and Development: A Critique", *Journal of Development Economics*, 41, pp. 19-43.

Atkinson, A. B., 1970, "On the Measurement of Inequality," *Journal of Economic Theory*, 2, pp. 244-263.

Atkinson, A.B. and Brandolini A. "Promises and Pittfalls in the Use of Secondary Data Sets: Income Inequality in OECD Countries", Forthcoming, *Journal of Economic Literature*.

Barro, J.B. 1999, "Inequality and Growth in a Panel of Countries", Mimeo, Harvard University.

Bruno, M., Ravallion, M. and Squire, L., (1998) "Equity and Growth in Developing Countries: Old and New Perspectives on the Policy Issues". In *Income Distribution and High Quality Growth*, Tanzi, V., and Chu, KY, eds., MIT Press, Cambridge, Mass.

Chen, S., and Ravallion, M. 1997 "How did the world's poorest fare in the 1990s?", World Bank Economic Review.

De Janvry and Sadoulet 2000, "Growth, Poverty and Inequality in Latin America: A Causal Analysis, 1970-94", *Review of Income and Wealth*, Series 46 (3), pp. 267-87.

Deaton, A., 2000 "Counting the World's Poor: Problems and Possible Solutions", Mimeo, Princeton University.

Deininger, K. and Squire, L. 1996 "Measuring Inequality: A New Data Base", *World Bank Economic Review*, 10:3, 565-91.

Dollar, D. and Kraay, A. 2000 "Growth Is Good for the Poor" Mimeo, World Bank, Washington, D.C.

Foster, J.E. and Sen, A., 1997, "On Economic Inequality after a Quarter Century," in *On Economic Inequality* (A. Sen), expanded edition, Oxford University Press, Oxford.

Foster, J.E. and Shneyerov, 1999, "A General Class of Additively Decomposable Inequality Measures," *Economic Theory*, 14, pp. 89-111.

Foster, J.E. and Shneyerov, A.A. , 2000, "Path Independent Inequality Measures," *Journal of Economic Theory*, 91, pp. 199-222.

Foster, J.E., and Shorrocks, A.F., 1988, "Poverty Orderings," Econometrica

Foster, J.E., and Székely, M., 2001, "Using the General Means to Evaluate Inequality and Changes in Inequality," Mimeo, Research Department, Inter American Development Bank, Washington DC.

Gallup, J.L., Radelet, S. and Warner, A. 1999 "Economic Growth and the Income of the Poor" CAER II Discussion Paper No. 36, HIID, Harvard.

Li., H., L. Squire and F. Zou, 1998, "Explaining International and Intertemporal Variations in Income Inequality", *The Economic Journal*, Vol. 108, No. 446, pp. 26-43.

Lundberg, M. and Squire, L. 2000 "The Simultaneous Evolution of Growth and Inequality", Mimeo, World Bank, Washington DC.

Morley, S. 2000 "*La Distribución del Ingreso en America Latina y el Caribe*", Fondo de Cultura Economica.

Panizza, U., 2001 "Income Inequality and Economic Growth: Evidence From American Data", Mimeo, Research Department, Inter American Development Bank, Washington DC.

Pyatt, G. 1999 "The Distribution of Living Standards Within Countries: An Historical Perspective on a New Data Base" The Hague, Netherlands. Institute of Social Studies, Mimeo.

Ravallion, M 2000, "Growth and Poverty: Making Sense of the Current Debate", Mimeo, World Bank, Washington DC.

Ravallion., M. and Chen, S. 1997 "What Can New Survey Data Tell Us About Recent Changes in Distribution and Poverty?", ", *World Bank Economic Review*, 11:2, 357-82.

Roemer, M. and Gugerty, M. 1997 "Does Economic Growth Reduce Poverty?" CAER II Discussion Paper No. 4, HIID, Harvard.

Sen, A. 1999, *Development as Freedom*, Alfred A. Knopf, New York.

Smolensky, E., Plotnick, R., Evenhouse, E. and Reilly, S., 1994 "Growth, Inequality and Poverty: A Cautionary Note", *Review of Income and Wealth*, Series 40 (2), pp. 217-22

Székely, M. and M. Hilgert, 1999 "What's Behind the Inequality we Measure?: An Investigation Using Latin American Data", Working Paper *WP-409, Research Department, Inter American Development Bank, Washington DC.*

Székely, M., Lustig, N., Cumpa, M., and Mejía, J.A., 2000 "Do We Know How Much Poverty There is?", Working Paper *WP-437, Research Department, Inter American Development Bank, Washington DC.*

Timmer, P. 1997 "How Well Do the Poor Connect to the Growth Process? CAER II Discussion Paper No. 17, HIID, Harvard.

Figure 1

Comparison of Living Standards in the
USA, UK and Sweden

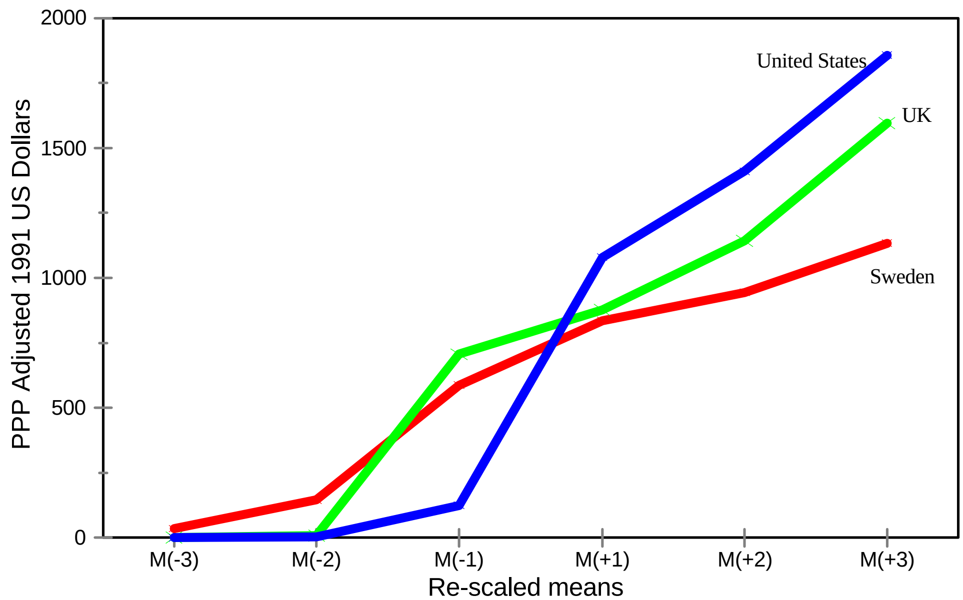

Figure 2

Change in Standard of Living in
Mexico and Costa Rica

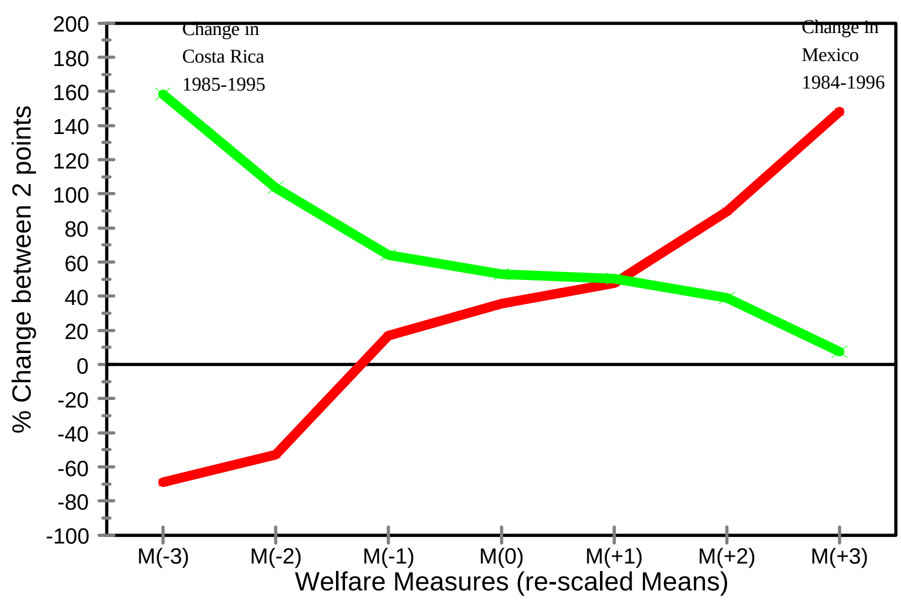

Table 1

Growth Elasticity of General Means
(Independent Variable is Growth in Mean Incor

Dependent Variable	Full Sample
General Mean with parameter = 0	**1.08** *8.11*
General Mean with parameter = -1	**0.93** *4.56*
General Mean with parameter = -2	**0.77** *1.58*
General Mean with parameter = -3	**0.36** *0.33*
General Mean with parameter = -4	**0.33** *0.22*
Number of Observations in Each Regression	123

Source: Authors' calculations.
*Each of the elasticities reported
is estimated from a separate regression.

Table 2

Growth Elasticity of Varios Welfare Measures
(Independent Variable is Growth in Mean Incon

Dependent Variable	Full Sample
Average Income Poorest Quintile	**1.03** *9.21*
Average Income Poorest Decile	**0.92** *7.34*
Average Income Poorest 30%	**1.06** *11.76*
Head Count Ratio	**-1.49** *-5.10*
Poverty Gap Index	**-2.09** *-5.28*
Number of Observations in Each Regression	123

Source: Authors' calculations.
*Each of the elasticities reported
is estimated from a separate regression.

Table 3

Growth Elasticity of General Means Using Different Samples
'(Independent Variable is Growth in Mean Income)

Dependent Variable	Sample Definition		
	Less Developed Countries	Latin America	Positive Growth Episodes
General Mean with parameter = 0	**1.07**	**1.14**	**1.08**
	5.95	*5.39*	*9.46*
General Mean with parameter = -1	**0.97**	**1.01**	**0.95**
	1.22	*3.77*	*2.79*
General Mean with parameter = -2	**0.70**	**0.66**	**0.63**
	1.40	*0.15*	*0.77*
General Mean with parameter = -3	**0.21**	**0.20**	**0.57**
	0.15	*0.45*	*0.57*
General Mean with parameter = -4	**0.17**	**0.21**	**0.35**
	0.60	*0.50*	*0.64*
Number of Observations in Each Regression	101	74	94

Source: Authors' calculations.
*Each of the elasticities reported
is estimated from a separate regression.

Table 4

Elasticity of General Means Using Different Estimation Techiques
'(Independent Variable is Growth in Mean Income)

Dependent Variable	Estimation Technique			
	3-Year Intervals Between Observations	Unadjusted Incomes	Instrumental Variables	Arellano-Bover Correction
General Mean with parameter = 0	**1.01**	**0.82**	**0.94**	
	6.17	*2.05*	*3.49*	
General Mean with parameter = -1	**0.88**	**0.59**	**0.85**	
	5.00	*1.13*	*2.27*	
General Mean with parameter = -2	**0.66**	**0.43**	**0.73**	
	1.27	*1.04*	*0.80*	
General Mean with parameter = -3	**0.18**	**0.22**	**0.33**	
	1.00	*1.02*	*0.52*	
General Mean with parameter = -4	**0.16**	**0.20**	**0.28**	
	0.85	*1.02*	*0.08*	
Number of Observations in Each Regression	123	123	123	

Source: Authors' calculations.
*Each of the elasticities reported
is estimated from a separate regression.

Appendix Table A

Household Surveys

Country	# Surveys	Years	Survey
Argentina	3	1980, 96,98	Encuesta Permanente de Hogares
Bolivia	7	1986	Encuesta Permanente de Hogares
		1990, 93, 95	Encuesta Integrada de Hogares
		1996, 97	Encuesta Nacional de Empleo
		1999	Encuesta Continua de Hogares (condiciones de vida)
Brazil	11	1981, 83, 86, 88	Pesquisa Nacional por Amostra de Domicilios
		1992, 93, 95, 96, 97,98,99	Pesquisa Nacional por Amostra de Domicilios
Chile	6	1987, 90, 92, 94, 96, 98	Encuesta de Caracterización Socioeconómica Nacional
Colombia	6	1991, 93, 95, 97, 98,99	Encuesta Nacional de Hogares - Fuerza de Trabajo
Costa Rica	10	1981, 83, 85	Encuesta Nacional de Hogares - Empleo y Desempleo
		1987, 89, 91, 93, 95, 97, 98	Encuesta de Hogares de Propósitos Múltiples
Dominican Republic	2	1996	Encuesta Nacional de Fuerza de Trabajo
		1998	Encuesta Nacional Sobre Gastos e Ingresos de los Hogares
Ecuador	2	1995, 98	Encuesta de Condiciones de Vida
El Salvador	3	1995, 97, 98	Encuesta de Hogares de Propósitos Múltiples
Guatemala	1	1998	Encuesta Naional de Ingresos y Gastos Familiares
Honduras	6	1989, 92, 96, 97, 98,99	Encuesta Permanente de Hogares de Propósitos Múltiples
Mexico	7	1977	Encuesta de Ingreso y Gasto de los Hogares
		1984, 89, 92, 94, 96,98	Encuesta Nacional de Ingreso y Gasto de los Hogares
Nicaragua	2	1993, 98	Encuesta Nacional de Hogares Sobre Medicion de Niveles de Vida
Panama	6	1979	Encuesta de Hogares - Mano de Obra (EMO)
		1991, 95, 97, 98,99	Encuesta Continua de Hogares
Paraguay	2	1995	Encuesta Nacional de Empleo
		1998	Encuesta Integrada de Hogares
Peru	5	1985, 91, 94, 97,2000	Encuesta Nacional de Hogares sobre Medición de Niveles de Vida
		1996	Encuesta Nacional de Hogares sobre Niveles de Vida y Pobreza
Uruguay	6	1981, 89	Encuesta Nacional de Hogares
		1992, 95, 97,98	Encuesta Continua de Hogares
Venezuela	8	1981, 86, 89, 93, 95, 97,98,99	Encuesta de Hogares por Muestra
United States	23	1976 - 1998	Current Population Survey
Thailand	8	1975,81,86,88,90,92,94,96	Socio - Economic Survey
Taiwan	21	1976 - 1996	Survey of Family Income and Expenditure