

Foster, James E.; Székely, Miguel

Working Paper

Is Economic Growth Good for the Poor? Tracking Low Incomes Using General Means

Working Paper, No. 453

Provided in Cooperation with:

Inter-American Development Bank (IDB), Washington, DC

Suggested Citation: Foster, James E.; Székely, Miguel (2001) : Is Economic Growth Good for the Poor? Tracking Low Incomes Using General Means, Working Paper, No. 453, Inter-American Development Bank, Research Department, Washington, DC

This Version is available at:

<https://hdl.handle.net/10419/88055>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



*Inter-American Development Bank
Banco Interamericano de Desarrollo (BID)
Research department
Departamento de investigación
Working Paper #453*

Is Economic Growth Good for the Poor? Tracking Low Incomes Using General Means

By

James e. Foster*
Miguel Székely**

***Vanderbilt University**

**** Inter-American Development Bank**

June 2001

**Cataloging-in-Publication data provided by the
Inter-American Development Bank
Felipe Herrera Library**

Foster, James E.

Is economic growth good for the poor? : tracking low incomes using general means /
James E. Foster, Miguel Székely.

p. cm. (Research Department Working paper series ; 453)
Includes bibliographical references.

1. Poor--Latin America. 2. Poor--Latin America--Effect of income on. I. Székely,
Miguel. II. Inter-American Development Bank. Research Dept. III. Title. IV. Series.

339.46 F359-dc21

82001

Inter-American Development Bank
1300 New York Avenue, N.W.
Washington, D.C. 20577

The views and interpretations in this document are those of the authors and should not be attributed to the Inter-American Development Bank, or to any individual acting on its behalf.

The Research Department (RES) produces the *Latin American Economic Policies Newsletter*, as well as working papers and books, on diverse economic issues.

To obtain a complete list of RES publications, and read or download them please visit our web site at: <http://www.iadb.org/res/32.htm>

Abstract*

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, which are called "general means" here. We provide a new characterization of general means that justifies their use in this context. 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 is: To what extent is growth in the ordinary mean accompanied by growth in the general mean? 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 usual mean. Our empirical analysis estimates this growth elasticity 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.

Keywords: Poverty, economic growth, welfare.

JEL codes: I3, O1, E6.

* Foster is a Professor of Economics at the Department of Economics at Vanderbilt University, and Székely is at the Research Department of the Inter-American Development Bank. We thank Ricardo Fuentes for excellent research assistance, as well as Santiago Levy, Sam Morley, Ugo Panizza, Graham Pyatt, Martin Ravallion, John Dunn Smith, John Waymar and seminar participants at the Asian Development Bank, the Inter-American Development Bank, and the Symposium on Poverty Measurement in Mexico for helpful comments. The views expressed here are the authors' and do not necessarily reflect those of the institutions with which the authors are affiliated.

1. Introduction

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 benefit “trickling down” to the lower income groups? The answers to these questions have 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 conditions of the poor. The latter outcome variable clearly has many dimensions besides income. 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 somewhat narrower: How does economic growth affect the incomes of the poor?¹

There are many ways of tracking the incomes of the poor empirically. The literature has focused mainly on two options, both of which 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?). The first approach employs a purely relative definition of the poor as all persons in the lowest quintile (or decile), and then aggregates poor incomes using the most common income standard, the mean. The central empirical question concerns the relationship between economic growth (in the economy-wide mean) and growth in the income standard of the poor, and whether the so-called “growth elasticity” of this income standard exceeds, equals or falls below unity. The earlier papers in this strand of literature, including Adelman and Morris (1973), Ahluwalia (1976) and Ahluwalia, Carter and Chenery (1979), were primarily interested in the growth-inequality relationship (with one inequality measure being the income *share* of the poor group), but they also asked whether the poorest 20 percent of the population shared the

¹ It is important, therefore, to be careful in the interpretation of results. Increased income is but one goal of development policies. See Sen (1999).

benefits of growth proportionally. They concluded that the income share of the poor tends to decline in the early stages of development but increases in the long run.² The approach has received renewed attention recently, and two different views on the magnitude of the growth elasticity can be found. Roemer and Gugerty (1997), Gallup, Radelet and Warner (1999) and Dollar and Kraay (2000) argue that the growth elasticity of the incomes of individuals at the bottom quintile is practically equal to one. Timmer (1997) obtains a more modest elasticity of around 0.8. Interestingly, these four studies use the same data and similar econometric techniques, but they disagree on whether growth in average income leads to a one-to-one (or proportional) increase in the incomes of the poor, or whether the gains for this group are considerably smaller.

The second approach tracks income poverty levels using an absolute poverty line and a standard poverty measure. Recent papers by Ravallion (2000), Ravallion and Chen (1997), and Bruno, Ravallion and Squire (1998) employ absolute poverty lines of \$1 and \$2 a day to identify the poor and then aggregate, using the most common measures of poverty, the headcount ratio and the per capita poverty gap. These studies find that the growth elasticity of the headcount ratio is typically below -2, or in other words that, when average income increases by 1 percent, the proportion of poor declines by more than 2 percent. Other authors, such as Morley (2000), De Janvry and Sadoulet (2000) and Smolensky, Plotnick, Evenhouse *et al.* (1994), report an elasticity of around -1, 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 growth elasticity of poverty ranges from -2.59 to -0.69, depending on whether the threshold is established at 50 percent or 100 percent of the average income observed at the initial period of observation.

² 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 (1993a) argue that if the specification is improved, the inverted “U”-shaped relationship between inequality and growth vanishes. Bruno, Ravallion and Squire (1998), Deininger and Squire (1996), Li, Squire and Zou (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 at 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.

Each of the two approaches is comprehensible and leads to results that inform the 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 non-poor. 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? Selecting a particular poverty standard is always arbitrary to some extent.³ This is true for the \$2 per day standard (why not \$2.1?) as well as the 20th percentile cutoff (why not the 21st?), 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 line 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 among the bottom fifth? A more defensible position might be to require progressively more weight to be placed on incomes further down the distribution.⁴

Finally, we observe that the specific income standard (the mean of the lowest fifth) employed in several of the studies to track poor incomes suffers from a form of inconsistency

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

⁴ A distribution-sensitive poverty measure, such as those proposed by Sen (1976) or Foster, Greer, and Thorbecke (1984), might go part of the way to addressing this concern. However, this awaits the establishment of a consistent framework for constructing poverty lines for cross-country evaluations.

that may seriously cloud its relevance for policy prescriptions. Specifically, an increase in this standard for a country as a whole is entirely consistent with a *decrease* in the standard for *every* region in the country. (An example of this is given in Section 2 below.) Clearly, it would be preferable to use a “subgroup-consistent” income standard that rules out the possibility of such contradictory evaluations.⁵

This paper makes two main contributions. First, we propose and justify an alternative methodology to track low incomes based on Atkinson’s (1970) family of “equally distributed equivalent income” functions, called “general means” here. Each is an income standard that emphasizes the incomes of the 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.” Each member of this family satisfies a basic collection of properties for income standards, as well as “subgroup consistency.” Moreover, as we demonstrate below, the general means are the *only subgroup-consistent income standards* satisfying the basic properties, which provides a compelling justification for their use.

The second contribution is empirical. We use the above methodology to estimate the growth elasticity of the general mean by accessing the micro data from 144 household surveys in 20 countries spanning over a quarter of a 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 an income standard of the poor and argue that the growth elasticity of this income standard is unity. We confirm that the main reason we obtain a different result is our methodology for tracking the incomes of the poor.

⁵ See, for example, the related discussion in Foster and Shorrocks (1991) and Foster and Sen (1997).

The rest of the paper is organized as follows. Section 2 discusses the usefulness of the general means as income standards and provides an axiomatic characterization in terms of subgroup consistency. Section 3 presents our empirical evidence, while Section 4 concludes.

2. General Means as Income Standards

Given the problem of setting a consistent poverty line for cross-country analysis, and the sensitivity of poverty levels (and growth elasticities of poverty levels) to the specific methodology employed, we will use the income standard approach rather than the poverty measure approach in our analysis. But the specific income standard commonly employed in related studies is also subject to criticism due to its reliance on an arbitrary cutoff and other conceptual difficulties. We therefore broaden consideration to other potential income standards. The motivating question of this section is: What functional form should be used to evaluate the incomes of the poor?

A. A Characterization of the General Means

We begin by presenting several definitions and a general framework for defining and evaluating income standards. An *income distribution* is a vector of the form $x = (x_1, \dots, x_n)$ where $x_i > 0$ is the income of the i th person and n is the population size, which is a positive integer. Denote by $D^n = \mathbb{R}_{++}^n$ the set of all n -person income distributions, and let $D = \bigcup_{n=1}^{\infty} D^n$ be the set of all income distributions (where population size n can vary across all positive integers). We say that x is *completely equal* if $x_i = x_j$ for all i, j . We say that x is a *permutation* of y if both are distributions having the same population size n , and $x = Py$ for some permutation matrix P of order n ; in other words, x has the same collection of incomes as y , but potentially in a different order. We say that x is a *replication* of y if there are integers $n \geq 1$ and $m \geq 2$ such that x is in D^{nm} and y is in D^n with $x = (y, y, \dots, y)$; in other words, every income in y has m “clones” in x . The *per capita* or *mean* income will be denoted by $\mu = \mu(x) = (x_1 + \dots + x_n)/n$.

We are interested in finding an appropriate income standard $f: D \rightarrow \mathbb{R}$ to track the incomes of the poor. Consider the following natural properties for f , where x and y are distributions in D :

Symmetry: If x is a permutation of y , then $f(x) = f(y)$.

Replication Invariance: If x is a replication of y , then $f(x) = f(y)$.

Homogeneity: If $x = ky$ for some scalar $k > 0$, then $f(x) = kf(y)$.

Normalization: If x is completely equal, then $f(x) = x_1$.

Continuity: f is continuous on each D^n .

Symmetry ensures that all incomes are treated symmetrically by the income standard; replication invariance makes the income standard coherent across population sizes; homogeneity requires that if all incomes are doubled, the income standard must double as well; normalization specifies that the income standard of a completely equal distribution is simply the common level of income; continuity ensures that the income standard does not abruptly change as incomes are altered. Notice that the usual mean income μ satisfies each of these natural properties, as does the mean income of the poorest fifth of the population.

A final property ensures consistency between the aggregate level of the income standard and the levels in population subgroups. Consider the following, where x , y , x' , and y' are distributions in D :

Subgroup Consistency: Suppose that $f(x') > f(x)$ and $f(y') = f(y)$, where x' has the same population size as x , and y' has the same population size as y . Then $f(x', y') > f(x, y)$.

In other words, if the income distributions in two population subgroups change in such a way that the income standard rises in one and is unchanged in the other, then the overall standard must rise. Repeated application of the property ensures that the same conclusion obtains when both subgroup standards rise, or when there are multiple subgroups. Surprisingly, the mean income of the bottom fifth—the income standard employed in many studies, including Dollar and Kraay (2000)—does *not* satisfy subgroup consistency. This is immediately seen with the help of a numerical example. Suppose that each of the distributions x , y , x' and y' has ten incomes, with the lowest three incomes being (4, 8, 12), (2, 6, 8), (2, 11, 12), and (3, 6, 8), respectively. Then, the mean of the lowest fifth (or lowest two incomes) is 6 for x and 4 for y , with a level of 5 for the combined distribution (x, y) . However, while the income standards rise to 6.5 for x' and 4.5 for y' , the overall standard *falls* to 4.75 for the combined distribution (x', y') .

This problem arises because of the endogenous nature of the set of the poor implicit in this standard of living. Recall that in the above example, the second lowest income in x rises significantly enough to compensate for the decline in the lowest income. However, this increase goes unnoticed in the combined distribution since the rising income is elevated outside the set of poor incomes, allowing the remaining decrement to dominate.

One family of income standards satisfying all of these basic properties is the class $\mu_\alpha(x)$ of *general means*, defined by $\mu_\alpha(x) = [(x_1^\alpha + \dots + x_n^\alpha)/n]^{1/\alpha}$ for all $\alpha \neq 0$ and by $\mu_\alpha(x) = (x_1 \cdots x_n)^{1/n}$ 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*. That the general means satisfy all six properties is immediately apparent. What is not so obvious is that they are the *only* income standards to do so.

Theorem: An income standard satisfies symmetry, replication invariance, homogeneity, normalization, continuity and subgroup consistency if and only if it is a generalized mean.

Proof: In the Mathematical Appendix.

Any income standard, whether it emphasizes the lower or the upper end of the distribution, can be expected to satisfy the basic requirements of symmetry, replication invariance, homogeneity, normalization and continuity. Subgroup consistency is a compelling additional property that ensures sufficient coherence between the overall and subgroup levels of the income standard. The above theorem identifies the single-parameter family of the general means as the only income standards satisfying all the requirements simultaneously.⁶ It is a powerful axiomatic justification for focusing exclusively on income standards drawn from this class.

B. Evaluating Income Distributions Using General Means

While the general means have been indirectly employed in the evaluation of income distributions since Atkinson (1970), their independent usefulness as income standards has not received much

⁶ Other characterizations of the general means can be found in the comprehensive survey of Diewert (1993).

attention.⁷ We now briefly explore them in greater detail and provide examples of their empirical usefulness.

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. 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 $\alpha > 1$.⁸ 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 the 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.

A comparison between the values of $\mu_\alpha(x)$ for the United States, the United Kingdom and Sweden illustrates this interpretation. Figure 1 plots $\mu_\alpha(x)$ for $\alpha = -3, -2, -1, 0, 1, 2$ and 3 , respectively, for each country.⁹ The figure shows that the United States has a considerably higher mean income than the other two countries (see μ_1). However, the fact that for $\alpha < 1$ the United States ranks lower than Sweden, while for all $\alpha > 1$ it ranks much higher, reveals that the distribution of income is more unequal in the United States, which is a well known fact.¹⁰ Additionally, the figure shows that even though the average income in the US is higher, the incomes of individuals at the bottom of the US distribution are 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

⁷ The paper of Blackorby, Donaldson and Auersberg (1981) is one notable exception. See also Anand and Sen (1996).

⁸ See, Foster and Shneyerov (1999) or Marshall and Olkin (1979, p. 54).

⁹ The values are computed from household survey data for 1995 accessed through the Luxemburg Income Study. 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's *World Development Indicators*, 2000).

¹⁰ Strictly speaking, inequality comparisons require proportional shifts in the graphs until they intersect at the same level of μ_1 .

measure (apart from the Gini) is a function of a ratio of two general means, or a limit of such functions.¹¹ 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, which can be defined as $A_\alpha = (\mu - \mu_\alpha)/\mu = 1 - \mu_\alpha/\mu$ for $\alpha < 1$, so that inequality is viewed as the gap between the 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 other words, poorer incomes grow less rapidly than the average income.

A similar conclusion holds for the generalized entropy measures I_α , which for $\alpha < 1$ are increasing transformations of the 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 I_α 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 the average income.

The connection between changes in inequality and in μ_α are illustrated in Figure 2. Calculations from household survey data show that mean incomes in Mexico and Costa Rica increased by virtually the same proportion during 1984-1996 and 1985-1995, respectively (see μ_1 at the center of the figure). However, there was a substantial decline in μ_{-3} , μ_{-2} and μ_{-1} in Mexico, while general means for $\alpha < 1$ grew much more than μ_1 in Costa Rica. The converse is true for μ_3 and μ_2 . Thus, in Mexico inequalities increased due to faster growth among higher-income strata, while in Costa Rica the poor gained much more than the average individual and inequality declined. The value of the Theil Entropy measure actually increased from 0.43 to 0.55 in Mexico during this period and decreased from 0.42 to 0.39 in Costa Rica.

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

¹¹ This includes the generalized entropy measures, the Atkinson measures and the variance of logarithms. See Foster and Sen (1997) for formal definitions of these measures.

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 all the properties of an income standard as presented above. With this interpretation of $\mu_\alpha(x)$, we 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 $\mu = \mu_\alpha(\mu, \dots, \mu)$), 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.

C. Tracking Low Incomes Using General Means

Low incomes, and how they are altered in the course of economic growth, are the focus of the present analysis. Our analysis therefore restricts consideration to general means with parameter values below 1, and we interpret $\mu_\alpha(x)$ as an income standard of the poor. Note, though, that since general means are functions of all incomes, this might raise questions about their suitability for tracking the incomes of the poor. 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. Even the geometric mean can be a bit too sensitive to upper incomes, as can be seen by comparing the levels of μ_0 for (1, 2, 10), and (1, 2, 100). 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 ($\alpha = -1$) 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). The effect is even more pronounced for lower values of α , with very small changes in low incomes having much larger impact on the standard than very large changes in middle and upper incomes. While general means are a function of all incomes, and no external poverty standard is being imposed, their functional form ensures that they focus on the lower incomes of the distribution.

In the next section we use general means to explore how growth affects the incomes of the poor. In particular, we evaluate how rapidly the income standard μ_α for $\alpha < 1$ changes when

there is a change in the mean income. The result provides an answer to the question: To what extent do the poor share in economic growth? 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 benefits the poor more than proportionally. The 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. 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 growth favors the richer incomes, it also includes the poor to some extent. However, a non-positive elasticity is a strong indicator of growth that does not benefit the poor.

3. Is Economic Growth Good for the Poor?

Our methodology is now used to estimate the growth elasticity of the incomes of the poor. Before discussing the 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 the underlying micro data in household surveys in order to apply a weight to each individual in the distribution. Furthermore, within country and cross-country comparability in the Deininger-Squire data is not guaranteed. Therefore, for this paper we construct our own database by directly accessing household survey records for as many countries as possible.

We impose four conditions for including a household survey in our analysis. First, the household survey has to be nationally representative. The only exceptions we make are Argentina and Uruguay, where household surveys are restricted to urban areas but still include more than 80% and 90% of the country's population, respectively. Nationally representative surveys are not available for these countries. Second, the survey questionnaire has to include a

breakdown of income by source, with at least three separate questions on income that identify labor income, profits, and capital rents separately. This is to assure lower measurement error in incomes. Third, the recall period for incomes has to be the same (the previous month) in each survey.¹² Fourth, the central purpose of the survey must be to collect information on the standard of living of the population. This last requirement assures us that obtaining accurate information on incomes is an objective of the survey.

We are able to access 144 household surveys fulfilling these requirements. The surveys cover 20 countries—17 from Latin America, 2 from Asia, and the United States—ranging between 1976 and 1999. Each of these surveys is used to compute the general means. The number of surveys per country ranges from 2 for the Dominican Republic, Ecuador, Nicaragua and Paraguay to more than 11 data points for Brazil, Costa Rica, Taiwan, and the United States. Appendix Table A1 provides additional details on the household surveys included in our sample. All in all, the 144 surveys include 4.1 million household records, corresponding to 15.2 million individuals. The average sample size across surveys is 26,435 and 100,576 households and individuals, respectively.

Since our interest is in changes in incomes at the bottom of the distribution, we compute general means for each household survey for parameter values of $\alpha = -1$, $\alpha = -2$, $\alpha = -3$, $\alpha = -4$, and then link each one of these measures to the growth in average income from the same survey. For illustration purposes we also include results for the geometric mean ($\alpha = 0$), although, as we have seen above, this indicator is not particularly bottom-sensitive. In our implementation we adjust all survey incomes to make the aggregate per capita income equal to PPP-adjusted GDP per capita from the National Accounts. Consequently, our results are more comparable to other elasticities reported in the literature.¹³

The central equation of interest is therefore:

$$(1) \quad \log \mathbf{m}_a(x)_{i,t} - \log \mathbf{m}_a(x)_{i,t-1} = c + \log \mathbf{m}_1(x)_{i,t} - \log \mathbf{m}_1(x)_{i,t-1} + \mathbf{e}_i$$

where t is the year in which a household survey for country i is available and \mathbf{g} is an error term.

¹² Mexico is the country with the longest recall periods. The household survey questionnaire asks about income in each of the previous six months, but we only use information on the previous month for consistency with the other countries.

¹³ PPP GDP per capita figures are from *World Development Indicators* 2000. For the Latin American countries with data for 1999 and 2000, the GDP data is from ESDB at the Inter American Development adjusted to PPP

Having direct access to each of the 144 household surveys and each survey questionnaire allows us to produce a data set with a high degree of comparability across observations, which minimizes measurement error in the dependent variable.¹⁴ For each country, we examine the survey questionnaire for each year and construct a household per capita income measure that is exactly the same for all years for that country. To do so we identify the minimum common denominator to ensure consistency. This implies discarding information on the income sources that do not appear in all surveys, which entails some loss of information, but we believe that there are larger gains from reducing the noise-to-signal ratio of the series. So, we are able to assure that the income concept is comparable *within* each country over time. The lack of comparability *across* countries that inevitably remains becomes irrelevant when regressions are estimated in first differences, as we do below.

The timing of the data 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.¹⁵ As shown in Appendix Table A1, 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 years and estimating the growth- poverty relationship with a reduced data set with observations spaced 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

accordingly. 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.

¹⁴ Although the Deininger-Squire data is a major improvement over the information available to authors such as Ahluwalia, it still has some limitations. Székely and Hilgert (1999), Atkinson and Brandolini (1999) and Pyatt (1999) provide a more thorough discussion of the limitations of secondary data sets. Panizza (2001) shows that the relation between inequality and economic growth changes substantially when strictly comparable data is used.

¹⁵ 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.

are corrected. Some authors acknowledge that successive spells within countries have one survey in common, and are therefore not independent observations.¹⁶ To produce our results we estimate our base regression in first differences, but in all cases we report robust standard errors to address this issue. Additionally, we perform all estimations using the Huber iteration to reduce the potential effect of outlier observations.

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.¹⁷ Although our central results refer to standard first difference estimations, we also test whether our conclusions are robust to the use of these techniques.¹⁸

B. Empirical Results

Table 1 presents our main results. 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. μ_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 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 not statistically significant. 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.

¹⁶ See especially Ravallion and Chen (1997).

¹⁷ Specifically, Dollar and Kraay (2000) use the Arellano and Bover estimator, which is similar to GMM estimators but does not include fixed effects.

¹⁸ 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.

The conclusions from Table 1 are at odds with the recent papers by Roemer and Gugerty (1997), Gallup, Radelet and Warner (1999) and Dollar and Kraay (2000), which argue that the poor gain one-for-one from growth in mean income.¹⁹ 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 10, 20 and 30 percent of the distribution and estimate equation (1) by using each of these as dependent variables.²⁰

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 elasticities of 1.019 and 0.92 reported by Dollar and Kraay (2000) and Roemer and Gugerty (1997), respectively, but which is lower than the elasticity of 1.16 in Gallup, Radelet and Warner (1999). Some of these authors report more than one estimate, but the comparison is only with those 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.

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 as those obtained by Ravallion and Chen (1997), and Bruno, Ravallion and Squire

¹⁹ 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 as the authors of these other papers.

²⁰ Household incomes are adjusted to match PPP GDP per capita.

(1998) with respect to the effect of growth on these two poverty measures.²¹ Our estimates of the growth elasticity of the headcount ratio and the poverty gap are -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 to 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 headcount 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. However, the same data, applied to the alternative methodology of using the general means to track the incomes of the poor, leads to a different conclusion about the relation between poverty and growth.

C. Robustness Tests

This section performs three sets of tests to check whether the conclusion stated above—that the growth elasticity of general means with low values of α is significantly smaller than one—holds under different situations. The first set of tests verifies if the results change when modifying the sample of countries and years used in the estimation. The second explores whether changing estimation techniques and modifying the methodology for estimating some of the variables has any implications. The third tests whether including control variables changes the value of the growth elasticity of general means.

Changes in Sample Composition

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

²¹ To compute the headcount ratio and the poverty gap we use a poverty line of 2-dollars-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.

are substantially smaller when the United States is excluded. So, the incomes of the poor are relatively less responsive to growth in our sample of developing countries.

The second column in Table 3 restricts the sample to the 17 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 μ_0 and μ_{-1} have higher growth elasticity in these countries.

Regressions 3 to 5 in Table 3 check the robustness of our conclusions to restricting the sample to smaller time periods. In the third column we present the elasticity from a sample that drops information from consecutive years. The data is modified so that there are at least three years between observations. 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 of observations. The growth elasticity derived from the restricted sample leads to exactly the same conclusions as in Table 1.

Column 4 asks whether the growth elasticity of general means differs when we use data from the 1990s only. This is to verify whether growth during the past decade has had a different impact on the incomes of the poor than growth in the whole period under observation. With the exception of the growth elasticity of μ_{-4} , which is much smaller than before, the results are practically the same as in Table 1. The difference in the elasticity of μ_{-4} suggests that during the 1990s the incomes of the poorest of the poor were less responsive to growth in mean income than in previous decades.

The last column in Table 3 presents the elasticity estimated with a sample that includes only episodes of positive growth in mean income. If during periods of positive growth governments “invest” more in the poor or the poor had better prospects for increasing their living standards, then the growth elasticity of the general means that emphasize more the incomes of the poor, would be larger than those in Table 1. However, the coefficients are very similar to those in Table 1, which suggests that the poor suffer from income contractions in a way similar to that in which they gain with expansion in the mean.

Changes in Estimation Procedures

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 first column in Table 4 presents the results. Our central conclusions are not modified, and the value of each growth elasticity is very similar to those obtained by applying OLS to the same data.

Estimations in first differences such as those performed so far have the advantage of reducing the problem of omitted variable bias by controlling for all time-invariant country characteristics, but this comes at the cost of reducing the time-series variation of our data. Perhaps the best solution to this problem is to estimate the Arellano-Bover System Estimator, which uses information on the levels and changes in the data as well as instrumental variables constructed with lagged variables. This technique reduces the problems of measurement error, endogeneity and omitted variables bias.²² The results are presented in the second column of Table 4 and lead to exactly the same conclusions as before: the incomes of the poor grow with positive changes in the ordinary mean, but they do so much less than the change in the average. We use the Arellano-Bover System Estimator only as a robustness check because of its well-known caveats of reducing standard errors excessively and relying on an ad-hoc lag structure for constructing the instrumental variables.

A criticism to which several of the papers in this strand of literature have been subject is that it makes no sense to compare growth in incomes from a household survey to changes in GDP from National Accounts. So far we have adjusted household incomes to match GDP in order to make our estimates more comparable with those in the literature, but we think this is a genuine concern. To address it we include a set of estimations that use the original household incomes (in real terms) to compute the general means and the ordinary mean. The third column in Table 4 presents the results. Each elasticity is considerably smaller than our base estimates in Table 1, but the general conclusion holds. The main difference is that even for μ_0 the growth elasticity is smaller than 1 (although not significantly smaller). This suggests that the growth elasticity 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.

²² See Arellano and Bover (1995).

One of the common concerns in empirical work using household survey data is that lower incomes are normally subject to larger measurement error. Thus, it is possible that standard of living measures such as μ_{-2} , μ_{-3} , μ_{-4} , which place more weight on lower incomes, could be subject to larger noise-to-signal ratios. There are many ways of reducing measurement error for low incomes, and perhaps the most drastic solution is to truncate the distribution at some point to eliminate them altogether. In order to test whether eliminating the incomes that are more subject to measurement error could modify our conclusions, we go back to the original micro data in each household survey, truncate the poorest 5 percent of the distribution, and then re-compute each general mean. The last regression in Table 4 presents the results analogous to those in Table 1, but where the general means are computed with truncated distributions. Each growth elasticity is somewhat larger than that in Table 1, but the conclusion still holds that for μ_{-2} , μ_{-3} , μ_{-4} the elasticity is significantly lower than one.

Changes in Empirical Specification

Table 5 includes the results of three regressions that explore whether variations in the way in which changes in the standard mean enter into the regression, modify any of our conclusions.²³ The first regression includes the change in the log of the standard mean as the rest of the regressions but adds the growth rate of the mean lagged five years. This is to test if growth takes some time to have an impact on the poor, for instance through trickle down. The coefficient estimates are very similar to those in Table 1, and the coefficients for the lagged growth variable (not reported here) turned out to be very small and non-significant for all parameter values. Thus, our evidence does not support the hypothesis of delayed trickle-down effects from growth.

The second regression uses the standard mean as dependent variable and adds the value of the standard mean (in levels) lagged five years. This is to ask whether richer countries, which presumably have greater possibilities of alleviating poverty, actually have a different growth-poverty relationship. Differences between these results and our original ones are not apparent in this case, either.

The last two columns report the results from a regression that introduces changes in the log of the standard mean as independent variable, as before, but adds a quadratic term. This is to

²³ Recently, Bourguignon (2000) argued that the appropriate specification to estimate the growth elasticity of poverty should also include some inequality measure as regressor. We have not followed this option here since the general means are the basis for computing a wide set of inequality measures, as discussed above.

test for the non-linearity of each elasticity and to confirm if the average elasticity is low for more negative values of α only because in high income countries changes in poverty are smaller due to the initial lower levels of poverty and vice versa. As can be seen in Table 5, the quadratic term is never significant in statistical terms. The coefficients for the linear term lead to the same conclusions as before.

Finally, although the central conclusions of this and of related papers in the literature are derived from specifications such as equation (1), there is also interest in knowing if there are specific policies that are associated with pro-poor growth. This is an important question from the policy perspective but it is also central to our discussion because of the omitted variable bias inherent in equation (1). Estimation in first differences controls for time-invariant country characteristics but not for time-variant ones. The rate of growth of GDP per capita is an outcome of more fundamental time-variant variables that are not included in the equation, but that might affect the incomes of the poor differentially.

To address this issue to some extent, we present in Table 6 regression results from a specification that includes the rate of growth of mean income as independent variable, as in our previous estimates, but adds terms that represent policy choices. Similarly to other papers we include a measure of imports plus exports as a share of GDP, which is a proxy for the degree of trade openness of a country, a bounded measure of inflation, and the share of government consumption in GDP.²⁴ We also include the share of M2 over GDP as a proxy for the degree of financial depth, under the argument that credit is a mechanism for social mobility, and access to it should create opportunities for all sectors of the population (all independent variables are in changes of log values). Each of the five regressions in the table refers to the use of one of five dependent variables: μ_0 , μ_{-1} , μ_{-2} , μ_{-3} , and μ_{-4} , respectively. Since the growth of mean income is also included, the coefficient for each of the control variables corresponds to the effects net of the impact on overall growth.²⁵

The first line presents the growth elasticity of each of the dependent variables. As before, we conclude that the elasticity is considerably lower at lower values of α . The growth elasticity

²⁴ All variables were obtained from the World Bank *World Development Indicators* 2000.

²⁵ We also estimated specifications including the Polity IV index of the degree of democracy in each country. The justification is that if democracy affects growth prospects and if democracies lead to more egalitarian societies, excluding this variable from the regression would aggravate omitted variable bias. The results from these regressions are not reported here because the democracy variable was never found to be statistically significant, and

of our proxy for trade openness is positive, statistically significant, and quite stable across the different regressions. A ten percent increase in the value of this indicator leads to an increase of around 2 percent over the general means.²⁶

The third line in Table 6 shows that inflation has a negative effect on the general means, but the effect is much more detrimental for poorer individuals. A ten percent increase in the inflation measure leads to reductions of almost 4 percent and 2 percent in the value of μ_{-4} , μ_{-3} , and μ_{-2} , respectively, but leaves μ_{-0} and μ_{-1} practically unchanged. This result is consistent with the argument that the poor have scarcer means to protect their incomes from inflation.

Government consumption as a share of GDP has a negligible effect on general means with $\alpha = 0$, $\alpha = -1$ and $\alpha = -2$ but is associated with income gains for the poorest individuals. Finally, Table 6 reports the elasticities to changes in financial depth. The elasticities for all general means are quite large, which suggests that policies oriented towards increasing the availability of credit in the economy tend to have a positive impact on the incomes of the poor, in addition to the effect that they may have on growth in mean incomes. The elasticity, however, declines with lower values of α , indicating that the poorest individuals benefit relatively less from credit expansion. One interpretation is that the poorest individuals are usually disconnected from the economic system and may find it harder to access formal financial markets.

4. Conclusions

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 in the economics literature as "equally distributed equivalent income" functions in Atkinson (1970), and this paper offers a characterization of the general means that provides considerable justification for their use in the present context. We suggest a new methodology based on general means for evaluating the relationship between economic growth and the incomes of the poor.

it was of low and similar value for all five general means. Excluding this variable from the regressions has only very marginal effects on the coefficients of the other policy variables in Table 6.

²⁶ Dollar and Kraay (2000) arrive at a similar conclusion, but Lundberg and Squire (2000) find that trade openness reduces the incomes of the poor. However, the latter result is based on a different openness measure and the application of different estimation techniques, so it is not strictly comparable to our results.

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.

Reinterpreting the general mean as a measure of social welfare makes it possible to provide a useful normative interpretation of our results. An elasticity less than one, such as those obtained here for low values of the parameter α , indicates that economic growth is somewhat less effective in generating an increase in social welfare as indicated by our income standards μ^* . As for inequality in this framework, an elasticity lower than one for parameter values $\alpha < 1$ implies that the associated Atkinson inequality measure is increasing. Hence, we have a method of evaluating the impact of growth on welfare and inequality, in addition to our original concern with low incomes.

Bibliography

- Aczel, J. 1966. *Lectures on Functional Equations and their Applications*. New York, United States: Academic Press.
- Adelman, I., and C. Morris. 1973. *Economic Growth and Social Equity in Developing Countries*. Stanford, United States: Stanford University Press.
- Ahluwalia, M. 1976. "Inequality, Poverty and Development." *Journal of Development Economics* 3: 307-42.
- Ahluwalia, M., N. Carter and H. Chenery. 1979. "Growth and Poverty in Developing Countries." *Journal of Development Economics* 6: 299-341.
- Anand, S. and R. Kanbur. 1985. "Poverty under the Kuznets Process." *Economic Journal* 95: S42-S50.
- . 1993a. "The Kuznets Process and the Inequality-Development Relationship." *Journal of Development Economics* 40: 25-52.
- . 1993b. "Inequality and Development: A Critique." *Journal of Development Economics* 41: 19-43.
- Anand, S., and A. Sen. 1996. "Gender Inequality in Human Development: Theory and Measurement." Occasional Paper 19. New York, United States: United Nations Development Programme.
- Arellano, M., and O. Bover. 1995. "Another Look at the Instrumental-Variable Estimation of Error-Components Models." *Journal of Econometrics* 68: 29-52.
- Atkinson, A.B. 1970. "On the Measurement of Inequality." *Journal of Economic Theory* 2: 244-263.
- Atkinson, A.B., and A. Brandolini. 1999. "Promises and Pitfalls in the Use of Secondary Data Sets: Income Inequality in OECD Countries." Forthcoming in *Journal of Economic Literature*.
- Barro, J.B. 1999. "Inequality and Growth in a Panel of Countries." Cambridge, United States: Harvard University. Mimeographed document.
- Blackorby, C., D. Donaldson and M. Auersperg. 1981. "A New Procedure for the Measurement of Inequality Within and Among Population Sub Groups." *Canadian Journal of Economics* 14: 665-685.

- Blackorby, C., D. Primont and R. Russell. 1978. *Duality, Separability, and Functional Structure: Theory and Economic Applications*. New York, United States: North-Holland.
- Bourguignon, F. 2000 "The Pace of Economic Growth and Poverty Reduction." Washington, DC, United States: World Bank. Mimeographed document.
- Bruno, M., M. Ravallion and L. Squire. 1998. "Equity and Growth in Developing Countries: Old and New Perspectives on the Policy Issues." In: V. Tanzi and K-Y. Chu, editors. *Income Distribution and High Quality Growth*. Cambridge, United States: MIT Press.
- De Janvry, A. and E. Sadoulet. 2000. "Growth, Poverty and Inequality in Latin America: A Causal Analysis, 1970-94." *Review of Income and Wealth* 46(3): 267-87.
- Deaton, A. 2000. "Counting the World's Poor: Problems and Possible Solutions." Princeton, United States: Princeton University. Mimeographed document.
- Deiningner, K., and L. Squire. 1996. "Measuring Inequality: A New Data Base." *World Bank Economic Review* 10(3): 565-91.
- Diewert, W.E. 1993. "Symmetric Means and the Choice Under Uncertainty." In: W.E. Diewert and A. Nakamura, editors. *Essays in Index Number Theory*. Volume 1. Amsterdam, The Netherlands: Elsevier.
- Dollar, D., and A. Kraay. 2000. "Growth Is Good for the Poor." Washington, DC, United States: World Bank. Mimeographed document.
- Foster, J., J. Greer and E. Thorbecke. 1984. "A Class of Decomposable Poverty Indices." *Econometrica* 52: 761-766.
- Foster, J.E., and A. Sen. 1997. "On Economic Inequality after a Quarter Century." In: A. Sen, editor. *On Economic Inequality*. Expanded edition. Oxford, United Kingdom: Oxford University Press.
- Foster, J.E., and A. Shneyerov. 1999. "A General Class of Additively Decomposable Inequality Measures." *Economic Theory* 14: 89-111.
- . 2000. "Path Independent Inequality Measures." *Journal of Economic Theory* 91: 199-222.
- Foster, J., and A. Shorrocks. 1988. "Poverty Orderings." *Econometrica* 56(1): 173-177.
- . 1991. "Subgroup Consistent Poverty Measures." *Econometrica* 59(3): 687-709.
- Foster, J.E., and M. Székely. 2001. "Using the General Means to Evaluate Inequality and Changes in Inequality." Washington, DC, United States: Inter-American Development Bank, Research Department. Mimeographed document.

- Gallup, J.L., S. Radelet and A. Warner. 1999. "Economic Growth and the Income of the Poor." CAER II Discussion Paper No. 36. Cambridge, United States: Harvard University, Harvard Institute for International Development.
- Gorman, W.M. 1968. "The Structure of Utility Functions." *Review of Economic Studies* 35: 357-390.
- Li, H., L. Squire and F. Zou. 1998. "Explaining International and Intertemporal Variations in Income Inequality." *Economic Journal* 108(446): 26-43.
- Lundberg, M., and L. Squire. 2000. "The Simultaneous Evolution of Growth and Inequality." Washington, DC, United States: World Bank. Mimeographed document.
- Marshall, A.W., and I. Olkin. 1979. *Inequalities: Theory of Majorization and its Applications*. New York, United States: Academic Press.
- Morley, S. 2000 *La Distribución del Ingreso en America Latina y el Caribe*. Santiago, Chile: 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, The Netherlands: Institute of Social Studies. Mimeographed document.
- Ravallion, M. 2000. "Growth and Poverty: Making Sense of the Current Debate." Washington, DC, United States: World Bank. Mimeographed document.
- Ravallion., M., and S. Chen. 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 M. Gugerty. 1997 "Does Economic Growth Reduce Poverty?" CAER II Discussion Paper No. 4. Cambridge, United States: Harvard University, Harvard Institute for International Development.
- Sen, A. 1976. "Poverty: An Ordinal Approach." *Econometrica* 44: 219-31.
- . 1999. *Development as Freedom*. New York, United States: Knopf.
- Smolensky, E., R. Plotnick, E. Evenhouse and S. Reilly. 1994. "Growth, Inequality and Poverty: A Cautionary Note." *Review of Income and Wealth* 40(2): 217-22.
- Székely, M., and M. Hilgert. 1999. "What's Behind the Inequality We Measure? An Investigation Using Latin American Data." Research Department Working 409.

Washington, DC, United States: Inter-American Development Bank, Research Department.

Timmer, P. 1997 “How Well Do the Poor Connect to the Growth Process?” CAER II Discussion Paper No. 17. Cambridge, United States: Harvard University, Harvard Institute for International Development.

World Bank. 2000. *World Development Indicators*. Washington, DC, United States: World Bank.

Figure 1

**Comparison of Living Standards in the
USA, UK and Sweden**

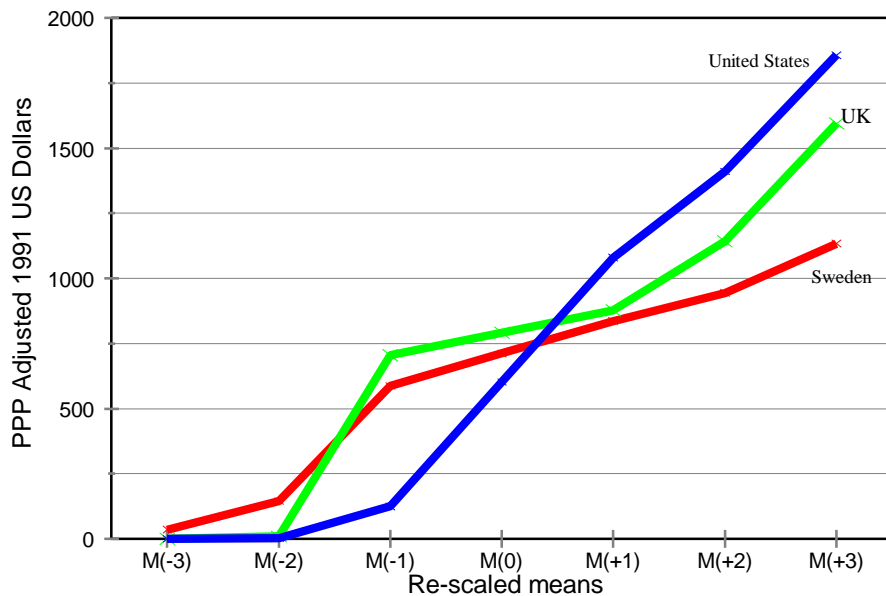


Figure 2

**Change in Standard of Living in
Mexico and Costa Rica**

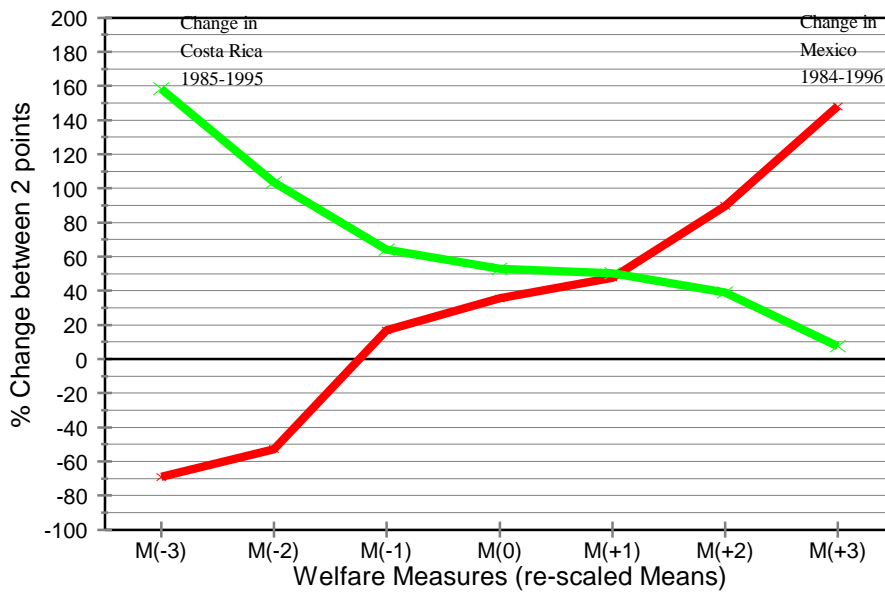


Table 1**Growth Elasticity of General Means****(Independent Variable is Growth in Mean Income)**

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

Source: Authors' calculations.

*Each of the elasticities reported

is estimated from a separate regression.

'T' Statistics appear under each coefficient

Table 2**Growth Elasticity of Various Welfare Measures****(Independent Variable is Growth in Mean Income)**

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

Source: Authors' calculations.

*Each of the elasticities reported

is estimated from a separate regression.

'T' Statistics appear under each coefficient

Table 3
Growth Elasticity of General Means Using Different Samples
(Independent Variable is Growth in Mean Income)

Dependent Variable	Regional Differences		Timing Differences		
	Less Developed Countries	Latin American Countries	Observations Spaced Every 3 Years	Observations for 1990s Decade	Positive Growth Episodes
	(1)	(2)	(3)	(4)	(5)
General Mean with parameter = 0	1.07 <i>5.95</i>	1.14 <i>5.39</i>	1.01 <i>6.17</i>	1.03 <i>13.07</i>	1.08 <i>9.46</i>
General Mean with parameter = -1	0.97 <i>1.22</i>	1.01 <i>3.77</i>	0.88 <i>5.00</i>	0.97 <i>3.96</i>	0.95 <i>2.79</i>
General Mean with parameter = -2	0.70 <i>1.40</i>	0.66 <i>0.15</i>	0.66 <i>1.27</i>	0.83 <i>1.05</i>	0.63 <i>0.77</i>
General Mean with parameter = -3	0.21 <i>0.15</i>	0.20 <i>0.45</i>	0.18 <i>1.00</i>	0.34 <i>0.22</i>	0.57 <i>0.57</i>
General Mean with parameter = -4	0.17 <i>0.60</i>	0.21 <i>0.50</i>	0.16 <i>0.85</i>	0.03 <i>0.13</i>	0.35 <i>0.64</i>
Number of Observations in Each Regression	101	74	88	82	94

Source: Authors' calculations.

*Each of the elasticities reported is estimated from a separate regression.

"T" Statistics appear under each coefficient

Table 4
Elasticity of General Means Using Different Estimation Techniques and Modified Variables
(Independent Variable is Growth in Mean Income)

Dependent Variable	Estimation Technique		Modified Variables	
	Instrumental Variables	Arellano-Bover System Estimator**	Unadjusted Incomes (Original Survey Incomes)	Top and Bottom Truncated Distributions
	(6)	(7)	(8)	(9)
General Mean with parameter = 0	0.94 <i>3.49</i>	0.89 <i>2.96</i>	0.82 <i>2.05</i>	1.06 <i>5.33</i>
General Mean with parameter = -1	0.85 <i>2.27</i>	0.74 <i>2.55</i>	0.59 <i>1.13</i>	1.03 <i>2.40</i>
General Mean with parameter = -2	0.73 <i>0.80</i>	0.45 <i>0.36</i>	0.43 <i>1.04</i>	0.80 <i>1.49</i>
General Mean with parameter = -3	0.33 <i>0.52</i>	0.28 <i>1.08</i>	0.22 <i>1.02</i>	0.69 <i>1.23</i>
General Mean with parameter = -4	0.28 <i>0.08</i>	0.13 <i>0.18</i>	0.20 <i>1.02</i>	0.55 <i>1.12</i>
Number of Observations in Each Regression	123	47	123	123

Source: Authors' calculations.

*Each of the elasticities reported is estimated from a separate regression.

** P-values associated with the test of veridentifying restrictions are 0.46, 0.41, 0.37, 0.37, and 0.36 for each regression, respectively.

"T" Statistics appear under each coefficient

Table 5
Elasticity of General Means Using Additional Independent Control Variables
(Independent Variable is Growth in Mean Income, plus an additional variable)

Dependent Variable	Includes Growth of Mean Income Lagged Five Years+ (10)	Includes GDP per Capita in First Period (12)	Regression With Quadratic Term	
			Coefficient Linear Term (11)	Coefficient Quadratic Term (11a)
General Mean with parameter = 0	0.94 <i>4.05</i>	0.92 <i>6.11</i>	0.86 <i>4.08</i>	0.81 <i>0.55</i>
General Mean with parameter = -1	0.71 <i>2.15</i>	0.88 <i>4.37</i>	0.68 <i>1.28</i>	0.66 <i>0.45</i>
General Mean with parameter = -2	0.68 <i>1.37</i>	0.73 <i>1.22</i>	0.39 <i>1.18</i>	0.30 <i>0.15</i>
General Mean with parameter = -3	0.40 <i>0.52</i>	0.52 <i>0.27</i>	0.33 <i>0.28</i>	0.32 <i>0.77</i>
General Mean with parameter = -4	0.43 <i>0.43</i>	0.17 <i>0.46</i>	0.24 <i>0.27</i>	0.20 <i>0.32</i>
Number of Observations in Each Regression	123	123	123	123

Source: Authors' calculations.

*Each of the elasticities reported is estimated from a separate regression.

'T' Statistics appear under each coefficient

+Regression includes log differenced current growth as control variable.

Table 6
Growth Elasticity of General Means Using Aggregate Control Variables

Dependent Variable (All variables in changes of log values)	Dependent Variable: General Mean				
	Parameter	Parameter	Parameter	Parameter	Parameter
	Value alpha=0 (12a)	Value alpha=-1 (12b)	Value alpha=-2 (12c)	Value alpha=-3 (12d)	Value alpha=-4 (12e)
Mean Income	1.00 <i>6.66</i>	0.86 <i>2.71</i>	0.65 <i>1.14</i>	0.34 <i>0.86</i>	0.28 <i>1.19</i>
(Exports+Imports)/GDP	0.21 <i>1.64</i>	0.25 <i>0.68</i>	0.23 <i>0.95</i>	0.15 <i>0.94</i>	0.20 <i>0.91</i>
Inflation (bounded)	-0.01 <i>-1.39</i>	-0.01 <i>-0.36</i>	-0.21 <i>-1.15</i>	-0.21 <i>-0.63</i>	-0.38 <i>-0.88</i>
Government Consumption/GDP	0.06 <i>0.56</i>	0.04 <i>0.12</i>	0.06 <i>1.08</i>	0.11 <i>1.48</i>	0.23 <i>1.48</i>
M2/GDP	0.94 <i>2.10</i>	0.93 <i>1.94</i>	0.76 <i>1.60</i>	0.57 <i>0.35</i>	0.49 <i>0.78</i>
Constant	-0.29 <i>-1.66</i>	0.14 <i>0.30</i>	-3.58 <i>-2.15</i>	-14.29 <i>-4.61</i>	-17.32 <i>-4.45</i>
F-Test	0.0000	0.0000	0.0020	0.0002	0.0002
R-Squared	0.506	0.111	0.167	0.179	0.173
Number of Observations	104	104	104	104	104

Source: Authors' calculations.

*Each of the elasticities reported is estimated from a separate regression.

'T' Statistics appear under each coefficient

Appendix Table A1

Household Surveys			
Country	# Surveys	Years	Survey
Argentina	3	1980, 96,98	Encuesta Permanente de Hogares
Bolivia	7	1986 1990, 93, 95 1996, 97 1999	Encuesta Permanente de Hogares Encuesta Integrada de Hogares Encuesta Nacional de Empleo Encuesta Continua de Hogares (condiciones de vida)
Brazil	11	1981, 83, 86, 88 1992, 93, 95, 96, 97,98,99	Pesquisa Nacional por Amostra de Domicilios 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 1987, 89, 91, 93, 95, 97, 98	Encuesta Nacional de Hogares - Empleo y Desempleo Encuesta de Hogares de Propósitos Múltiples
Dominican Republic	2	1996 1998	Encuesta Nacional de Fuerza de Trabajo 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 Nacional de Ingresos y Gastos Familiares
Honduras	6	1989, 92, 96, 97, 98,99	Encuesta Permanente de Hogares de Propósitos Múltiples
Mexico	7	1977 1984, 89, 92, 94, 96,98	Encuesta de Ingreso y Gasto de los Hogares 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 1991, 95, 97, 98,99	Encuesta de Hogares - Mano de Obra (EMO) Encuesta Continua de Hogares
Paraguay	2	1995 1998	Encuesta Nacional de Empleo Encuesta Integrada de Hogares
Peru	5	1985, 91, 94, 97,2000 1996	Encuesta Nacional de Hogares sobre Medición de Niveles de Vida Encuesta Nacional de Hogares sobre Niveles de Vida y Pobreza
Uruguay	6	1981, 89 1992, 95, 97,98	Encuesta Nacional de Hogares 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

Mathematical Appendix

Proof: Each of the general means m_a for $a \in R$ clearly satisfies symmetry, replication invariance, homogeneity, normalization, continuity and subgroup consistency. Hence one direction of the verification is immediate.

Alternatively suppose that the income standard $f: D^{\otimes} R$ satisfies these six properties. Select any $n \geq 3$ and let f_n be the restriction of f to D^n . Consider any two distribution x and y having the same population size $m < n$ and denote $m' = n - m$.

We can show that

$$(A1) \quad f_n(x, x') \geq f_n(y, x') \Rightarrow f_n(x, y') \geq f_n(y, y') \quad \text{for all } x', y' \in D^{m'}$$

Indeed suppose that $f_n(x, x') \geq f_n(y, x')$. Subgroup consistency rules out the possibility that $f(x) < f(y)$. Consequently, $f(x) \geq f(y)$. Now if this inequality is strict, then subgroup consistency immediately ensures that $f_n(x, y') > f_n(y, y')$. If, on the other hand, it holds with equality, then subgroup consistency rules out $f_n(x, y') < f_n(y, y')$. Since $f(x, y, y') < f(y, y', x)$ is a violation of symmetry. Thus $f_n(x, y') \geq f_n(y, y')$ as desired.

By (A1), we conclude that the continuous function f_n is strictly separable in each partition of incomes, and hence by standard results on separability (see, for example, Gorman, 1968 or Blackorby, Primont and Russel, 1978) f_n may be written as

$$(A2) \quad f_n(x) = F_n\left(\sum_{i=1}^n \Phi_n(x_i)\right) \quad \text{for } x \in D^n$$

where $\Phi_n : R_{++} \otimes R$ is continuous and F_n is continuous and strictly increasing. By the normalization property, we know that $s = F_n\left(\sum_{i=1}^n \Phi_n(s)\right)$ for $s > 0$, and hence $F_n^{-1}(s) = n\Phi_n(s)$.

Setting $G_n := F_n^{-1}(s)$, we can therefore express f_n as

$$(A3) \quad f_n(x) = G_n^{-1}\left(\frac{1}{n} \sum_{i=1}^n G_n(x_i)\right) \quad \text{for all } x \in D^n$$

where each $G_n : R_{++} \rightarrow R$ is strictly increasing and continuous. Clearly (A3) holds for any $n \geq 3$.

Replication invariance of f will allow us to replace each G_n with a single function G and to extend (A3) to the cases $n = 1$ and $n = 2$. Indeed, set $G := G_4$ and select $m = 4n$ where n is a positive integer. Applying G to (A3) yields

$$(A4) \quad G(f_m(x)) = G \left[G_m^{-1} \left(\frac{1}{m} \sum_{i=1}^m G_m(G^{-1}(G(x_i))) \right) \right] \\ = H_m^{-1} \left(\frac{1}{m} \sum_{i=1}^m H_m(G(x_i)) \right) \text{ for all } x \in D^m.$$

where $H_m(t) := G_m(G^{-1}(t))$ is continuous and strictly increasing in $t \in G(R_{++})$. In addition, we clearly have $H_4(t) = t$ for all $t \in G(R_{++})$.

Now consider any $y \in D^2$ with replication $y' \in D^4$ and $y'' \in D^m$. By replication invariance

$$G(f(y'')) = H_m^{-1} \left(\frac{1}{m} \sum_{i=1}^m H_m(G(y_i'')) \right) = H_m^{-1} \left(\frac{1}{2} H_m(u_1) + \frac{1}{2} H_m(u_2) \right) = G(f(y')) \\ = H_4^{-1} \left(\frac{1}{2} H_4(u_1) + \frac{1}{2} H_4(u_2) \right) = \frac{1}{2} (u_1 + u_2)$$

where $u_i := G(y_i)$ for $i = 1, 2$. Consequently, H_m must satisfy

$$(A5) \quad \frac{1}{2} H_m(u_1) + \frac{1}{2} H_m(u_2) = H_m \left(\frac{1}{2} (u_1 + u_2) \right) \text{ for all } u_1, u_2 \in G(R_{++})$$

The solution to this Jensen equation (Aczel, 1966, p.43) is $H_m(t) = a_m t + b_m$ for some constants $a_m, b_m \in R$, which along with (A4) implies

$$(A6) \quad G(f(x)) = \frac{1}{m} \sum_{i=1}^m G(x_i) \text{ for all } x \in D^m$$

for the cases where $m = 4n$ and n is some positive number.

Now for arbitrary $n \geq 1$ and $x \in D^n$ consider the replication $x' \in D^{4n}$ and note that

$$(A7) \quad G(f(x)) = G(f(x')) = \frac{1}{4n} \sum_{i=1}^{4n} G(x_i') = \frac{1}{n} \sum_{i=1}^n G(x_i)$$

making use of (A6) and replication invariance. Therefore,

$$(A8) \quad f(x) = G^{-1}\left(\frac{1}{n} \sum_{i=1}^n G(x_i)\right) \text{ for all } x \in D^n, n \geq 1,$$

where $G : R_{++} \rightarrow R$ is continuous and strictly increasing.

To determine form of G , consider any $x \in D^2$ and $I > 0$. By homogeneity and (A8) we have

$$(A9) \quad G\left[IG^{-1}\left(\frac{1}{2}G(x_1) + \frac{1}{2}G(x_2)\right)\right] = \frac{1}{2}G(Ix_1) + \frac{1}{2}G(Ix_2)$$

Define $u_i := G(x_i)$ so that $G^{-1}(u_i) = x_i$ for $i = 1, 2$. Then

$$(A10) \quad G\left[IG^{-1}\left(\frac{1}{2}u_1 + \frac{1}{2}u_2\right)\right] = \frac{1}{2}G(IG^{-1}(u_1)) + \frac{1}{2}G(IG^{-1}(u_2))$$

and hence

$$(A11) \quad G^I\left[G^{-1}\left(\frac{1}{2}u_1 + \frac{1}{2}u_2\right)\right] = \frac{1}{2}G^I(G^{-1}(u_1)) + \frac{1}{2}G^I(G^{-1}(u_2)),$$

where $G^I(s) := G(Is)$ for $I, s > 0$. Consequently, where $H^I(t) := G^I(G^{-1}(t))$, equation (A11) becomes

$$(A12) \quad H^I\left(\frac{1}{2}u_1 + \frac{1}{2}u_2\right) = \frac{1}{2}H^I(u_1) + \frac{1}{2}H^I(u_2) \text{ for all } u_1, u_2 > 0,$$

another Jensen equation with solution $H^I(t) = a^I t + b^I$ for some $a^I, b^I \in R$. Substituting $G(s) = t$ into this solution yields $G(Is) = a G(s) + b$, or more explicitly,

$$(A13) \quad G(Is) = a(I) G(s) + b(I) \text{ for all } s, I > 0,$$

for real-valued functions a and b . The solution to (A13), as given by Eichorn (1978, pg. 42) implies in particular that

$$(A14) \quad G(s) = \begin{cases} p \ln s + q & \text{for } \mathbf{a} \neq 0 \\ ps^{\mathbf{a}} + q \end{cases}$$

where $p \neq 0$ and $q \in R$ are constants. Consequently substituting (A14) into (A8) yields

$$f(x) = \begin{cases} \prod_{i=1}^n x_i^{1/n} & \text{for } \mathbf{a} \neq 0 \\ \left(\frac{1}{n} \sum_{i=1}^n x_i^{\mathbf{a}} \right)^{1/\mathbf{a}} & \text{for } \mathbf{a} = 0 \end{cases} \quad \text{for all } x \in D^n \text{ and } n \geq 1,$$

which immediately implies that $f(x) = \mathbf{m}_{\mathbf{a}}(x)$ for some $\mathbf{a} \in R$. This completes the proof.