

The Unfairness of (Poverty) Targets

Melanie Allwine

Jamele Rigolini

Luis F. López-Calva

The World Bank
Latin America and the Caribbean Region
Office of the Chief Economist
and Poverty Reduction and Economic Management Unit
February 2013



Abstract

Adopted on September 8, 2000, the United Nations Millennium Declaration stated as its first goal that countries “...[further] resolve to halve, by the year 2015, the proportion of the world’s people whose income is less than one dollar a day and the proportion of people who suffer from hunger...” Each country committed to achieve the stated goal, regardless of their initial conditions in terms of poverty and inequality levels. This paper presents a framework to quantify *how much* initial conditions affect poverty reduction, given a level of “effort” (growth). The framework used in the analysis allows for the growth elasticity of poverty to vary according to changes in the income distribution along the dynamic path of growth and redistribution, unlike previous examples in the literature where this is assumed to be constant. While wealthier countries did perform better in reducing poverty in the last decade and a half (1995–2008), assuming equal initial conditions, the situation reverses: the paper finds a statistically

significant *negative* relation between initial average income and poverty reduction performance, with the poorest countries in the sample going from the worst to the best performers in poverty reduction. The analysis also quantifies how much poorer countries would have scored better, had they had the same level of initial average income as wealthier countries. The results suggest a remarkable change in poverty reduction performance, in addition to the reversal of ranks from worst to best performers. The application of this framework goes beyond poverty targets and the Millennium Development Goals. Given the widespread use of targets to determine resource allocation in education, health, or decentralized social expenditures, it constitutes a helpful tool to measure policy performance toward all kinds of goals. The proposed framework can be useful to evaluate the importance of initial conditions on outcomes, for a wide array of policies.

This paper is a product of the office of the Chief Economist and Poverty Reduction and Economic Management Unit, Latin America and the Caribbean Region. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at <http://econ.worldbank.org>. The author may be contacted at jrigolini@worldbank.org.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

The Unfairness of (Poverty) Targets

Melanie Allwine

George Washington University

Jamele Rigolini

World Bank

Luis F. López-Calva¹

World Bank

¹ We would like to thank James Foster, Norman Loayza, Pedro Olinto, Renos Vakis and seminar participants at George Washington University for useful early discussions. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

Introduction

Adopted on September 8 of 2000, the United Nations Millennium Declaration stated as its first goal that countries “...[further] resolve to halve, by the year 2015, the proportion of the world’s people whose income is less than one dollar a day and the proportion of people who suffer from hunger...” (General Assembly resolution 55/2). The resolution was adopted by all 189 Member States of the United Nations, over 140 of them represented directly by their Head of State. Each country committed to achieve the stated goal, regardless of their initial conditions in terms of poverty and inequality levels.

The target of halving extreme poverty between 1990 and 2015 is on track to being met. The proportion of people living on less than \$1.25 a day in purchasing power parity (PPP) terms has already declined from 47 percent in 1990 to 24 percent in 2008, a reduction of more than 2 billion to less than 1.4 billion (United Nations, 2012). Yet, the UN progress report shows enormous differences across continents. The target will be met mainly due to the impressive poverty reduction achievements in Southeast Asia. Between 1990 and 2008, Southeast Asia reduced extreme poverty from 45 percent of the population to 17 percent. Looking at China alone, the progress is even more remarkable, with extreme poverty falling from 60 percent in 1990 to 13 percent in 2008. On the other extreme, during the same period, Sub-Saharan Africa reduced its extreme poverty by a modest 9 percentage points, from 56 to 47 percent.

Heterogeneity in achievements tends to be associated with differences in the rates of economic growth. Poverty plummeted alongside China’s stunning economic performance of the past three decades; while the more modest growth performance of Sub-Saharan countries led to less impressive rates of poverty reduction. Now that growth has picked up in Sub-Saharan Africa, the widespread expectation is that poverty reduction will show more dramatic results.

This paper does not question the fundamental role that growth has played, and always will as a necessary condition, in achieving sustainable poverty reduction. It shows, however, that poverty achievements are substantially affected by the way that poverty is measured and targets are set. While the Millennium Declaration encouraged all developing countries to pursue the goal of cutting poverty in half, it did not account for the fact that countries start off with different initial conditions, specifically in terms of where the poverty line is located with respect to the initial distribution of income. In fact, countries are evaluated according to the same overarching target, independently of where they started.

But the same policies and growth rates could have a dramatically different impact on poverty reduction, depending on where the poverty line stands with respect to the distribution of income. This paper

brings forth two elements to prior analysis on the subject of how initial conditions matter. It quantifies the effect of initial conditions (that is, average income per capita and inequality in the reference year) on outcomes given a level of “effort” (i.e. growth in per capita income). The underlying motivation supporting such a distinction is that, in any benchmarking exercise, countries and policymakers cannot be “held accountable” for the initial distribution of income, while they should (at least to some extent) be held accountable for the *evolution* of the income distribution. This evolution is characterized, in our exercise, by the growth in per capita income.²

The framework used in the analysis allows for the growth elasticity of poverty to vary according to changes along the dynamic path of growth and redistribution, unlike previous mechanical examples in the literature where this is assumed to be constant. The framework is also important in the context of aid allocation decisions (both domestic and international) that aim at rewarding “effort” and maximizing the poverty impact of transfers.

Specifically, the analysis quantifies *how much* poorer countries would have scored better, had they had the same level of initial per capita income as wealthier countries, by attributing to each country the median initial per capita income measured across countries. The results suggest a remarkable change in poverty reduction performance: while wealthier countries did perform better in reducing poverty in the last decade and the half (1995-2008), assuming equal initial conditions, we find the poorest countries in the sample going from the worse to the best performers in poverty reduction. The reversal of the relationship, and the large magnitude of such reversal extends to other dimensions of poverty beyond the headcount index, such as the poverty gap.

To measure whether the MDG targets may have been markedly unfavorable towards the poorest countries, the simulations are repeated for those countries which in the benchmark year of the MDGs (1990) had rates of extreme poverty above 10 percent. Interestingly, the difference between actual poverty reduction rates and the counterfactual ones becomes less striking, a result of this subgroup of countries being more similar in initial per capita income. Nonetheless, substantial differences subsist. For example, if China would have had in the 1990s the same per capita income as Ghana, with the same growth performance, it would have faced half of a percentage point less in annual poverty reduction. Accumulating this difference over 20 years, China would have had in 2008 extreme poverty rates almost 10 percentage points higher than the actual ones. As such, initial conditions appear to be an important

² While changes in inequality can also be characterized as “effort,” they impact less poverty reduction in the long run, and remain more challenging to simulate (see the discussion in the methodology section).

factor in the success of the poverty reduction efforts. This is particularly relevant given that, by disentangling policymakers' efforts from initial conditions out of their control, achievements can be benchmarked more accurately.

With the elements presented in this paper, we aim to contribute to the discussion of the new definition of objectives in the post-2015 MDG agenda. Moreover, the application of such a framework goes beyond the Millennium Development Goals. Given the widespread use of targets to determine resource allocation, in education, health, or decentralized social expenditures, it can constitute a helpful tool to measure policy performance towards all kinds of goals. The proposed framework can be useful to evaluate the importance of initial conditions on outcomes, for a wide array of policies. While the present analysis focuses on poverty, the methodology can be used as a principle to evaluate performance indicators in a more general context.

Basic Framework and Related Literature

The fact that initial per capita income and inequality levels matter in poverty reduction for a given a level of growth has been discussed previously in the literature. It has been first analyzed by Bourguignon (2002), who looks, similarly to us, at the growth elasticity of poverty to explain heterogeneity in poverty reduction across countries. The research analyzes the identity that links poverty reduction, mean income growth, and distributional change. The growth elasticity of poverty is found to be a decreasing function of the development level of a country and of the degree of inequality of the income distribution, under the assumption that income is log-normally distributed. Per the basic identity analyzed, a permanent redistribution of income plays two roles in poverty reduction: an instantaneous reduction through the distribution effect; and (as it contributes to a permanent increase in the growth elasticity of poverty) an acceleration of poverty reduction for a given rate of economic growth. Initial conditions thus, can play a crucial role in the transformation of growth into poverty reduction, with poorer countries having in general a lower elasticity of poverty to growth. The paper, however, fails to quantify these effects, and also to allow the elasticity of poverty reduction to vary along the growth and inequality path.


Ravallion (2012) analyzes the implications of initial distribution and high initial poverty to explain why poverty convergence is not observed across countries. Using a recent dataset for 100 developing countries, his findings indicate that, despite the so-called *advantages of backwardness and growth* in the development process, countries starting with higher poverty rates do not see higher proportionate

rates of poverty reduction. Two “poverty effects” work against mean-convergence. A high level of initial poverty has an unfavorable impact on growth, such that countries with a higher initial incidence of poverty are likely to experience lower rates of growth, controlling for the initial mean. At the same time, a high poverty rate makes economic growth less effective in reducing poverty. In other words, the advantage of starting out with a low mean income is wasted for many poor countries, given their high poverty rates. His analysis, however, focuses more on causal links, and less on the “mechanical” impact of initial conditions. Moreover, by using regression methodologies, the analysis only accounts for a fraction of the impact of initial conditions, and fails to capture nonlinear effects.

Along this line of reasoning, Easterly (2008) emphasizes how initial conditions have slowed down poverty reduction in Africa. He makes the point that the continent’s progress towards the poverty reduction Millennium Development Goal is unfairly benchmarked against that of other regions, given Africa’s poorer initial conditions. Given a log-normal distribution of income, low poverty elasticity of growth is found in a country with a low initial per capita income. In this sense, a higher growth of mean income would be required to achieve the same percentage reduction in poverty than in a country with a higher per capita income. By having the lowest per capita income of any region, Africa is disadvantaged in its goal of cutting poverty in half. The continent would need a higher economic growth than other regions’ to attain the goal, in order to compensate for its low-poverty elasticity. The argument lies at the heart of our analysis. The discussion in Easterly (2008), however, remains theoretical, as he makes no attempt to quantify the extent to which Sub-Saharan African countries may remain disadvantaged. In fact, in the context of the MDGs, initial conditions seem to matter *less* than one would have thought.

In order to analyze this process, the dynamic links between initial conditions, growth and poverty reduction should be understood. Figure 1 summarizes the main theoretical argument on which the present analysis is based. Consider two countries with identical distributions of income, but with different means: average income in country A is lower than average income in country B. Because the two distributions intersect the poverty line at different places, given equal growth rates in average income, poverty reduction rates will look differently as different numbers of people will “cross” the poverty line. In our example, the poverty line crosses both distributions on the right of their apex (i.e., we consider the case of two countries where a majority of the population is poor), and therefore, for equal growth rates poverty achievements will be more marked in the richer country. But, as Figure 1 shows, this does not always have to be the case, in particular if we compare countries with large differences in average income.

Figure 1: Distribution of income and poverty reduction



In setting universal targets of poverty reduction, many countries could thus be penalized simply due to their initial distribution of incomes. Uganda, for instance, had in 1990 a GDP per capita of 563 dollars (in PPP terms), less than half the GDP per capita of China (1,100 dollars). Because of these different initial conditions, in order to achieve the same poverty reduction rates as China, Uganda's growth rates would have had to be even higher than the Chinese ones. Setting equal poverty reduction targets for China and Uganda may thus not be fair.

As discussed, we are not the first to make the argument, but previous paper did not quantify these effects. Moreover, as Figure 1 shows, the growth elasticity of poverty varies *along* the growth path, hence by just considering initial elasticities (or average ones, as in Ravallion, 2012)), cross-country differences could be grossly over- or under-estimated. In this paper we aim at relax the constant elasticity assumption, and look for the actual growth realizations, to find out how much have actual poverty reduction achievements been affected by differences in initial conditions.

Data and Methodology

The main data source of our analysis is the World Bank's *Povcal* database (<http://iresearch.worldbank.org/PovcalNet>). The *Povcal* database provides information on mean expenditures or income (in 2005 USD PPP terms) and the Lorenz curve of expenditures/income distributions for specific countries and years, both of which are estimated from nationally representative household surveys. The *Povcal* database includes data on 131 countries. Since no parametric information is provided for Latin America, we also draw information directly from nationally representative household surveys for 18 Latin American countries.

The *Povcal* database reports Lorenz curves using two approximation methods: the *General Quadratic* and the *Beta* Lorenz curve approximations. Combined with information on mean expenditures or income, these approximations allow estimating poverty headcount and gap indexes quite accurately, as well as Gini coefficients (see Datt, 1998, and the Appendix for details on all approximations used in the analysis). For each country in *Povcal*, we use the approximation that best fits the actual poverty headcounts (see below). Additionally, for the 18 Latin American countries, we first convert income into 2005 USD PPP using conversion factors reported by the *International Comparison Program*, and then estimate ourselves the best parametric approximation of the income distribution.

To adequately represent poverty in all the countries analyzed, we use a 2.5 dollars a day poverty line. We base our estimation for each country either on income or expenditures, as reported by *PovCal* or the surveys (given the focus of the analysis on poverty change *within* countries, we do not apply any form of correction for harmonizing income and expenditure data). The timeframe we consider in our simulations is circa 1995 to 2008, which allows us to capture the largest number of countries. To correct for the fact that, for some countries, the time frame slightly diverts from these two reference years, we report results in terms of average annual changes. Moreover, to achieve comparability with the Millennium Development Goals (MDG) reports, we also repeat the analysis for a subset of countries with high initial rates of extreme poverty (more than 10 percent of the population with income/expenditures below 1.25 dollars a day), and look at changes between 1990 and 2008.

In our main analysis, out of the 84 countries available in the *Povcal* (including the Latin American ones), we exclude those that had poverty rates below 5 percent in 1995, since our approximation would capture with high errors the actual poverty changes in these countries. We also exclude Georgia from the analysis, as with poverty rates more than doubling in the period, it is a clear outlier. We are thus left with 74 countries (see Table 1). For the same reason, in the analysis focusing on the MDGs, we only

consider countries that had a rate of extreme poverty in 1990 above 10 percent. For this simulation we are left with 38 countries (see Table 2).

Accuracy of the parametric approximation

The core of our analysis is based on numerical simulations that apply changes to parameterized income distributions. The accuracy of such numerical approximations is therefore of central importance. Accordingly, in what follows, we compare the accuracy of each parametric approach using data for Latin America, where we can use the survey data to obtain direct estimates of the poverty rates.

Figure 2: Actual vs. Predicted headcount ratios (Latin America, 1995)




Figure 2 presents four different estimates of the headcount ratio (at a threshold of 2.5 dollars a day in PPP terms). The first is the actual poverty headcount ratio estimated directly from household surveys (our estimate of reference for the comparison). The second estimate represents the headcount ratio computed from the best fit between the General Quadratic and the Beta Lorenz approximations. The third estimate is the headcount ratio computed from an income distribution parameterized with a lognormal approximation. Finally, the last estimate is the headcount ratio computed from a fully nonparametric approximation of the shape of the income distribution. Among the three parameterizations, the one that deviates significantly from the actual poverty estimate is the lognormal

approximation. This is because the lognormal relies only on one parameter (the standard deviation) to capture the shape of the income distribution, while the other approximations rely on two or more, and are thus better able to capture the tails of the distributions. Whenever possible, we will therefore use the best fit of the General Quadratic and the Beta Lorenz curves (as reported by the *PovCal* database, and, for LAC by the comparison with actual data), and leave the use of the lognormal approximation only for simulations that will not allow for this.

Simulations

The heart of our analysis lies in simulating counterfactual poverty changes between 1995 and 2008, where we keep country-specific “performances,” but attribute equal “initial conditions” to countries. These two concepts remain specific to the way we structure the simulations, and deserve some attention.

Let us consider a space of income distributions $f(y, \mu, G)$ characterized by two parameters: *average income* μ , and an *inequality index* G . We then consider a *poverty index* $H(z, \mu, G)$ that can be mapped to each income distribution as follows:

$$H(z, \mu, G) = \int_0^{\infty} h(y, z) f(y, \mu, G) dy \quad (1)$$

where z can be thought of as an absolute reference threshold, which we shall denote as the *poverty line*, and $h(y, z)$ a function that can be integrated. Observe that common indexes, such as the poverty headcount and the poverty gap, can all be expressed in a form compatible with (1).

Consider then, two income distributions, $f(y, \mu_0, G_0)$, and $f(y, \mu_1, G_1)$. Think of the first set of parameters as characterizing the distribution of income in a given country in period zero (i.e., 1995 for our main simulations), and of the second set as characterizing the distribution of income in period one (i.e., 2008). For a poverty line z , changes in poverty between the two periods can then be expressed as: $H(\mu_1, G_1) - H(\mu_0, G_0)$.

Observe that, up to this point, we could run the comparison using actual poverty rates. The purpose of the parameterization, however, is to allow us to simulate counterfactual changes in poverty, so we proceed as follows. In the simulation, we aim to distinguish between *initial conditions* (μ_0, G_0) and *policy performance* $(\Delta\mu/\mu_0, \Delta G/G_0)$. The underlying motivation supporting such a distinction is that, in

any benchmarking exercise, countries and policymakers cannot be “held accountable” for the initial distribution of income, while they should (at least to some extent) be held accountable for the *evolution* of the income distribution. This evolution is characterized, in our exercise, by the growth in mean income and changes in inequality.

The idea is then to isolate performance from initial conditions in poverty reduction. To do this, in computing counterfactual poverty changes, we attribute to each country the median $(\bar{\mu}_0, \bar{G}_0)$ of the initial parameters in our sample of countries, but keep, for each country, the country-specific performance $(\varepsilon_\mu^c = \Delta\mu^c / \mu_0^c, \varepsilon_G^c = \Delta G^c / G_0^c)$. Our counterfactual simulated poverty change can therefore be expressed as follows:

$$\Delta\hat{H}^c = H(\varepsilon_\mu^c \cdot \bar{\mu}_0, \varepsilon_G^c \cdot \bar{G}_0) - H(\mu_0^M, G_0^M) \quad (2)$$

The simulated poverty change considers therefore the country-specific improvements in per capita income and inequality $(\varepsilon_\mu^c, \varepsilon_G^c)$ that the countries actually experienced, but applies them to median initial conditions, which are kept equal across countries.

We conclude with an observation about the alternative parameterizations that drive our simulations. Observe that our simulations are based on the assumption that it is possible to capture the shape of the income distribution – which drives inequality – by a single inequality parameter, G . This is a rather restrictive assumption, which may hamper the quality of the numerical approximation. For instance, the lognormal approximation, which summarizes an income distribution based on its average and standard deviation, tends to approximate poorly the tails of a distribution, which can generate substantial errors in approximating poverty rates (Figure 2), and even worse errors when approximating *changes* in poverty (Figure 3). This is the reason why both the *General Quadratic* and the *Beta* approximations used in *Povcal*, as mentioned above, are based on multiple parameters that are calibrated to achieve a good fit of the tails.


Accuracy notwithstanding, from the point of view of our simulations, the problem with using more than one parameter in approximating the shape of the income distribution is that it breaks the bijective relation between the parameter “ G ” in the poverty index (1), and any economic index of inequality that has economic meaning. Consider, for instance, measuring progress in inequality by the Gini coefficient,

and the case of the *General Quadratic* approximation. Under such an approximation, the shape of an income distribution is approximated by three parameters, a , b and c (see the Appendix). Any change in the Gini coefficient can thus be associated with infinite combinations of changes in these three parameters; but these combinations have different impacts on the ultimate change in poverty, as they affect the shape of the income distribution differently. With any approximation that uses two or more parameters to approximate the shape of an income distribution, it is therefore not possible to conduct counterfactual simulations as in (2).

In running the simulations we are thus faced with two choices: either we run the simulation using the lognormal approximation – where we can simulate changes both in average per capita income and in the Gini coefficient – or we only simulate changes in average per capita income (keeping inequality constant at its initial level) using, however, a more accurate approximation. Neither approach is perfect, hence, the decision of which one to use ultimately boils down to minimizing approximation errors. For these purposes, in Figure 3 we compare for Latin America (where we have the actual household surveys) the actual changes in poverty against three simulation alternatives that all use the lognormal parameterization:³ the first approach takes into consideration the actual changes in per capita income, but keeps income inequality at its initial level; the second approach takes into consideration actual changes in inequality, but keeps per capita income at its initial level; and, finally, the third approach takes into consideration both actual changes in per capita income and inequality.

³ The simulations use the first and second moments of the actual income distributions as the parameters for the lognormal approximation.

Figure 3: Contribution to poverty reduction of mean income growth vs. changes in inequality, Latin America, ca. 1995-2008



Two facts emerge. First, while the lognormal is already a poor parameterization to approximate poverty *levels*, regarding approximating poverty *changes*, its performance is even worse: in some cases, the approximation misrepresents changes in poverty by up to 30 percent of the actual value. The use of the lognormal approximation should thus be avoided, when possible. Second, for countries where poverty declined significantly between 1995 and 2008, the difference between considering both changes in per capita income and inequality, and only changes in per capita income, remains relatively small. This finding suggests that, for countries that faced large reductions in poverty, changes in per capita income preponderantly capture the variation in poverty. Given our focus on long term changes in poverty, for the sake of greater accuracy, in what follows we thus use the more accurate *General Quadratic* and *Beta* approximations, and run all the simulations maintaining income inequality constant at its country-specific initial level.


Main Findings

Our first exercise is to assess how *ranks* in poverty reduction performance are affected by the assumption of equal initial conditions across countries. The simulated ranks present the great advantage of depending only on the assumption of giving equal initial levels of average income μ_0 to all countries,

and not on the specific value of the initial average per capita income that we choose for the simulation. The reason is simple. Observe that the poverty index (1) decreases monotonically with average income. Thus, once initial conditions have been equalized, there is a bijective relation between growth in average income, and the actual poverty reduction performance. Hence, in the simulations, only growth but not the value chosen as an equalizing initial condition, affects the rank.

Figure 4 shows the actual ranks in poverty reduction performance (measured as the annual percentage change in poverty between 1995 and 2008), and the simulated rank under the assumption of equal initial average income across countries. The counterfactual picture we obtain changes dramatically the performance of poorer vs. richer countries. Before assuming initial conditions, we observe an increasing (and statistically significant) relation between initial average income, and the rank in poverty reduction performance (measured as percentage change in poverty). The poorest countries in the sample, such as Malawi, Madagascar and Burundi score around the worse in the sample, while many of the wealthiest countries, such as Russia, Chile and Malaysia score high. It would be tempting to conclude that richer countries are better managed to reduce poverty, but the conclusion would be misleading. When we give each country equal initial conditions, the picture *reverses*: we observe a statistically significant *negative* relation between initial average income, and poverty reduction performance. For instance, Mali, Nepal, and Niger – which previously performed poorly – are now among the best performers in poverty reduction.

Figure 5: Actual and simulated poverty reduction (rank; annual percentage points), ca. 1995-2008



Note: the shaded area represents the 95 percent confidence bands.

Looking only at ranks, however, only provides a partial picture. Ranks do not allow quantifying *how much* poorer countries would have scored better, had they had the same level of initial average income as wealthier countries. To quantify the change in the performance of poorer countries, we attribute next to each country the median initial income per capita in our sample, which happens to be that of Jamaica (4.1 dollars in 2005 PPP terms). Figure 6 shows the results. Not only do ranks get reversed (and both actual and simulated relationships with initial per capita income remain significant), but the magnitude of change in poverty reduction performance is quite impressive. If Mali, Nepal, Niger and Tajikistan would have had the initial per capita income of Jamaica in 1995, they would have virtually eradicated poverty by now. On the other side of the spectrum, Chile and Russia, which by now have poverty levels below 1 percent, would have been facing in 2008 poverty rates of 39 and 17 percent, respectively.

Figure 6: Actual and simulated annual poverty changes, ca. 1995-2008





Figure 7 shows that the reversal of the relationship, and the large magnitude of such reversal, extends to other dimensions of poverty, such as the poverty gap. Again, had Mali, Nepal, Niger and Tajikistan had the same initial per capita income as Jamaica in 1995, they would have literally closed the poverty gap by now. In Chile, on the other hand, due to its high income inequality, a significant amount of the population would still remain well below the poverty line (i.e., a poverty gap of 0.14).

Figure 7: Actual and simulated annual poverty gap changes, ca. 1995-2008




Do initial conditions matter for the Millennium Development Goals?

The previous analysis comprised a quite heterogeneous group of countries, from low to middle income ones. The equalization of initial average income across all countries can thus imply, for the poorest and richest countries, major shifts of where the poverty thresholds crosses the income distribution.

It is then natural to ask whether the drastic reversals of trends observed in the previous figures would hold for more homogenous groups of countries. Such a question is essential in gauging whether the MDG targets may have been excessively unfavorable towards the poorest countries.

Towards this purpose, we repeat the simulations for those countries in our database that in 1990 (the benchmark year of the MDGs) had rates of extreme poverty above 10 percent (that is, 10 percent or more of the population in 1990 living on less than 1.25 dollars a day). To capture the largest time span, we run the simulations for the period 1990-2008, but since the timeframe varies significantly across countries at times, we report average annual changes.

Figure 8: Actual and simulated annual changes in extreme poverty, ca. 1990-2008



Note: poverty figures are population weighted.

Among the subset of countries, we have chosen as initial condition the median average income in 1990, which corresponds to the daily per capita income of Ghana (1.6 dollars a day) for the simulation. For the sake of clarity, and to enhance similarities with the MDG reports, we present in [Figure 8](#) the regional poverty reduction averages.⁴

Because this subgroup of countries is more similar in initial average income than the larger group of countries from the previous analysis, the difference between actual poverty reduction rates and the ones that countries would have experienced under equal initial conditions is less striking. Nonetheless, substantial differences remain. If China would have had in the 1990s the same average income than Ghana, with the same growth performance, it would have faced half of a percentage point less in annual poverty reduction. Given that these differences accumulate over a time span of almost twenty years,

⁴ To obtain regional averages, we weighted the annual poverty reduction in each country in our sample by the initial population of that country. Our comparison does not mirror exactly the MDG reports for several reasons: (i) the sample of countries differ; (ii) we chose to show poverty reduction in India and China separately; and (iii) because of limited data availability for some regions, we have grouped countries differently (see Table 2).

China would have thus had in 2008, extreme poverty rates almost 10 percentage points higher than the actual ones.

In full similarity with China, India would have also faced lower annual rates of extreme poverty reduction by a third of a percentage point under these conditions. Observe, however, that even if India and China would have had the same initial conditions as Ghana, because of their stunning growth performance the countries would still be facing much higher rates of extreme poverty reduction: almost twice (for India) and four times (for China) the average of Sub-Saharan African countries.

On the other side of the spectrum, had Latin American countries had the same average income as Ghana, their rates of extreme poverty reduction would have been more than double the actual ones. This is because Latin American countries (and in particular Brazil) started, on average, with a much higher income per capita. Thus, in the 1990s, the extreme poverty line was already crossing the income distribution past its mode, at a point where the income-growth elasticity of poverty reduction was already relatively low.

Finally, observe that, on average, the simulations suggest that poverty reduction would not have changed much for Sub-Saharan Africa if all its countries would have had the same average income as Ghana. This result, however, depends very much on the choice of the level of the initial conditions, and on the fact that for most Sub-Saharan African countries in the sample the initial level of average income remained relatively close to that of Ghana (relative to, say, that of China or Latin American countries), so that poverty reduction rates for Sub-Saharan countries remain relatively unaffected by the change in initial conditions.

Implications and Conclusions

This paper quantifies how much initial conditions – the position of the poverty line with respect to the distribution of income – affect poverty reduction, given a level of “effort” measured in terms of growth in per capita income. While wealthier countries did perform better in reducing poverty in the last decade and a half (1995-2008), assuming equal initial conditions, the situation reverses: we find a statistically significant *negative* relation between initial average income and poverty reduction performance, with the poorest countries in the sample going from being the worst to the best performers in poverty reduction. The fundamental result also holds if we consider absolute changes in poverty reduction vis-à-vis percentage ones.

Initial conditions appear to be an important factor in the success of the poverty reduction efforts. While policymakers can influence changes in the mean and shape of the income distribution, they cannot affect initial conditions. This information is particularly relevant given that, by disentangling policymakers' efforts from initial conditions out of their control, achievements can be benchmarked more accurately. Collier and Dollar (2002), for example, have analyzed the distribution of aid that would maximize poverty reduction considering the quality of policies and the initial levels of poverty. Their framework, however, fails to fully disentangle the effect of initial conditions to assess effort and it does not incorporate the dynamics in a way this framework does by allowing for a change of the growth elasticity of poverty over time.

While our analysis focuses on poverty, the proposed framework remains relevant for indicators beyond poverty – it can be applied to any indicator that divides an underlying distribution. Given the importance of target-based mechanisms for different purposes, including allocating resources, the methodology presented can be key to correctly assess the importance of initial conditions in evaluating policy performance towards different goals.

References

- Bourguignon, François (2002) "The growth elasticity of poverty reduction: explaining heterogeneity across countries and time periods," *DELTA Working Papers* 2002-03, DELTA (Ecole normale supérieure).
- Collier, P. and D. Dollar (2002), "Aid Allocation and Poverty Reduction", *European Economic Review* 46, 1475-1500.
- Datt, Gaurav, (1998) "Computational tools for poverty measurement and analysis," *FCND discussion papers* 50, International Food Policy Research Institute (IFPRI).
- Easterly, William (2009) "How the Millennium Development Goals are Unfair to Africa," *World Development*, Vol. 37, No. 1, pp. 26-35.
- Kakwani, N. (1980). "On a class of poverty measures." *Econometrica*. 48(2). 437–446.
- Ravallion, Martin (2012) "Why Don't We See Poverty Convergence?," *American Economic Review*, 102(1): 504–23.
- United Nations (2012) "The Millennium Development Goals Report 2012". New York: UN.
- United Nations General Assembly (2000). United Nations Millennium Declaration. GA Res. 55/2, UN GAOR, 55th session Sess., Supp. 2, U.N. Doc. A/RES/55/2 (2000) 5.
- Villasenor, J. and B. C. Arnold (1989). "Elliptical Lorenz curves." *Journal of Econometrics*. 40(2). 327–338.

Table 1: Poverty trends (USD 2.5 a day)

Country	Year	Period 1			Period 2			
		Average monthly per capita income (USD PPP)	Poverty headcount	Gini	Year	Average monthly per capita income (USD PPP)	Poverty headcount	Gini
Albania	1997	151	0.16	0.29	2008	174	0.11	0.35
Argentina	1995	378	0.11	0.49	2008	413	0.09	0.46
Armenia	1996	108	0.51	0.44	2008	127	0.25	0.31
Azerbaijan	1995	87	0.54	0.35	2008	201	0.07	0.34
Bangladesh	1995	42	0.91	0.33	2010	52	0.86	0.32
Bolivia	1997	214	0.35	0.53	2008	214	0.32	0.52
Brazil	1995	297	0.29	0.60	2008	356	0.17	0.51
Bulgaria	1995	282	0.05	0.31	2007	274	0.02	0.28
Burundi	1992	26	0.97	0.33	2006	29	0.96	0.33
Cambodia	1994	65	0.77	0.46	2007	78	0.71	0.40
Cameroon	1996	82	0.65	0.41	2007	115	0.43	0.39
Central African	1993	25	0.94	0.61	2008	51	0.86	0.56
Chile	1994	378	0.12	0.50	2009	494	0.01	0.42
China	1996	60	0.76	0.32	2008	110	0.47	0.35
Colombia	1996	303	0.22	0.46	2008	343	0.20	0.51
Costa Rica	1995	260	0.17	0.46	2008	372	0.10	0.50
Cote d'Ivoire	1995	81	0.62	0.37	2008	88	0.59	0.42
Croatia	1998	513	0.00	0.27	2008	88	0.59	0.42
Dominican Rep	1996	252	0.18	0.47	2008	246	0.19	0.49
Ecuador	1995	198	0.36	0.54	2008	254	0.21	0.51
Egypt	1995	98	0.46	0.30	2008	114	0.32	0.31
El Salvador	1995	194	0.28	0.47	2008	214	0.22	0.47
Ethiopia	1995	45	0.90	0.40	2005	51	0.88	0.30
Ghana	1998	63	0.74	0.41	2005	80	0.64	0.43
Guatemala	2000	186	0.35	0.45	2006	200	0.35	0.47
Guinea	1994	42	0.88	0.45	2007	57	0.80	0.39
Honduras	1995	121	0.54	0.53	2008	225	0.42	0.33
India	1993	47	0.90	0.36	2010	60	0.81	0.33
Indonesia	1996	52	0.87	0.31	2008	72	0.69	0.33
Jamaica	1993	125	0.32	0.36	2004	279	0.11	0.46
Jordan	1997	152	0.22	0.36	2008	200	0.08	0.34
Kazakhstan	1996	137	0.28	0.35	2008	333	0.00	0.26
Kenya	1994	78	0.66	0.42	2005	65	0.77	0.46
Kyrgyz Republic	1993	173	0.37	0.54	2008	134	0.31	0.37
Lao PDR	1997	49	0.88	0.35	2008	63	0.78	0.37
Macedonia	1998	192	0.07	0.28	2008	294	0.09	0.44
Madagascar	1993	36	0.93	0.46	2010	28	0.95	0.44
Malawi	1997	30	0.95	0.50	2004	34	0.94	0.39
Malaysia	1995	263	0.17	0.49	2009	400	0.06	0.46
Maldives	1998	205	0.43	0.63	2004	176	0.20	0.37
Mali	1994	24	0.96	0.51	2010	46	0.87	0.33
Mauritania	1995	79	0.62	0.37	2008	84	0.61	0.40
Mexico	1996	193	0.35	0.51	2008	312	0.15	0.48
Moldova	1997	95	0.50	0.37	2008	183	0.14	0.35
Mongolia	1995	81	0.58	0.33	2007	150	0.25	0.37
Morocco	1999	130	0.36	0.39	2007	161	0.24	0.41
Mozambique	1996	30	0.95	0.44	2008	47	0.88	0.46
Namibia	1993	147	0.68	0.74	2004	146	0.59	0.64
Nepal	1995	38	0.94	0.35	2010	68	0.72	0.33
Nicaragua	1993	107	0.59	0.57	2005	148	0.42	0.46
Niger	1994	30	0.95	0.42	2007	53	0.85	0.35
Nigeria	1996	39	0.91	0.47	2010	40	0.90	0.49
Pakistan	1996	47	0.91	0.29	2007	66	0.76	0.30
Panama	1995	275	0.28	0.58	2009	325	0.17	0.52
Paraguay	1995	286	0.27	0.52	2008	265	0.21	0.48
Peru	1997	203	0.32	0.54	2008	251	0.21	0.49
Philippines	1994	83	0.64	0.43	2009	104	0.53	0.43
Poland	1996	138	0.23	0.33	2008	370	0.01	0.34
Romania	1994	99	0.39	0.28	2008	227	0.04	0.31
Russia	1996	287	0.12	0.46	2008	471	0.00	0.42
Senegal	1994	50	0.87	0.41	2005	67	0.72	0.39
South Africa	1995	158	0.49	0.57	2009	257	0.40	0.63
Sri Lanka	1995	83	0.62	0.35	2006	119	0.43	0.40
Swaziland	1995	34	0.93	0.35	2009	80	0.69	0.51
Tajikistan	1999	43	0.92	0.29	2009	100	0.42	0.31
Tanzania	1991	33	0.95	0.34	2007	37	0.93	0.38
The Gambia	1998	42	0.87	0.50	2003	82	0.66	0.47
Tunisia	1995	154	0.30	0.42	2005	218	0.15	0.41
Turkey	1994	204	0.17	0.42	2008	305	0.07	0.39
Uganda	1996	40	0.91	0.37	2009	68	0.76	0.40
Ukraine	1995	205	0.14	0.39	2008	324	0.00	0.28
Venezuela	1995	177	0.29	0.47	2006	213	0.21	0.45
Vietnam	1993	40	0.91	0.36	2008	85	0.58	0.36
Zambia	1996	46	0.87	0.48	2006	42	0.87	0.55

Table 2: Poverty trends (USD 1.25 a day)

Country	Period 1				Period 2			
	Year	Average monthly per capita income (USD PPP)	Poverty headcount	Gini	Year	Average monthly per capita income (USD PPP)	Poverty headcount	Gini
Bangladesh	1991	34	0.70	0.28	2010	52	0.43	0.32
Brazil	1990	300	0.13	0.61	2008	356	0.08	0.55
Burundi	1992	26	0.84	0.33	2006	29	0.81	0.33
Cambodia	1994	65	0.44	0.46	2007	78	0.32	0.40
Central African	1993	25	0.83	0.61	2008	51	0.63	0.56
China	1990	41	0.60	0.29	2008	133	0.13	0.38
Cote d'Ivoire	1988	99	0.14	0.37	2008	88	0.24	0.42
Ecuador	1994	182	0.16	0.47	2008	254	0.08	0.51
El Salvador	1991	168	0.18	0.52	2008	214	0.07	0.47
Ethiopia	1995	45	0.61	0.40	2005	51	0.39	0.30
Ghana	1991	49	0.51	0.38	2005	80	0.28	0.43
Guatemala	2000	186	0.13	0.45	2006	200	0.14	0.47
Honduras	1991	91	0.35	0.51	2008	225	0.24	0.33
India	1988	45	0.54	0.31	2010	60	0.26	0.33
Indonesia	1990	43	0.54	0.29	2008	72	0.23	0.33
Kenya	1992	93	0.38	0.57	2005	65	0.44	0.46
Lao PDR	1992	43	0.56	0.30	2008	63	0.34	0.37
Madagascar	1993	36	0.72	0.46	2010	28	0.81	0.44
Mali	1994	24	0.86	0.51	2010	46	0.50	0.33
Mauritania	1993	71	0.43	0.50	2008	84	0.23	0.40
Mongolia	1995	81	0.19	0.33	2007	150	0.00	0.37
Namibia	1993	147	0.49	0.74	2004	146	0.32	0.64
Nepal	1995	38	0.68	0.35	2010	68	0.25	0.33
Nicaragua	1993	132	0.18	0.50	2005	148	0.17	0.46
Niger	1992	34	0.73	0.36	2007	53	0.44	0.35
Nigeria	1992	40	0.62	0.45	2010	40	0.68	0.49
Pakistan	1990	38	0.65	0.33	2007	66	0.21	0.30
Panama	1991	215	0.21	0.58	2009	325	0.04	0.52
Peru	1997	203	0.14	0.54	2008	251	0.06	0.49
Philippines	1991	81	0.31	0.44	2009	104	0.18	0.43
Senegal	1991	45	0.66	0.54	2005	67	0.34	0.39
South Africa	1993	172	0.24	0.59	2009	257	0.14	0.63
Sri Lanka	1991	76	0.15	0.32	2006	119	0.07	0.40
Swaziland	1995	34	0.79	0.35	2009	80	0.41	0.51
Tanzania	1991	33	0.73	0.34	2007	37	0.68	0.38
Uganda	1989	37	0.69	0.44	2009	68	0.38	0.40
Vietnam	1993	40	0.64	0.36	2008	85	0.17	0.36
Zambia	1991	47	0.63	0.60	2006	42	0.69	0.55

Appendix. Description of Parameterizations

The lognormal approximation assumes that the income distribution can be approximated by a lognormal distribution of the following form $f_x = \frac{1}{x\sigma\sqrt{2\pi}} e^{-\frac{(\ln x - \mu)^2}{2\sigma^2}}$, where x is income, f_x is the probability density function, μ is the mean and σ is the standard deviation. The lognormal distribution is thus fully described by its mean and standard deviation, implying that one only needs to estimate the mean and standard deviation of income to fully parameterize the distribution.

The General Quadratic approach of Villasenor and Arnold (1989) uses a general quadratic form to estimate the Lorenz curve. The general quadratic form of a Lorenz curve can be rewritten as $L(1 - L) = a(P^2 - L) + bL(P - 1) + c(P - L)$ where P is the cumulative proportion of the population and L is the cumulative share in aggregate consumption or income. To calculate the parameters a , b , and c we rename $t = L(1 - L)$, $u = (P^2 - L)$, $v = L(P - 1)$, and $w = (P - L)$ and run the ordinary least-squares regression of t on u , v , and w with no intercept. Datt (1998) provides details on parameter restrictions, and on how to derive Gini and poverty headcount coefficients.

The Beta Lorenz approach of Kakwani (1980) uses a beta distribution to approximate the Lorenze curve. This implies that the Lorenz curve can be written as $L = P - \theta P^\gamma (1 - P)^\delta$, where again P is the cumulative proportion of the population and L is the cumulative share in aggregate consumption. Then to estimate the parameters θ , γ , and δ one must simply rearrange and take the logarithm of the above in order to get $\ln(L - P) = \ln(-\theta) + \gamma \ln(P) + \delta \ln(1 - P)$, rename $t = \ln(L - P)$, $u = \ln(P)$, $v = \ln(1 - P)$ and then estimate the parameters by running the ordinary least squares regression of t on u and v . See also Datt (1998).

For the nonparametric approximation we use a nonparametric form to approximate the income distribution. In practice, we use a univariate kernel density estimation, where the density of x can be estimated as: $\hat{f}_K(x; h) = \frac{1}{W} \sum_{i=1}^n \frac{w_i}{h} K\left(\frac{x-x_i}{h}\right)$ where $W = \sum_{i=1}^n w_i$, $K(z)$ is a kernel function, and h is the smoothing parameter (ie the bandwidth). We use an Epanechnikov kernel function.

The goodness of fit criterion defined in Datt (1998) is used to determine which of the parameterizations of the Lorenz curve (the General Quadratic or the Beta Lorenz) is the best fit. This goodness of fit criterion is simply the sum of the squared errors between the actual Lorenz and projected Lorenz from the respective approximation up to the estimated head count ratio. Because the Lorenz curve is

continuous, the percentage of the population is separated into 1,000 bins. Datt (1998) describes this mathematically as follows:

$$SSE = \sum_{i=1}^k (\hat{L}_i - L_i)^2$$

where L_i is the estimated Lorenz curve from the specified parameterization, L_i is the actual Lorenz curve, and $k = \lceil k | \sum_{i=1}^{k-1} p_i \leq \hat{H}_1 \leq \sum_{i=1}^k p_i \rceil$ where H_i is the estimated Head Count Ratio. The sum of squared errors (SSE) is calculated for each specification and the parameterization with the lowest SSE is used.