

World Bank

From the Selected Works of Martin Ravallion

December, 2012

How Long will it Take to Lift One Billion People out of Poverty?

Martin Ravallion, *World Bank*



Available at: https://works.bepress.com/martin_ravallion/26/

How Long will it Take to Lift One Billion People out of Poverty?

Martin Ravallion¹

Alternative scenarios are considered for reducing by one billion the number of people living below \$1.25 a day. The low-case, “pessimistic,” path to that goal would see the developing world outside China returning to its slower pace of growth and poverty reduction of the 1980s and 1990s, though with China maintaining its progress. This path would take another 50 years or more to lift one billion people out of poverty. The more optimistic path would maintain the (impressive) progress against poverty since 2000, which would instead reach the target by around 2025-30. This scenario is consistent with both linear projections of the time series data and non-linear simulations of inequality-neutral growth for the developing world as a whole.

Keywords: Development goals, global poverty, growth, inequality, targets

JEL: I32, O15

At the time of writing (in 2012) the available data indicate that 1.2 billion people in the world live in poverty, as judged by the frugal standards used for defining poverty in the world’s poorest countries. This paper draws on recent research on global poverty to assess how long it might take to lift one billion people out of such extreme poverty. The paper identifies both “pessimistic” and “optimistic” paths to that target. It is not claimed that extreme poverty, as conventionally defined in terms of command over commodities, is the only development goal of interest, though it is clearly a very important one.

This is not a standard forecasting problem but rather a task of “benchmarking” future progress. International development institutions and almost all developing countries track poverty measures, and the results are keenly watched. But less analytic attention is given to that

¹ The author is the Edmond D. Villani Professor of Economics at Georgetown University; his email address is mr1185@georgetown.edu. At the time of writing, the author was Director of the World Bank’s research department. The author is grateful to Jim Yong Kim for encouraging him to think about this issue and for helpful discussions with Kaushik Basu, Francois Bourguignon, Shaohua Chen, Stefan Dercon, Shanta Devarajan, Francisco Ferreira, Hiroki Uematsu, Sri Mulyani Indrawati, Peter Lanjouw, Ernesto May, Johan Mistiaen, Mahmoud Mohieldin, David Rosenblatt, Jaime Saavedra, Warrick Smith, Hans Timmer, Jos Verbeek, Nobuo Yoshida, Dominique van de Walle and seminar participants at the Overseas Development Institute, London, the UK Department of International Development, London, and the World Bank.

task of benchmarking performance.² When do we say that performance is “good” or “bad”? Setting goals can motivate achieving them. The *Millennium Development Goals* were clearly conceived with that aim in mind. As Hume (2009, p.4) puts it, the MDGs aimed to “stretch ambitions and mobilize political commitment and public support.” (The MDGs were not the first such effort at goal setting; Hume identifies a number of antecedents including the UN Declaration of Human Rights.) The first Millennium Development Goal (MDG1) was to halve the developing world’s 1990 “extreme poverty rate” by 2015. Using the \$1.25 a day poverty line at 2005 purchasing power parity (PPP) for defining “extreme poverty,” MDG1 was attained in 2010, five year ahead of the goal (Chen and Ravallion, 2012). However, that still leaves over one billion people living in extreme poverty.

At the time MDG1 was set, there were very few time series observations to draw on for benchmarking performance. That has changed dramatically. While the data are still far from ideal, we now have a firmer empirical basis for considering alternative goals. The benchmarks proposed here rest on explicit scenarios for future growth and distributional change, informed by knowledge of recent past performance against poverty and current expectations about growth prospects across the developing world. While there must (inevitably) be a large dose of conjecture in setting targets, it is hoped that this effort will help inform public discussions on development goals, guide assessments of performance over the coming years, and help mobilize efforts against poverty.

The first of the two main benchmark trajectories proposed here assumes that the faster pace of poverty reduction we have seen in the developing world since the 1990s is not maintained except for China. Rather, a series of policy reversals and crises entail that the developing world outside China returns to the slower pace seen in the 1980s and 1990s. On this trajectory it would take 50 years to lift one billion people out of poverty.

One can speculate on potentially many more optimistic paths. At one extreme one might imagine that that poverty could be eliminated tomorrow, and for ever after, by using perfectly targeted transfers to close the aggregate poverty gap—bringing everyone up to \$1.25 a day. However, that has never happened, and appears unlikely to ever happen (including because of

² The only careful attempt to benchmark performance at country level appears to be Newman et al. (2010) who draw on time series evidence across developing countries. They use the empirical distribution of absolute changes in the poverty rate as data for informing the setting of performance benchmarks for countries in Latin America and the Caribbean.

the likely incentive effects of perfect targeting, which would create 100% marginal tax rates on poor people). Here the focus is on an optimistic trajectory that is within the range of past experience, rather than well outside that range.

The optimistic benchmark assumes that the developing world's recent trajectory of success against poverty is maintained. This would entail that China continues on its path toward eliminating extreme poverty, and that the developing world outside China is able to stay on the new path of more rapid poverty reduction that we have seen since 2000. This trajectory can be expected to lift one billion people out of extreme poverty by 2025-30. Such an ambitious trajectory would require that the developing world is successful along all relevant dimensions, notably in assuring continued, reasonably rapid, pro-poor growth and in avoiding major crises (whether financial and agro-climatic)—in short, pro-poor policies (including protection from shocks) plus a measure of good luck in avoiding crises.

After reviewing what we know about recent growth, distributional change and poverty reduction in the developing world, the paper discusses the proposed benchmark paths for poverty reduction. Two main approaches are used, one based on time series evidence and one using simulations. The final section concludes.

Recent trends in growth, redistribution and poverty reduction

The measure of poverty obtained for a given distribution of consumption (or income) depends on the mean of that distribution (relative to the poverty line) and the extent of “inequality” in the distribution. Strictly speaking, it is not “inequality” as normally defined, but rather a more complex function of relative distribution, as discussed in Datt and Ravallion (1992). However, for the present purposes it is defensible to refer to this as “inequality.”

Consider the mean. For global poverty monitoring, Chen and Ravallion (2012) provide estimates for “reference years” spanning 1981-2008 at three yearly intervals, updated to include 2010 for this paper. These estimates cannot be based solely on the surveys, which do not line up in time. Chen and Ravallion (2010, 2012) use an interpolation method that also uses growth rates in private consumption from the national accounts (NAS) when surveys are unavailable. So the means implicit in the poverty measures are essentially a hybrid of surveys means (when available) and NAS data. The means are also a mixture of consumption and income, depending on the surveys used. For the Chen-Ravallion poverty measures, two-thirds of the surveys used

consumption, which is taken to be the preferred welfare metric. When consumption is not available, income is used instead.

On calculating the overall survey mean for each of the Chen-Ravallion reference years used by, one finds that the trend rate of growth over 1981-2010 was 1.9% per annum with a standard error of 0.3%. (The trend rate of growth is defined here as the regression coefficient of the log mean on time.) However, there is a marked break in the trajectory around the year 2000. The trend growth rates prior to 1999 was 0.9% per annum (standard error=0.1) while it was 4.3% (0.2%) from 1999 onwards.³ (The difference in trends is statistically significant; $t=3.32$; $\text{prob.}=0.009$.) It is also notable that this break reflects the higher growth rate of the developing world outside China, which rose from 0.6% (0.1%) prior to 1999 to 3.8% (0.3%) from 1999 onwards.

Using nationally representative surveys, the evolution of mean household consumption will naturally depend in part on national income per capita, as measured in the NAS. As is well known, the developing world as a whole has been maintaining a growth rate for GDP of around 6% over most of the last decade, though dipping substantially (and temporarily) in 2008-09 due to the global financial crisis. This is a full 2% points higher than the average growth rate of about 4% from the 1960s through to the mid-1990s.

At the time of writing, the World Bank's current growth projections assume that a full-blown Euro crisis will be avoided and the annual growth rate of GDP for the developing world will still be 6% in the coming few years (having fallen slightly this year) (World Bank, 2012b). Given current population projections, a 6% growth rate in the GDP over the coming decade would represent a 4.9% rate for GDP per capita. (The compound rate of population growth over 2012-17 and 2012-22 is 1.1% per annum.) While the recent growth has not been even across all regions, the three regions that account for the bulk of absolute poverty—East Asia, South Asia and Sub-Saharan Africa (SSA)—have also seen strong GDP growth rates in recent years, namely around 8% in East Asia, 7% in South Asia and 5% in SSA. Current expectations are that these rates will be maintained, though serious risks are faced that further crises emanating in the “rich

³ Two further tests were used. First, the calculation was repeated dropping the 2010 observation, given that weaker survey coverage means that the estimate is more affected by NAS consumption growth rates used for interpolation. However, the growth rate for the recent period was almost identical (4.3% with standard error of 0.4%). Second, these calculations do not constrain the predicted values of the log mean to be identical in 1999. If one prefers to impose that constraint then the trend growth rates for the two sub-periods are slightly lower, at 0.7% per annum (0.1%) and 4.1% (0.2%) respectively.

world” will spill over significantly into growth and poverty reduction in the developing world, as discussed in World Bank (2012a,b).

However, one should not presume that growth as measured in the NAS will automatically be translated into similar growth in the mean for the distribution of household consumption or income on which poverty measures are based. Indeed, at least prior to 2000, the above observations suggest sizable divergence between the two series. There is invariably a gap between mean consumption or income measured from the household surveys used to measure poverty and these national accounts aggregates. And the gap is also found between the level of mean household consumption, as measured from surveys, and the private consumption component in the NAS. There are a number of reasons. The way the NAS are constructed in practice means that there is nothing exactly corresponding to household consumption as measured in surveys, so full agreement should never be expected. Measurement errors in both sources also play a role, as do differences in accounting periods and sampling problems. An important source of discrepancies between the two data sources is under-reporting of incomes or consumptions in sample surveys, or selective compliance in the randomized assignments used in implementing the surveys. These survey problems are unlikely to leave inequality unchanged. The rich will undoubtedly have a stronger incentive than the poor to under-report their incomes or consumptions, and also be less likely to be available for interview. Indeed, in one assessment (using data for the US) it was found that selective compliance—whereby the rich are less likely to participate in surveys than the poor—resulted in a sizeable under-estimation of inequality using the unadjusted sample survey data but made little difference to poverty measures (Korinek et al., 2006).

What about inequality? The concept of “global inequality” relevant to measuring global poverty pools all residents of all countries, and measures the inequality amongst them as if it was one country. Figure 1 plots one such a measure of inequality, namely the Mean Log Deviation (MLD). This is a theoretically sound measure with the useful property of exact decomposability by population subgroups (Bourguignon, 1979). Thus we can cleanly separate the “between-country” component of total inequality from the “within-country” component. Figure 1 also gives this breakdown of total inequality in the developing world.

We see that there has been a trend decrease in total inequality; over the period as a whole there is a small but statistically significant negative trend, at -0.002 per year (with a standard

error of 0.001), as compared to a mean MLD of 0.57. However, the bulk of the decline in overall inequality was in the period up to the late 1990s. There is a sign of rising overall inequality since 2005. (Note that these are aggregate indices for the developing world as a whole. Different regions have seen different patterns in how average inequality across countries has been evolving; Ravallion and Chen, 2012, provide details by regions.)

The future evolution of overall inequality will be crucial to the trajectories for overall poverty measures. It might be conjectured that higher growth will (at least initially) put upward pressure on inequality within low- and possibly middle-income developing countries, as predicted by the famous Kuznets Hypothesis (Kuznets, 1955). However, that conjecture is not consistent with the evidence, which indicates that inequality within growing developing countries falls about as often as it rises (Ravallion, 2001; Ferreira and Ravallion, 2009). A number of high-inequality growing developing countries have succeeded in attenuating and even reducing inequality. The evidence leads one to doubt that higher inequality is simply the “price” for higher growth and lower absolute poverty (Ravallion, 2005).

However, as is evident from Figure 1, inequality between countries matters more to the evolution of total inequality. Recent economic growth in India and (especially) China has played an important role in the evolution of the between-country component. Given their initially low average incomes, that growth has been a force for inequality reduction globally. That will change in due course, when these countries reach the overall mean for the developing world. (This has not yet happened, but it will soon for China where mean consumption in 2010 was 95% of the mean for the developing world as a whole of \$5.03 per person per day.) A continuing improvement in Africa’s growth performance would work in the opposite direction, putting downward pressure on overall inequality.

Turning now to poverty, following past practice, “extreme poverty” is defined here as being poor by the standards used to define “poverty” in the poorest countries of the world. On that basis the World Bank’s current international poverty line is \$1.25 per person per day at 2005 Purchasing Power Parity (PPP) (Ravallion et al., 2009).

The latest (survey-based) estimate of the proportion of the population of the developing world living below \$1.25 a day is 21% for 2010 (representing 1.2 billion people), down from

33% in 2000 and 43% in 1990; Figure 2 plots the full series over time.⁴ The overall poverty rate (“headcount index”) has been falling at a robust 1% point per year over 1981-2010. The regression coefficient of the poverty rate on the year in Figure 2 is -1.04 with a standard error of 0.05 (n=11). This implies a “\$1.25 a day” poverty rate for the developing world as a whole in the year of writing (2012) of 19% (with a standard error of 0.6%), representing 1.1 billion people.

The extent of the linearity in time is a striking feature of the series for the headcount index in Figure 2. This is plain looking at the graph, but it is confirmed statistically using a standard test for nonlinearity in the function form.⁵ This linearity is not what one would expect if the overall growth rate in the mean was constant and the elasticity of the poverty rate to the mean also remained constant; then one would expect the annual rate of decline in the poverty rate to fall over time as the poverty rate fell. However, these conditions have plainly not held. Instead, we have seen higher growth rates over time in developing countries and (for all except the last five years) falling overall inequality, which would tend to increase the elasticity of poverty to growth (Ravallion, 1997). It seems that these forces have been strong enough to keep the pace of poverty reduction roughly constant.

Figure 2 also gives the series for the poverty gap (PG) index. This is the mean gap below the poverty line, as a proportion of the line, expressed as a percentage. Equivalently, the PG index is the product of the headcount index and the income gap-ratio, given by the difference between the poverty line and the mean consumption or income of those living below the line, expressed as a percentage of the line. The PG index fell from 21% to 6% over the period 1981-2010; the regression coefficient of the index on year is -0.46 with a standard error of 0.04 (n=11). The income gap ratio also fell, from 0.41 to 0.31.

Progress against poverty has been uneven across regions. Table 1 gives the regional breakdown for selected years. Figure 3 gives the series for three regions accounting for 95% of those living below \$1.25 a day (in 2008). For East Asia, the rate of poverty reduction has been about double the average, at slightly over 2% points per year; for South Asia, it has been slightly less than 1%, and for Sub-Saharan Africa it has been about zero over the period as a whole. (The

⁴ The estimates of absolute poverty measures used here are from Chen and Ravallion (2012), which discusses data and methods. (The update for 2010 gives 20.8% below \$1.25 a day, representing 1227 million people.) A fuller discussion of the absolute measures (with various tests of robustness) can be found in Chen and Ravallion (2010).

⁵ The Ramsey RESET test using the squared fitted values indicated that one could not reject the null hypothesis of linearity; the t-statistic for the squared fitted values was -0.83. Nor could the null be rejected on adding a cubed term as well.

coefficients are -2.24%, with a standard error of 0.14, for East Asia, -0.88% (0.04) for South Asia and -0.09% (0.14) for Africa.)

China's success against absolute poverty has (understandably) attracted much attention. (For further discussion, including on the reasons for this success, see Ravallion and Chen, 2007.) However, less widely appreciated is the fact that the developing world outside China has moved to a steeper trajectory of poverty reduction since 2000 (in keeping with the survey means). There have been clear signs of a positive trend emerging in Africa in the 2000s, at a similar trend to South Asia. Chen and Ravallion (2012) show that the 2000s have seen progress against poverty—with both falling incidence and falling numbers of poor—in all six regions of the developing world (EECA, LAC and MENA, in addition to those in Figure 3).

If we focus on the developing world outside China the rate of poverty reduction (again using the regression coefficient on time and the headcount index for \$1.25 a day) over the period 1981-2010 was -0.6% points per year (standard error of 0.05; n=11). It was China's success that pulled the overall rate up to 1% point per year. However, there was a marked change in the trajectory for the developing world outside China around 2000, as can be seen in Figure 4. For the period up to the late-1990s, the trend rate of poverty reduction for the developing world outside China was -0.4 percentage points per year (with a standard error of 0.03; n=6). For the period from 1999 onwards it rose (in absolute value) to -1.0 percentage points per year (standard error=0.04; n=5).⁶ The difference in trends is significant (t=10.01, prob<0.00005). If the pre-2000 trajectory had continued then 29.6% of the population of the developing world outside China would have lived below \$1.25 a day in 2012 instead of 23.4%—representing an extra 280 million people who would otherwise lived below \$1.25 a day.

The clear signs of a break around 2000 in Figure 4 might be taken to suggest that the MDGs have helped, given that they were officially ratified that year at the *Millennium Summit*. It is also notable that we would expect the MDGs to have had their impact outside China, notably in Africa. However, attribution to the MDGs would have to await a fuller analysis, considering other possible explanations for the break in trajectories around 2000.

⁶ If one constrains the predicted values of the poverty rate to be identical in 1999 then the trend annualized rates of change become -0.4% per annum (0.02%) and -0.9% (0.04%).

Pessimistic and optimistic trajectories

In the following analysis, whether one is “optimistic” or not about future progress against poverty is seen to depend on whether one thinks that the new path for the poverty rate seen in the developing world outside China since 2000 will be maintained. Given current population projections imply that lifting one billion people out of poverty, measured against \$1.25 a day, would require a poverty rate of 3% (to the nearest integer).⁷ In describing the two paths to that goal the discussion will focus on two dates, 2022 (10 years from the time of writing), 2030 and the date at which the 3% target is reached.

The low-case trajectory assumes that it will not—that the developing world outside China returns to its pre-2000 pace of poverty reduction from 2012 onwards, though China remains on track. Projecting the series forward linearly from 2012, but at the pre-1999 rate of poverty reduction, this implies a poverty rate of 17.7% for the developing world outside China in 2022, and 14.7% by 2030.⁸ China’s poverty rate would have fallen to zero. So, given that current population projections indicate that China’s share of the developing world’s population will be 21% by 2022 (and 20% by 2030), the overall poverty rate would fall to 14.0% by that year and to 11.8% by 2030. The number of poor would fall from 1.1 billion in 2012 to 0.9 billion in 2022 and 0.8 billion in 2030. On this trajectory it would not be until 2060 that one billion people would be lifted out of poverty.

The guiding principle for the optimistic benchmark is that the recent success against extreme poverty is maintained. Two main ways are used here to quantify such a benchmark. The first is based on a time series projection of past experience. The next section turns to a second method using simulations.

A simple linear projection of the series in Figure 2 gives a poverty rate for 2022 of 8.6% (standard error = 1.0%), or 0.6 billion people, falling to 3.4% (1.2%), or 0.2 billion people, by 2027. The 95% confidence interval for the 2027 poverty rate is (1.0%, 5.8%). The precise date

⁷ The required poverty rate falls from 3.3% in 2012 to 2.8% in 2030. The Bank’s current population projections in millions for the developing world, as defined in 2012, are 5884 in 2012, 6230 in 2017, 6561 in 2022 and 7048 by 2030. Using the 1990 classification of developing countries (to be consistent with *PovcalNet*) the numbers are 6009, 6355, 6687 and 7173 respectively. This difference does not change the poverty counts reported in this paper at the stated number of significant figures.

⁸ The projected poverty rate for the developing world outside China in 2012 (based on the trajectory since 1999) is 21.5% (standard error=0.4%) and the rate of poverty reduction based on the pre-1999 trajectory is -0.38% points per year.

at which the 3% target is expected to be reached is 2027.4 (standard error = 1.16 years) and the 95% confidence for the year at which that target is reached is (2025, 2030).

While this discussion focuses on the headcount index it should be noted that the linear projection is an implausible specification for the PG index. The linear projection implies that the PG index goes to virtually zero by 2022; the projected value is 0.3%, which is not significantly different from zero (standard error=0.97%). A log specification for projecting the income gap-ratio clearly performs better and this gives a projected PG index of 3.9% for 2017, 2.3% for 2022 and 0.9% by 2027.⁹ Notice that these projections (allowing for non-linearity) imply only a modest decline in the income gap ratio under the optimistic trajectory, from 0.31 in 2010 to 0.26 in 2027; the main driving force for the lower PG index is the decline in the headcount index.

The method of linear projection based on the aggregate series for the headcount index in Figure 2 can be questioned as a means of setting the optimistic benchmark. The method may suffer from a form of “aggregation bias.” As we have seen, past progress has been quite uneven across regions (and countries). China has naturally had a high weight and we can reasonably expect the virtual elimination of extreme poverty there by 2022. (Indeed, that is implied by a linear projection of the time series for China and is assumed by the pessimistic scenario above.) Then we should recalculate the projection setting China’s poverty rate to zero in 2022. For 2022, the linear projection for the developing world outside China gives 18.5% living below \$1.25 a day. With a projected 79% of the developing world’s population living outside China by that date, this implies that we would see an overall poverty rate of 14.6%. However, this calculation ignores the clear acceleration we have seen in progress against absolute poverty in the developing world outside China since the mid-1990s. If we project forward from 1999, the poverty rate in the developing world outside China would reach 11.9% by 2022, implying an aggregate rate of 9.4%—only slightly higher than the simple linear projection.

An alternative method of checking for aggregation bias is to use simulation methods that apply country-level growth projections to baseline poverty measures, and then aggregate up to the global level. This is a more complicated calculation. (The base year distribution is projected forward for each country separately. A constant elasticity is not assumed.) However, reassuringly, it is in close accord with the calculation based on aggregate trends, at least in the

⁹ Regressing the log of the income-gap ratio (PG index divided by headcount index) on the year one obtains projections of 0.282 for 2017, 0.271 for 2022 and 0.261 for 2027. The projected PG indices are then obtained by applying these projections to the prior linear projections for the headcount index.

near term. Trajectories for poverty in 2015 by this method are reported in the 2012 *Global Monitoring Report* (GMR), drawing on the Bank's country-specific growth projections. This gives a poverty rate of 16.3% for 2015 (World Bank, 2012a).¹⁰ This is very close to that implied by the first method, namely 15.9%. Table 1 includes the forecasts by region for 2015 from the GMR. As is plain from Table 1, success in fighting poverty in South Asia and Sub-Saharan Africa will be crucial to overall progress; 80% of the projected extreme poverty count for 2015, consistent with the more optimistic scenario above, is in these two regions (about equally).

A further concern is that a constant annual percentage point decline cannot, of course, continue indefinitely—no matter how optimistic one is willing to be, the trajectory must slow down, becoming nonlinear in time. What is less clear is when this would happen. The next section will instead use a version of the simulation method (as described above for addressing aggregation bias) to see if whether the linear trajectory for the overall poverty rate is consistent with the expected growth rates under the optimistic path under explicit assumptions about distribution.

An alternative approach using simulations

We have seen that the linear projection of the time series of poverty measures for the developing world suggests that one billion people could be lifted out of poverty by 2027. However, the above discussion has also pointed to some concerns about whether the linearity will hold that long. This section approaches the problem of quantifying an optimistic trajectory in a rather different way, without relying on the time series evidence, to see whether it ends up at roughly the same conclusion.

The key assumption driving the following simulations of the optimistic path is that the overall level of inequality does not increase going forward. In other words, all income levels across the developing world as a whole grow at the same rate—maintaining overall inequality. As we saw in Figure 1, the period since the late 1990s has seen little change in overall inequality, although new forces may well emerge to put upward pressure on overall inequality and there is a sign of this in recent years. For example, in the near future, China's growth will start to put

¹⁰ These forecasts factor in historical gaps between growth rates in the survey means and growth rates in private consumption per capita in the national accounts. For most countries, about 90% of the national accounts growth rate is passed onto the survey means, but for India it was only about half, consistent with the larger gap between the two growth rates for India. In addition an allowance is made for rising inequality in both China and India. For further details see World Bank (2008).

upward pressure on overall inequality. Whether Africa can reach the average growth rates of the rest of the developing world will be key.

To quantify the implications of an inequality-neutral growth process one can use the baseline distribution and project this forward with a higher mean until any given poverty rate is met, and then see what growth rate is required. This automatically takes account of the nonlinearity in how the cumulative distribution function of consumption varies with the poverty line relative to the mean. (That nonlinearity can hold and yet the path of the poverty measures over time can come out to a roughly linear trajectory.) The question we ask is thus: will a continuation of the growth rates seen in recent times in the developing world be enough to lift one billion people out of poverty without a change in the relative distribution of income?

The results of this exercise are found in Table 2, Column (1), which gives the growth rate in household consumption per capita needed to bring the \$1.25 a day poverty rate down to various levels by 2027 assuming no increase (or decrease) in inequality within the developing world. The calculations in column (1) are anchored to the 2008 distribution—the last year for which the whole distribution can reliably be constructed from survey data. For each target poverty rate, one solves backwards to find the growth rate needed to reach that target, holding constant the 2008 Lorenz curve. Note that the accuracy of poverty estimates at low levels is questionable for those countries that use income, rather than consumption, as the welfare indicator. The data base in *PovcalNet* prefers consumption when it is available, which is for about two-thirds of the surveys. The only region that still uses income surveys extensively is Latin America.

This indicates that, without any increase in inequality, a growth rate of 4.5% in mean household consumption per capita—very close to the trend rate of growth in the mean of 4.3% since 1999—would get us close to the 3% poverty rate in 2027. This simulation suggests that the optimistic target of lifting about one billion people out of poverty by 2030 is attainable with a continuation of the relatively higher growth rates seen in the developing world in the last 15 years or so, as long as this does not come with rising inequality.

The target could also be comfortably achieved by 2030 even if recent growth is not maintained but overall inequality falls. And the drop in overall inequality would be within the range of recent experience. To illustrate the magnitudes, Column (2) of Table 2 gives the growth rates needed to attain each target poverty rate in 2027 but this time using the relative distribution

of 1999, which was the year of lowest total inequality in the series in Figure 1, with an inequality index of 0.52, as compared to 0.57 in 2008. This lower level of inequality will entail lower poverty but also allow a more pro-poor growth path. We now see that if mean consumption grew at 3.4% (with no other change in inequality) we would reach a 3% poverty rate by 2027.

Naturally, a higher growth rate would attain the target ahead of that date. For example, an extra 1.2% points on the growth rate at the lower inequality path of column (2) would lift the one billion people out of poverty five years ahead of time, by 2022. At the higher inequality path of column (1), this would require an extra 1.6% points.

Conclusions

The developing world as a whole has seen huge progress against absolute poverty. China has contributed greatly to that progress, but the recent progress goes well beyond China. While the developing world outside China was seeing slow progress against poverty until around the year 2000, it has done much better since then.

In thinking about how we might judge future progress, a natural choice for a low-case, “pessimistic,” trajectory is to assume that we return to the slow progress of the 1980s and ‘90s in the developing world outside China. This trajectory implies that the proportion of the developing world’s population living below \$1.25 a day would fall to 14% over the coming decade, from 19% in 2012. On this path it would take another 50 years or so to lift one billion people out of poverty. This would surely be judged a bad performance. The paper has also tried to suggest what good performance might look like.

To set goals that can motivate extra effort to do better than this pessimistic scenario, they must represent real progress in reducing the incidence of extreme poverty below this trajectory, but they cannot be so far outside the range of experience as to be considered impossible. The proposed “staying on the path” benchmark for the \$1.25 poverty rate in 10 years time is 9%. This would clearly be a very good performance—maintaining an impressive trajectory of poverty reduction going back to the early 1980s when the corresponding poverty rate was over 50%. In the 20 years from 1990 to 2010 the developing world halved its overall poverty rate, from 43% to 21%. Under this trajectory it would be halved again in a further 10 years.

The same calculations for the optimistic trajectory suggest that we can be confident that the target of lifting one billion people out of poverty (relative to the count of the number of poor

in 2010) would be reached by 2025-30, with 2027 as the most likely date. However, this assumes that the robust linear path we have seen for the poverty rate over time will be maintained. That will not be easy. Instead, it might be expected that the pace of poverty reduction will start to decline at low levels, making it harder to reach the target. From what we know, we cannot be confident now about when such a slow down might be expected.

The paper has also provided simulations that try to see what combinations of growth and distributional changes might allow the optimistic trajectory to be attained. The simulations suggest that a continuation of the recent growth performance of the developing world as a whole can be expected to lift one billion people out of poverty by 2027 provided that there was no further deterioration in overall inequality. While overall inequality in the developing world has been fairly stable since the 1990s, there have been signs of a recent rise in inequality. If this continues, then even higher growth rates than those seen since 2000 will be needed to reach the target. By contrast, lower inequality could allow this target to be reached with lower growth than we have seen in recent times, by directly reducing poverty and making subsequent growth more pro-poor.

In choosing amongst the multiple solutions for lifting one billion people out of poverty, the sustainability of poverty reduction efforts is clearly important. We do not want to reach the target only to fall back in the subsequent years. On an encouraging note, recent research has suggested that lower initial levels of absolute poverty foster higher subsequent rates of growth average living standards in developing countries, and also help assure that the growth is more poverty reducing (Ravallion, 2012). This suggests that a “virtuous cycle” could be anticipated that would help assure sustainability of the reduction in poverty.

Environmental sustainability is also important. A growth path that consumed all of a country’s natural resources in 15 years may well attain the target, but the poverty rate will soon bounce back. However, existing measures of poverty do not tell us whether the changes observed are sustainable—they relate only to the “current” time period, which might be quite short. That suggests that we would need a separate check on the sustainability of observed poverty reduction. The best data we currently have for that check would appear to be the “adjusted net savings” rate based on Hamilton and Clemens (1999). (The latest estimates by country are found in World Bank, 2012c, Table 4.11.) The sustainability of a reduction in poverty that came with reduced (or negative) adjusted net savings would be questionable.

The best sustainable route will naturally vary from country to country. The growth projections underlying these benchmark trajectories are grounded in the economic realities in the countries concerned and the global economy. However, the policy challenges in assuring that poor households share sufficiently and sustainably in that growth at country level will need to be addressed in depth.

The bulk of the work that would be needed to reach such global targets will be done at country level. A similar benchmarking exercise for individual countries would be desirable. Naturally this would reflect the specifics of each country. One should be cautious in trying to infer what any one country can do by cross-country comparisons, even when these control for observable exogenous differences. It should ideally also come with a reasonably clear plan of how the targets would be achieved. Various tools of economic analysis exist that can inform such plans, ranging from computable general equilibrium models to micro-simulations tools, all with both their strengths and weaknesses.¹¹

Monitoring performance against these benchmarks poses a number of serious data challenges. There has been huge progress in collecting the primary household survey data. When the World Bank's current global poverty monitoring effort began in 1990 the estimates used 22 surveys for 22 countries (Ravallion, et al., 1991). Today around 900 surveys are used, spanning 125 countries—over six per country; the latest estimates use a “global” sample of 2.1 million households. However, many problems remain. There are persistent lags and uneven coverage. (The surveys used here cover 90% of the population of the developing world as a whole in 2008, but this varies from 94% in East Asia to only 50% in Middle-East and North Africa.) There are continuing concerns about the comparability of the surveys over time and across countries. And there are concerns about under-reporting and selective compliance in household surveys; the rich are hard to interview, and that task is not getting any easier. The weak integration of “macro” and “micro” data is also a long-standing concern, warranting more attention than it has received. Better data will help determine how close we are to reaching our development goals, including poverty reduction, and in assessing the efficacy of alternative policies for attaining those goals.

¹¹ A useful compendium of the tools available can be found in Bourguignon et al. (2008). On microeconomic simulation methods see Ferreira and Leite (2003).

References

- Bourguignon, Francois. 1979. "Decomposable Income Inequality Measures," *Econometrica* 47(4): 901-920.
- Bourguignon, Francois, Maurizio Bussolo and Luiz Pereira da Silva (eds). 2008. *The Impact of Macroeconomic Policies on Poverty and Income Distribution. Macr-Micro Evaluation Techniques and Tools*. Washington DC., Palgrave Macmillan and the World Bank.
- Chen, Shaohua and Martin Ravallion. 2010. "The Developing World is Poorer than we Thought, but no Less Successful in the Fight Against Poverty," *Quarterly Journal of Economics* 125(4): 1577-1625.
- Chen Shaohua and Martin Ravallion. 2012. "More Relatively-Poor People in a Less Absolutely-Poor World," [Policy Research Working Paper 6114](#), World Bank.
- Datt, Gaurav and Marti Ravallion. 1992. "Growth and Redistribution Components of Changes in Poverty: A Decomposition with Application to Brazil and India," *Journal of Development Economics* 38: 275-295.
- Ferreira, Francisco and Phillippe Leite. 2003. "Meeting the Millennium Development Goals in Brazil: Can Microeconomic Simulations Help?," *Economica*, Spring: 235-279.
- Ferreira, Francisco and Martin Ravallion. 2009. "Poverty and Inequality: The Global Context", in Wiemer Salverda, Brian Nolan and Tim Smeeding (eds), *The Oxford Handbook of Economic Inequality*, Oxford: Oxford University Press.
- Hamilton, Kirk and Michael Clemens. 1999. "Genuine Savings Rates in Developing Countries," *World Bank Economic Review* 13(2): 333-356.
- Hume, David. 2009. "The Millennium Development Goals (MDGs): A Short History of the World's Biggest Promise," BWPI Working Paper 100, Brooks Institute, University of Manchester.
- Korinek, Anton, Johan Mistiaen and Martin Ravallion. 2006. "Survey Nonresponse and the Distribution of Income," *Journal of Economic Inequality* 4(2): 33-55.
- Kuznets, Simon. 1955. "Economic growth and income inequality," *American Economic Review* 45:1-28.
- Lustig, Nora and Luis Felipe López-Calva. 2010. *Declining Inequality in Latin America: A Decade of Progress?* Brookings Institution, Washington DC.

- Newman, John L., Joao P. Azevedo, Jaime Savedra and Ezequiel Molina. 2010. "The Real Bottom Line: Benchmarking Performance in Poverty Reduction in Latin America and the Caribbean," mimeo, World Bank, Washington DC.
- Ravallion, Martin. 1997. "Can High Inequality Developing Countries Escape Absolute Poverty?," *Economics Letters* 56: 51-57.
- _____. 2001. "Growth, Inequality and Poverty: Looking Beyond Averages," *World Development* 29(11): 1803-1815.
- _____. 2005. "A Poverty-Inequality Trade-Off?" *Journal of Economic Inequality* 3(2): 169-182.
- _____. 2012. "Why Don't We See Poverty Convergence?" *American Economic Review* 102(1): 504-523.
- Ravallion, Martin, Gaurav Datt and Dominique van de Walle. 1991. "Quantifying Absolute Poverty in the Developing World," *Review of Income and Wealth*, 37(4): 345-361.
- Ravallion, Martin and Shaohua Chen. 2007, "China's (Uneven) Progress Against Poverty," *Journal of Development Economics* 82(1): 1-42.
- _____ and _____. 2012. "[Monitoring Inequality](#)," Blog Post on Let's Talk Development, World Bank.
- Ravallion, Martin, Shaohua Chen and Prem Sangraula. 2009. "Dollar a Day Revisited," *World Bank Economic Review* 23(2): 163-184.
- World Bank. 2008. *Global Economic Prospects*. Development Economics, World Bank.
- _____. 2012a. *Global Monitoring Report 2012*. Development Economics, World Bank, in collaboration with the International Monetary Fund.
- _____. 2012b. *Global Economic Prospects June 2012*. Development Economics, World Bank.
- _____. 2012c. *World Development Indicators 2012*, World Bank.

Table 1: Regional breakdown of poverty incidence and the projection for 2015

	1990	1999	2008	2015
Poverty rate (% of the population living below \$1.25 a day)				
East Asia and Pacific	56.2	35.6	14.3	7.7
Europe and Central Asia	1.9	3.8	0.5	0.3
Latin America and Caribbean	12.2	11.9	6.5	5.5
Middle East and North Africa	5.8	5.0	2.7	2.7
South Asia	53.8	45.1	36.0	23.9
Sub-Saharan Africa	56.5	58.0	47.5	41.2
Total	43.1	34.1	22.4	16.3
Number of poor (millions living below \$1.25 a day)				
East Asia and Pacific	926.4	655.6	284.4	159.3
Europe and Central Asia	8.9	17.8	2.2	1.4
Latin America and Caribbean	53.4	60.1	36.8	33.6
Middle East and North Africa	13.0	13.6	8.6	9.7
South Asia	617.3	619.5	570.9	418.7
Sub-Saharan Africa	289.7	376.8	386.0	397.2
Total	1908.6	1743.4	1289.0	1019.9

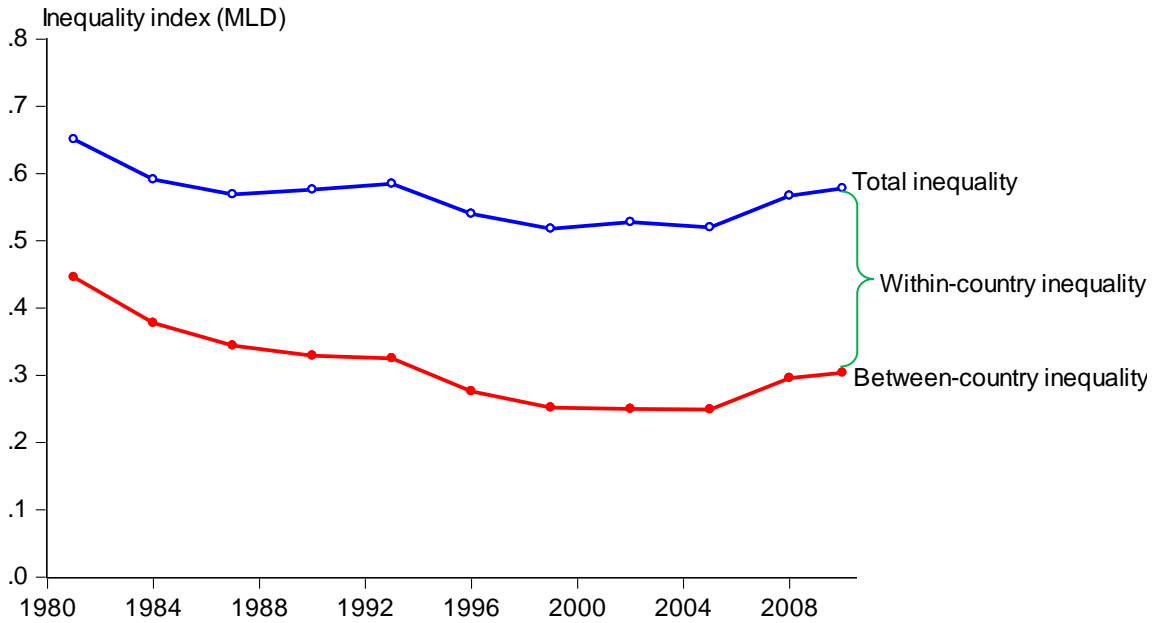
Source: Chen and Ravallion (2010) and World Bank (2012a).

Table 2: Growth rates for the developing world as a whole required for various target poverty rates at fixed level of overall inequality within the developing world

Target poverty rate for 2027 (\$1.25 a day)	(1) Required annual growth rate in household consumption per capita at 2008 level of inequality	(2) Required annual growth rate in household consumption per capita at 1999 level of inequality
1%	7.6%	6.1%
2%	5.5%	4.0%
3%	4.5%	3.4%
4%	3.7%	2.8%
5%	3.3%	2.3%
6%	2.9%	2.0%


Note: These are the required annual growth rates for survey means over the period 2008-2027. The calculations in Col. (1) are based on the distribution of consumption for the developing world in 2008, while those in Col. (2) use anchored to the 2008 mean but using the Lorenz curve for 1999, which gave the lowest overall inequality during the period 1981-2008. For each target poverty rate, the author solved backwards to find the growth rate needed to reach that target holding constant the Lorenz curve at that of either 2008 or 1999. Source: Author's calculations using [PovcalNet](#).

Figure 1: Inequality in the developing world




Source: update of Ravallion and Chen (2012a)

Figure 2: Poverty measures for the developing world 1981-2010




Note: Both measures use a poverty line of \$1.25 a day at 2005 PPP
 Source: Chen and Ravallion (2012) and author's calculations

Figure 3: Poverty rates for the three poorest regions



Source: Chen and Ravallion (2012)

Figure 4: Poverty rates for the developing world outside China



Source: Author's calculations