



THE FINAL COUNTDOWN: PROSPECTS FOR ENDING EXTREME POVERTY BY 2030

LAURENCE CHANDY

Fellow, Development Assistance and Governance Initiative

NATASHA LEDLIE

Research Assistant, Development Assistance and Governance Initiative

VERONIKA PENCIAKOVA

Research Analyst, Development Assistance and Governance Initiative

APRIL 2013

THE BROOKINGS INSTITUTION
1775 MASSACHUSETTS AVE., NW
WASHINGTON, DC 20036

INTRODUCTION

Between 1990 and 2010, the share of the population of the developing world living in extreme poverty was cut in half. This fulfilled the first and foremost Millennium Development Goal (MDG1a) five years ahead of schedule. Moreover, it would appear to bring the end of extreme poverty within reach. If the same rate of global progress is sustained, extreme poverty will be eliminated in 2030. In other words, if it took 20 years to lift half the world's poor above the poverty line, a further 20 years should take care of the other half. This has rightfully got people excited. In the words of Bono, "The zero zone... is just around the corner".¹

What are reasonable expectations for how poverty will unfold over the coming years? How feasible is it for extreme poverty to be eliminated by 2030? And what factors will determine progress toward this goal?

This paper attempts to answer these questions. We develop a number of scenarios based on some carefully derived assumptions to better understand the future trajectory of global poverty. Predicting poverty is a fool's errand so it would be wrong to place too much stock on poverty estimates derived from any one scenario. Nevertheless, we are able to make some reasonably strong statements about the prospects for ending extreme poverty and its sensitivity to different factors.

Extreme poverty is defined here by the narrowly construed, but widely used, consumption-based poverty line of \$1.25 a day measured in 2005 international prices. This is the poverty line used in MDG1a.

In the remainder of this paper, all references to poverty refer to extreme, consumption-based poverty. We are acutely aware that the two are not the same, but drop the prefixes for the sake of brevity.


THE SCENARIOS

The prevalence of poverty in a country can be derived from two pieces of information: the mean level of consumption per person, and the distribution of consumption around that mean (i.e., how consumption levels across the population vary from the average). We develop our baseline scenario based on the best available information for how these two factors might change in each developing country between now and 2030.

Projections for the growth in consumption per person are readily available from the Economist Intelligence Unit. Changes in distribution over time follow no discernible trend and have no established correlates so we assume the distribution of consumption in each country remains unchanged from the most recent available data. Poverty levels are then calculated for each country and aggregated to obtain a global total for each year. (See Technical Appendix for further details).

The baseline scenario is illustrated in Figure 1 as a continuation of official poverty estimates which end in 2008. The global poverty rate falls from 22.4 percent in 2008 to 9.9 percent in 2020, to 5.4 percent in 2030. This brings the total number of people living in poverty down to approximately 386 million by the end of the period.

FIGURE 1: BASELINE SCENARIO



Source: World Bank PovcalNet database and authors' calculations

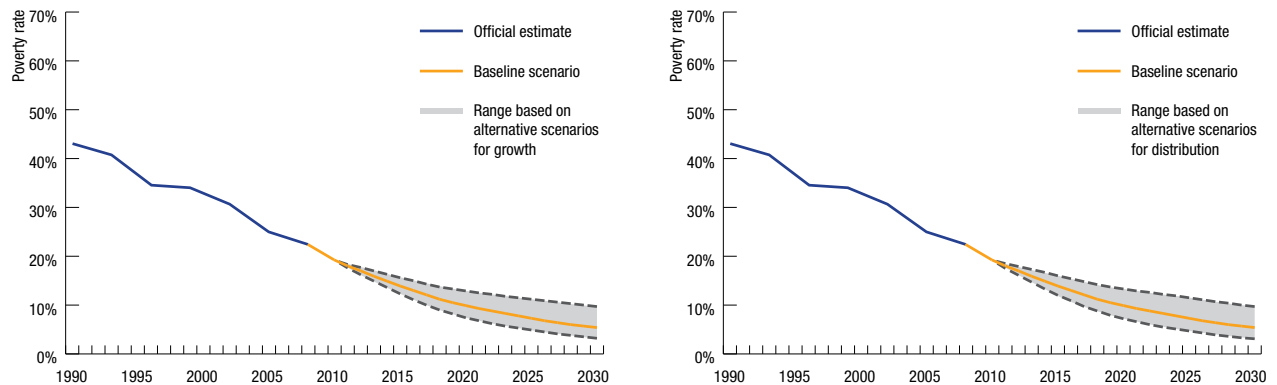
The baseline scenario represents our best bet for the future, but it is not a bet we would be inclined to take. While it is based on the best information we have, our knowledge about the future is necessarily limited. Our baseline assumptions and the poverty estimates which they generate are one of many plausible outcomes.

A more credible approach is to derive a range of plausible poverty outcomes based on alternative scenarios for how mean consumption per person and the distribution of consumption in each country might evolve. This range represents neither the limits of what is possible, nor confidence bands for our baseline, although it is intended to serve as a rough proxy for the latter.

Our alternative scenarios for consumption growth assume that the projections used in our baseline for each country are consistently wrong in one or other direction by 2 percentage points each year. This is based on the differences we observe historically between forecasts and reported data. Maintaining our original assumption regarding the distribution of consumption, these scenarios generate a range of poverty estimates around our baseline (Figure 2). If consumption growth consistently exceeds baseline projections, the global poverty rate

could fall to 3.1 percent by 2030. If instead consumption growth in each country persistently falls short of projections, 9.7 percent of the developing world could remain under the poverty line at the end of the period.

FIGURE 2: POVERTY RANGE BASED ON ALTERNATIVE SCENARIOS FOR CONSUMPTION GROWTH AND DISTRIBUTION, RESPECTIVELY



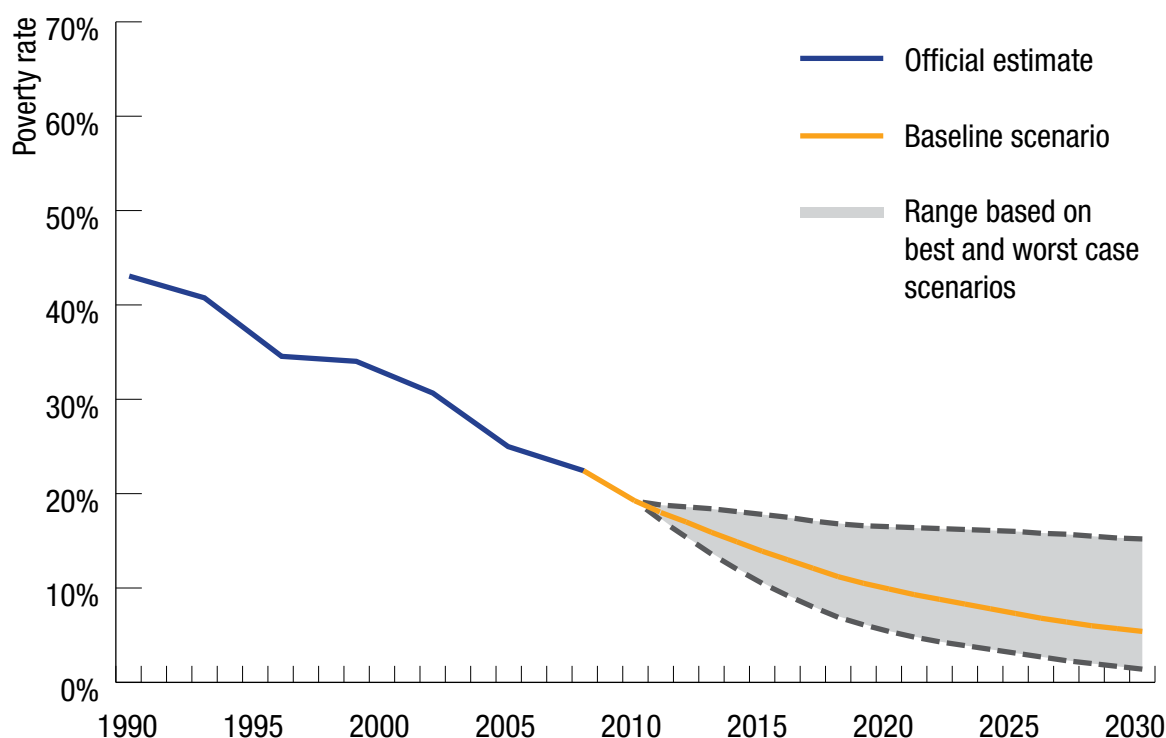
Source: World Bank PovcalNet database and authors' calculations

Modeling changes in the distribution of consumption is complicated as it requires specifying who gains at the expense of whom. One hypothesis which has recently received a lot of attention submits that differences in the distribution of consumption, both across countries and within the same country over time, follow a stylized pattern and can largely be explained by the competition for shares of total consumption in each economy between the poorest 40 percent and the richest 10 percent of the population. We are agnostic about the validity of this claim but find it useful as a basis for modeling distribution changes within each country in the absence of more persuasive alternatives.

For our scenarios, we assume that the shares of total consumption held by the poorest 40 percent and the richest 10 percent in each country steadily shift in favor of one or other group between 2010 and 2030. These shares move upward or downward in increments of a quarter of a percentage point per year, reflecting the scale of long term distribution trends observed in historical data. Figure 2 illustrates the range of poverty estimates that are generated assuming a persistent decline or improvement in the distribution of consumption in each country, while maintaining our baseline projections for consumption growth. Poverty numbers for 2030 range from 3.2 percent under a narrowing distribution to 9.7 percent under a widening distribution.

A final step is to combine our scenarios for how consumption and distribution will change in each economy. This generates a wider range of poverty estimates based on the interaction of changes in consumption and distribution (Figure 3). Under the worst case (low growth, worsening distribution) scenario, poverty in 2030 will remain prevalent at 15.2 percent. Under the best case (high growth, improving distribution) scenario, poverty falls tantalizingly close to zero, reaching 1.4 percent.

FIGURE 3: POVERTY RANGE BASED ON ALTERNATIVE SCENARIOS FOR CONSUMPTION GROWTH AND DISTRIBUTION COMBINED

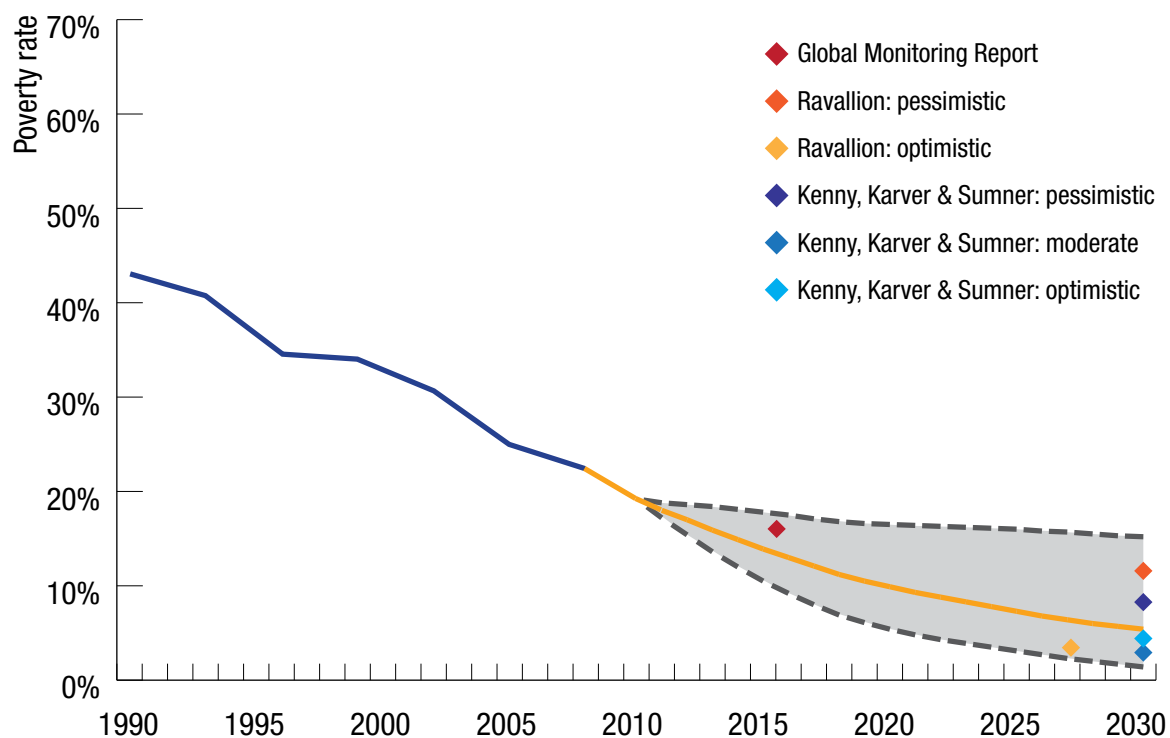


Source: World Bank PovcalNet database and authors' calculations

A few points are worth noting here. First, the range of poverty outcomes for 2030 is large. This confirms our earlier claim that the future trajectory of global poverty is highly uncertain. Indeed, the likelihood of our baseline scenario for 2030 being realized—to within, say, 50 million people—appears remote. Poverty numbers for 2030 are highly sensitive to changes in consumption and distribution between now and then. Second, the range is not symmetric around the baseline: depressed consumption growth and a deteriorating distribution imply greater downside risks (i.e., little poverty reduction) than buoyant growth and an improving distribution imply upside risks. We explore the reasons for this in the next section. Third, getting to the “zero zone”, which is defined here as a poverty rate of less than 3 percent, by 2030 is unlikely to occur through stronger than expected consumption growth or an improving distribution alone. Both factors are needed simultaneously.

We are not the first people to explore how the trajectory of global poverty may unfold. Figure 4 overlays our range with projections or scenarios developed by others. These largely fall within our range giving greater confidence to our findings.

FIGURE 4: COMPARISONS WITH OTHERS



Source: World Bank PovcalNet database; World Bank, 2012; Ravallion, 2013; Kenny, Karver and Sumner, 2012; authors' calculations

COUNTING DOWN TO ZERO

The scenarios described in the previous section are inspiring. Our baseline scenario implies that hundreds of millions of people can escape poverty over the next generation. The asymmetry of our range suggests that such progress is there for the taking; it would be a missed opportunity were this result foregone.

Some, however, will view our scenarios as a disappointment. The ongoing international dialogue over what should replace the MDGs after 2015 has given credence to the idea that poverty can be eliminated in a generation. Our scenarios do not disprove this notion, but the fact that zero poverty does not lie within our range of outcomes for 2030 demonstrates the scale of the challenge.

If this challenge is to be met, it is crucial to understand the factors that limit progress. Only then can an informed effort be made to overcome those factors and to move as close to zero as possible.

As described earlier, poverty will be eliminated if the trend rate of global poverty reduction from the past 20 years—a percentage point a year—is sustained. Figure 1 shows that our baseline scenario keeps up with this trajectory until around 2020, when the line flattens. Even in our best case scenario, progress on global poverty reduction loses ground after 2027.


This may seem counterintuitive. Our baseline scenario entails a record of robust consumption growth and stable distribution across the developing world that is at least equal to the performance of the past 20 years, while our best case scenario is far superior. Why then should progress against poverty slow down?

We unpack this problem in three steps. Each provides a unique and complementary viewpoint for observing the countdown to zero and identifying what is needed to move closer toward that goal.

The view from the poverty line

We begin by looking again at the pattern of poverty reduction but this time in terms of the movement of poor people across the poverty line. The probability density functions in Figure 5 allow such a perspective. In each case, the area under the curve represents the population of the developing world arranged according to individuals' level of daily consumption. The height of the curve at any given consumption level reflects the number of people who meet that exact standard of living. Thus the area under the curve to the left of the vertical poverty line (at \$1.25) captures the number of people living in poverty.

FIGURE 5: PROBABILITY DENSITY FUNCTIONS: 1990, 2000, 2010, 2020 (BASELINE), 2030 (BASELINE)



Source: Authors' calculations

From this viewpoint it is evident that the amount of poverty reduction in any given year is not merely a function of recent growth and distribution changes, but of the precise contours of the curve. Specifically, it is driven by the number of people living just below the poverty line—captured by the height of the curve at the poverty line—in the previous year and whether their consumption improves sufficiently to lift them over the poverty line, either as a result of consumption growth, or a redistribution of consumption toward them, or both.²

The steady rate of poverty reduction achieved over the period of the MDGs, as illustrated in Figure 1, gives the false impression that the factors enabling this progress have themselves been relatively constant. The truth is much more complex.

In the 1990s, progress in poverty reduction was made possible by a consistently large mass of people lining up behind the poverty line. The number of people in the world living between \$1.20 and \$1.25 a day stood at a little over 100 million each year. Average consumption growth per person in the developing world, weighted by the location of poor people, was relatively modest over the decade at 2.3 percent per year.

In the 2000s, a different dynamic emerged. The number of people lined up behind the poverty line began to fall at the start of the new millennium. By 2010, the number of people living between \$1.20 and \$1.25 a day stood at only 85 million. This was compensated for by stronger consumption growth per person, which rose to 3.1 percent per year on average for the decade.

Throughout both decades, the extent to which people immediately below the poverty line shared in their economies' consumption growth, on average, varied from year to year. However, in several years, these “immediate poor” appear to have lost out, resulting in more sluggish poverty reduction than would have been possible under equitable consumption growth (Figure 6).

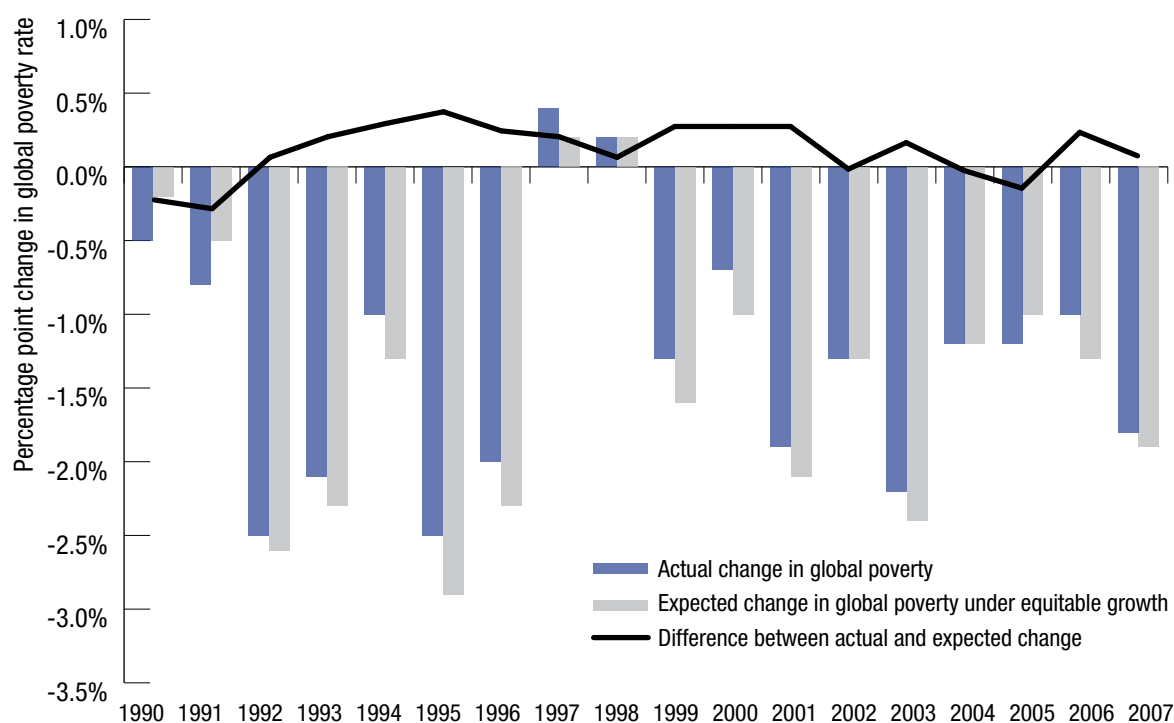
Over the next two decades, the dynamic is expected to change again, with important consequences for the speed of poverty reduction.

Today we stand at a momentary sweet spot, where the mode of consumption per person is at the \$1.25 poverty line. That is, there are more people living around the \$1.25 mark than at any other consumption level in the world, as illustrated by the peak of the probability density function. Equitable growth around the world will result in more movement of people across the poverty line than across any other level of consumption.

However, the height of the curve at the poverty line is diminishing precipitously, meaning that the number of immediate poor is continuing to fall. In our baseline scenario, the number of people living between \$1.20 and \$1.25 a day drops to 56 million in 2020—the year when our baseline falls behind the historical trajectory of poverty reduction—and 28 million in 2030. Despite assumptions of slightly stronger consumption growth per person between now and 2030 (averaging 3.4 percent), and growth being equitable, it is not possible to maintain the trend rate of poverty reduction with so many fewer individuals ready to cross the line.

In our best case scenario, consumption growth is stronger still and that growth is skewed toward the poor, but the number of people positioned just below the poverty line drops even lower. By 2028, when our best case

FIGURE 6: ACTUAL POVERTY REDUCTION VERSUS EXPECTED POVERTY REDUCTION UNDER EQUITABLE GROWTH



Source: Authors' calculations

scenario falls behind the trend trajectory of poverty reduction, the number of people living between \$1.20 and \$1.25 a day is less than 15 million. More than 10 times that number stand further back from the line.


Sustaining the trend rate of poverty reduction requires that each year a new set of individuals is primed to cross the poverty line. When one year's immediate poor crosses the \$1.25 a day threshold, those a little further from the poverty line must move up to take their place. The reduction in the number of immediate poor indicated by our analysis suggests that some of the poorest of the poor will struggle to make enough progress to be in a position to escape poverty over the next 20 years. Even in our best case scenario, many poor people remain some distance behind the poverty line in 2030.

The three giants

To delve deeper into these issues and to explore our scenarios from a different angle, we look beneath the global level to the performance of China, India and sub-Saharan Africa. Together, they account for three-quarters of the world's poor people, within a 3 percentage point margin, in each year since official poverty estimates began. An examination of their record and future prospects can therefore help in understanding the properties of global poverty aggregates.

Figure 7 shows the contribution of the three regions to the annual percentage point reduction in global poverty reduction between 1990 and 2030, combining official estimates with our baseline scenario. This illustrates again how the drivers of progress on poverty reduction are changing. The 40-year period resembles a relay race in which responsibility for leading the charge on global poverty reduction passes from China to India to sub-Saharan Africa.

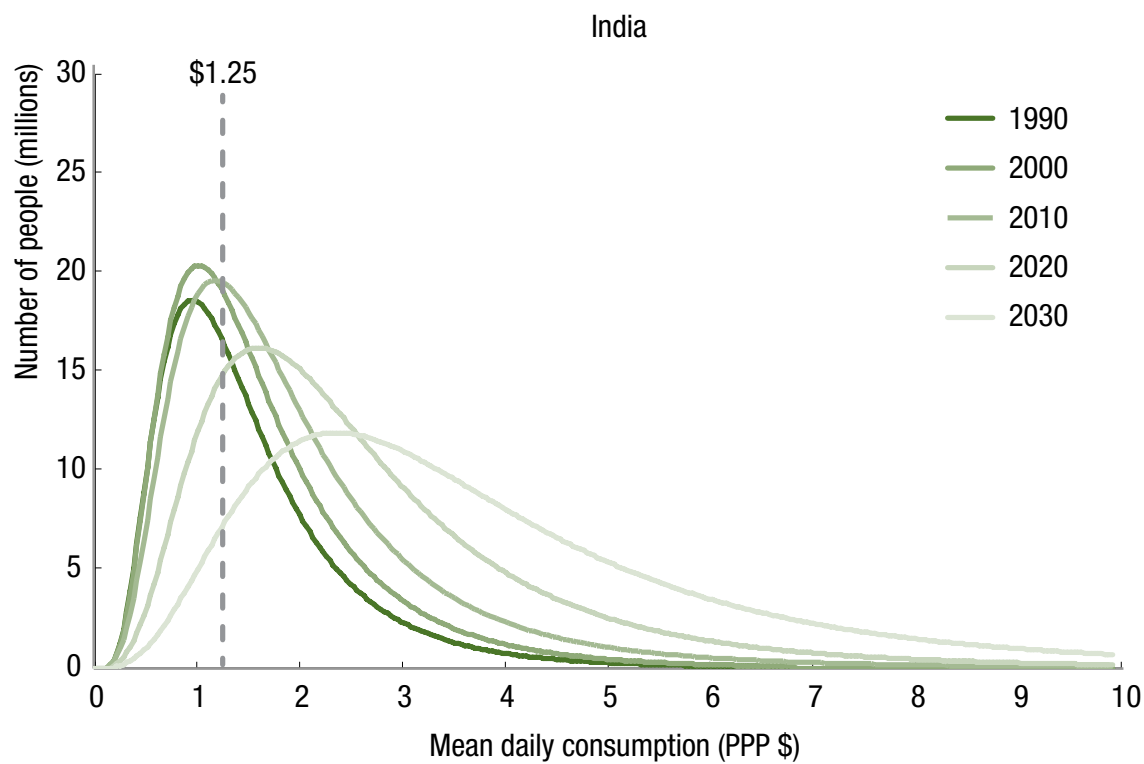
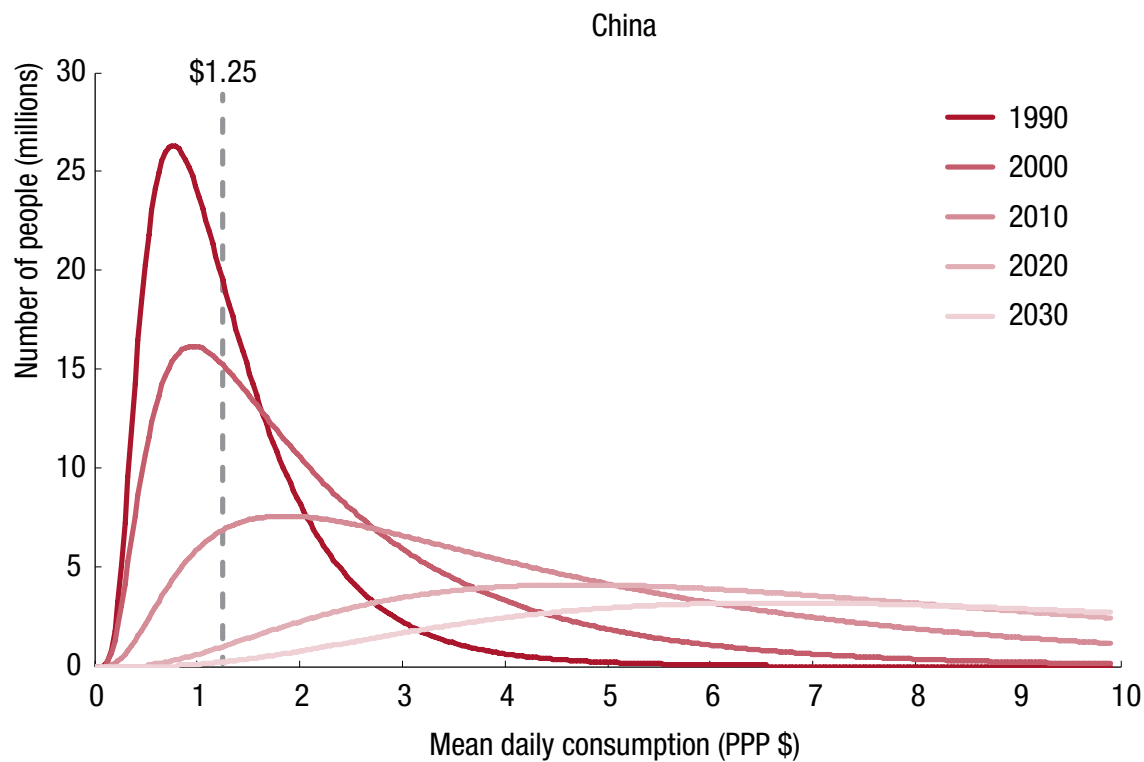
FIGURE 7: AVERAGE ANNUAL PERCENTAGE POINT REDUCTION IN GLOBAL POVERTY BETWEEN 1990 AND 2030 (OFFICIAL ESTIMATES AND BASELINE SCENARIO)




Source: Authors' calculations

This dynamic can be more richly observed using probability density functions (Figure 8). China's relay leg is the most striking. It undergoes a dramatic transformation in which its population, which is initially predominantly poor, disperses over a range of consumption levels beyond the poverty line, driven by rapid, though often inequitable, consumption growth. While the number of people located immediately below the poverty line in China drops off sharply from the mid-1990s onward, the rate of consumption growth is sufficiently fast that its rate of poverty reduction does not let up. In certain years, China alone contributes more than a percentage point reduction in global poverty, singlehandedly keeping the world on its trend path.

FIGURE 8: PROBABILITY DENSITY FUNCTIONS, CHINA, INDIA, SUB-SAHARAN AFRICA:
1990, 2000, 2010, 2020 (BASELINE), 2030 (BASELINE)





Source: Authors' calculations

With China's poverty rate now down in the single digits, the baton has been passed to India. India's mode, like the world's, is situated at the global poverty line. Large numbers of its population are on the precipice of escaping poverty, and others are not far behind, meaning that India has the capacity to deliver sustained progress on global poverty reduction over the next decade based on modest assumptions of equitable growth. India is unlikely to match China's record given its lower forecasts for consumption growth, but it alone can still muster between half and three-quarters of a percentage point reduction in global poverty per year, depending on how its growth and distribution unfold.

Once India's poverty is largely exhausted, it will be up to sub-Saharan Africa to run the final relay leg and bring the baton home. Herein lies the challenge.

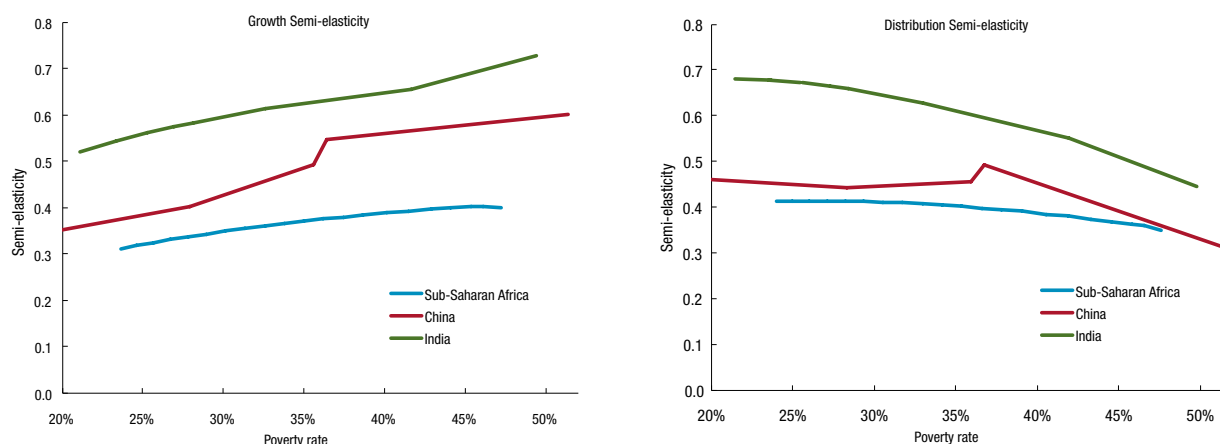
Sub-Saharan Africa has historically struggled in its fight with poverty. In recent years, progress has begun yet the region still contains hundreds of millions of poor people, most of whom are located far behind the poverty line. In our baseline scenario, the region records further progress over the next 20 years, but its mode remains to the left of the \$1.25 a day mark. Much of the movement in its curve occurs to the right hand side of the poverty line, implying that some portion of Africa's population could remain impervious to a further 20 years of global economic progress and integration. In this scenario, the region's poverty rate stands at 23.6 percent in 2030.

Stronger consumption growth and distributional change, especially in Africa's poorest countries, can overcome part of this challenge, but not all of it. In our best case scenario, poor people across the continent share more

fully in the continent's rising prosperity. But given the initial distance of many poor people from the poverty line, 2030 comes too soon for poverty to be eliminated.

Crucially, at no point in any of our scenarios are there as many immediate poor in sub-Saharan Africa as has been recorded by China and India, which would facilitate sharper reductions in poverty. As a result, the region is never expected to contribute more than a 0.3 percentage point reduction to global poverty in a year. This stylized fact, which directly follows from the region's distribution of consumption vis-à-vis the international poverty line, can be captured in poverty elasticity measures. The semi-elasticities in Figure 9 show that poverty is less responsive to growth and distribution changes in sub-Saharan Africa than in India or China at any given poverty rate.

FIGURE 9: SEMI-ELASTICITIES OF POVERTY TO GROWTH AND DISTRIBUTION: CHINA, INDIA AND SUB-SAHARAN AFRICA (BASELINE)



Source: Authors' calculations


Diminishing returns

Underlying the slowdown in the rate of poverty reduction in our scenarios is the property of diminishing returns. Diminishing returns are observed at two levels.

First, the rate of poverty reduction within each country is expected to slow as it approaches zero. This follows logically from the shape of a country's probability density function. Once a country's mode exceeds the poverty line, the number of people positioned immediately below the poverty line typically starts to fall, reflecting the slope of the left-hand tail. This makes it harder to sustain the rate of poverty reduction. Progress can be further hindered if those individuals that remain in poverty are partially or fully excluded from the growth of the economy, due to factors such as geography or discrimination. Research indicates that this slowdown sets in, on average, once countries lower their poverty rate below 10 percent.³ We observe this slowdown in our scenarios for both China and India, which explains why they cede responsibility for leading global poverty reduction efforts before their poverty is eliminated.

Second, as global poverty approaches zero, it becomes increasingly concentrated in countries where the record of and prospects for poverty reduction are weakest. Today, a third of the world's poor live in fragile states but this share is set to rise to half in 2018 and nearly two-thirds in 2030 according to our baseline scenario (Figure 10). Most of these countries can be found in sub-Saharan Africa, which explains why some of its poor appear to be cut off from progress in our baseline scenario.

FIGURE 10: NUMBER OF PEOPLE IN POVERTY: FRAGILE STATES VERSUS STABLE COUNTRIES, OFFICIAL ESTIMATES AND BASELINE SCENARIO, 1990-2030



Source: Authors' calculations

Moreover, there are reasons to suspect that the prospects for fragile states implied by our baseline may paint a picture that is too sanguine. The consumption projections that provide a foundation for our scenarios cannot adequately account for the negative shocks countries may face in the future—a possibility that is, by definition, heightened in fragile environments. Furthermore, the countries where our data is most patchy are among the most fragile in the world. There are five countries in the developing world with populations of more than a million people each for which we have no country data and so we rely instead on regional aggregates: Eritrea, Myanmar, North Korea, Somalia and Zimbabwe. Each of these countries is likely to contain large numbers of poor people.⁴

One of the implications of China, and then India, initially leading the charge on global poverty reduction is that while global poverty aggregates are today dominated by these two nations, poverty will likely be spread more

thinly across countries in the future. This may give the impression that the determinants of progress on global poverty will shift from country-specific factors to broader global economic prospects. The opposite is likely true. As poverty numbers near zero, the fortunes of individual fragile states will become increasingly critical to attaining further progress. Success will depend on how many fragile states can be brought onto a more stable and prosperous trajectory. The prospects for this are exceedingly hard to judge.

CONCLUSION

The goal of ending extreme poverty within a generation is both noble and exhilarating. At first glance, all that is required to achieve this objective over the next 20 years is a repetition or continuation of whatever transpired in the previous 20. Yet the events of the last 20 years were anything but uniform and cannot be simply replicated. To the extent that they can be—namely, matching the same record on growth and distribution change across the developing world—even this will not sustain the trend rate of global poverty reduction. Today we face a different set of challenges in furthering progress against poverty, which demand a different set of solutions.

Promisingly, the immediate future offers tremendous potential for recording progress. With the global mode located at the poverty line, equitable growth can lift hundreds of millions more out of destitution. India is well placed to take up the reigns from China and drive global poverty reduction over the coming decade. But soon after, sustaining the rate of global progress will become increasingly hard. Further reductions in global poverty will no longer be accounted for by individuals who today are nearest the poverty line in countries enjoying growing prosperity, but by those who currently stand furthest from the line in environments where the prospects for sustained improvements in living standards are most tenuous.

There is no magic ingredient for eliminating poverty. Rather it hinges on a complex recipe: better than expected consumption growth and distributional trends in favor of the poor; country-by-country progress in transitioning fragile and conflict-affected states onto a stable path; strengthening the resilience of vulnerable households and economies to other kinds of shocks; the incorporation of isolated or excluded sub-national populations into the orbit of their economies; more deliberate and efficient targeting of the poor, including the poorest of the poor, at a country and sub-national level.

While the future trajectory of global poverty is impossible to predict, our understanding of what it will take to eliminate poverty is growing. The challenge for the global community is to seize this knowledge so that the dream of achieving a poverty-free world becomes a reality.

TECHNICAL APPENDIX

Framework

The following methodology was employed for 126 of the 129 developing countries for which the World Bank provides household survey data on the PovcalNet database. Three countries—Marshall Islands, Micronesia and West Bank and Gaza—were excluded from our analysis due to the absence of projections for consumption growth.

We treat income surveys the same as consumption surveys, other than where explicitly stated. All references to consumption hereafter therefore refer to income for countries who adopt income surveys. In countries that have administered both consumption and income surveys, our analysis is based on consumption surveys alone.

Following the World Bank, we scale up our results at a regional level to account for the population of missing countries. Hence, the poverty rate in a country with no survey data is assumed to be the same as the regional average based on the available country coverage in that region. This allows our scenarios to be seen as a continuation of official estimates.

The accuracy of our analysis relies, inter alia, on the quality of the underlying data. Chandy (2013) provides a review of the weaknesses in global poverty aggregates. We would especially highlight the sensitivity of our results to updated PPP estimates from the 2011 International Comparison Program, which are due for release later this year.

Baseline scenario

Our baseline scenario comprises poverty estimates for each country between 2010 and 2030, using the global poverty line of PPP\$1.25 a day. Country poverty estimates in a given year are derived from two pieces of information: the mean level of consumption per capita; and the distribution of consumption around the mean.

Initial consumption levels are drawn directly from the most recent survey estimates available for each country, published on the World Bank's poverty data repository, PovcalNet. Consumption levels for subsequent years up to 2030 are obtained by applying private consumption growth estimates and projections from the Economist Intelligence Unit (EIU) to initial survey levels. Where projections are incomplete or missing, we supplement with GDP growth projections from the IMF's World Economic Outlook (WEO) database until 2017 and assume a four-year moving average from thereafter.

Given the well-documented discrepancy between national accounts and survey-based measures of consumption growth, we apply a discount to EIU and WEO growth estimates before applying them to survey-based consumption levels. Our discounts are obtained by running a regression of average annual survey consumption growth (between each pair of consecutive surveys from the same country) on average annual national accounts private consumption growth over the corresponding period. Following Deaton (2005), our regression is population weighted. Following Ravallion (2003), separate regressions are run for countries that conduct consumption and

income surveys. Outlier observations that exceed ± 5 standard deviations are excluded from the regression. The coefficient on the independent variable, and therefore our discount, is 0.81 for consumption surveys and 0.91 for income surveys.

The discrepancy between national accounts and survey-based measures of consumption growth is particularly egregious for India. Given its importance to global poverty aggregates, and following the precedent set by the World Bank, a separate discount is calculated for India. This is obtained by taking the geometric mean of the ratio of annual survey consumption growth to annual national accounts private consumption growth for each period bookmarked by two surveys. This generates a discount of 0.54.

The use of a common discount across countries (other than India) does not imply that the relationship between national accounts and survey-based measures of consumption growth is uniform and stable. On the contrary, we are limited to this broad-brush approach by our inability to discern the relationship more clearly across countries and over time.

Distribution data is similarly obtained from the most recent survey for each country. Since changes in distribution over time follow no discernible trend and have no established correlates, we assume the distribution of consumption in each country remains unchanged in our baseline.

To obtain poverty estimates, we follow the method described in Datt (1998) and employed by the World Bank. In this method, parameter estimates are obtained for two functional forms of the Lorenz curve (general quadratic and beta), headcount ratios are then calculated using each curve, and a choice is made between the two results depending on which Lorenz curve is valid or which performs better on a goodness of fit measure.

Headcount ratios are converted into headcounts using the United Nations' medium variant population projections. Since China, India and Indonesia conduct separate surveys for rural and urban areas, we develop our poverty estimates for rural and urban areas separately for these countries, and use the UN's urbanization projections to derive future rural and urban populations.

For Afghanistan, Mongolia and Uzbekistan, no PPP estimates are available. For these three countries, we derive initial values of per capita consumption in PPP dollars using survey distribution data and the reported poverty rate against the national poverty line, by assuming the national and international poverty line are the same.

Alternative scenarios

Our alternative scenarios for consumption growth assume that the projections used in our baseline for each country are consistently wrong in one or other direction by 2 percentage points each year. The 2 percentage point figure is based on the differences we observe historically between forecasts and reported growth data. Specifically, we examined the errors between estimated GDP growth and reported GDP growth over more than a decade of IMF projections published in the WEO. The average absolute population-weighted error observed in one-year forecasts is 2.1 percentage points.

The absence of data prevented us from analyzing historical EIU projections; consumption projections as opposed to GDP projections; and the accuracy of longer term projections. Regarding the latter, we recognize that short-term errors are an imperfect proxy for capturing longer-term uncertainty.

Our alternative scenarios for distribution build on the work of Palma (2006). He submits that differences in the distribution of consumption, both across countries and within the same country over time, follow a stylized pattern and can largely be explained by the competition for shares of total consumption in each economy between the poorest 40 percent and the richest 10 percent of the population.

For our scenarios, we assume that the shares of total consumption held by the poorest 40 percent and the richest 10 percent in each country steadily shift in favor of one or other group between 2010 and 2030. To determine the size of these changes, we examine the scale of long-term distribution trends observed in historical data. We begin by taking the first and last survey available for each country and calculate the annual percentage point change in the consumption shares held by the poorest four deciles and the top decile respectively over the period. We turn these into absolute values, drop outliers, and then calculate the standard deviation separately for the poorest four deciles and the top decile. The standard deviations are 0.21 and 0.33 respectively. On this basis we alter the consumption shares of the two groups in increments of a quarter of a percentage point per year. To ensure our scenarios remain plausible, we apply a floor of 7 percent below which the collective consumption share of the bottom four deciles in any country cannot fall, and a ceiling of 60 percent above which the consumption share of the top decile cannot rise. As a point of comparison, this brings about an increase or decrease in the gini coefficient for each country over the 20 year span of our scenarios of, on average, 7 percentage points.

Probability density functions

To generate probability density functions, we follow the method outlined in Sala-i-Martin and Pinkovskiy (2009) and Quah (2002).

We begin by assuming that consumption in each country is lognormally distributed. This means that the distribution of consumption can be fully explained by two parameters: a scale parameter and a location parameter. The scale parameter is derived from a country's gini coefficient, which we obtain from the Lorenz curve, following Datt (1998). The location parameter is a function of the scale parameter and the consumption mean. The probability density function can then be directly estimated.

Since the methodology for creating probability density functions assumes distribution lognormality for each country, but the methodology for generating our scenario estimates doesn't, the two do not precisely correspond.

Probability density functions for countries with separate rural and urban surveys, regions, and the world are obtained by aggregating individual probability density functions weighted by population.

Semi-elasticities

Semi-elasticities capture the responsiveness of absolute (or percentage point) changes in poverty to growth and distributional change. To obtain these, we follow the methodology outlined in Klasen and Misselhorn (2007). As with probability density functions, we require the assumption that consumption is lognormally distributed in each country. Semi-elasticities can then be straightforwardly derived from the consumption mean, the poverty line, and the standard deviation of the distribution which is calculated from the gini coefficient.

The same caveat applies as above for the assumption of distribution lognormality.

Semi-elasticities for countries with separate rural and urban surveys and regions are obtained by averaging semi-elasticities weighted by population.

Fragility classification

Our classification of fragile states follows the unofficial methodology employed by the OECD (2013), which combines the list of “alert” countries on the Fund for Peace’s Failed States Index and the harmonized World Bank and Asian Development Bank’s list of fragile and post-conflict countries. Classification is assumed to remain unchanged across years.

ENDNOTES

1. TED Talk, February 2013
2. Reality is slightly more complex than this. The patterns we describe capture net changes in consumption patterns, masking more varied changes in the fortunes of specific individuals. Evidence suggests that for every four people who move out of poverty, three fall back, on average, whereas in our analysis this would be illustrated simply by a reduction in poverty of one person. This “churn” is an important characteristic of poverty dynamics but does not directly affect our analysis.
3. Ravallion, 2012
4. In our methodology, the poverty rate in countries with no survey data is assumed to be the same as the regional average based on the available country coverage in that region. We suspect that this underestimates the prevalence of poverty in at least four of these countries.

REFERENCES

- Chandy, Laurence, "Counting the Poor" Brookings Institution and Development Initiatives, 2013.
- Datt, Gaurav, "Computational Tools for Poverty Measurement and Analysis" International Food Policy Research Institute, 1998.
- Deaton, Angus, "Measuring poverty in a growing world (or measuring growth in a poor world)" *The Review of Economics and Statistics*, Vol. 87 No. 1, February 2005.
- Karver, Jonathan, Charles Kenny, Andy Sumner, "MDGs 2.0: What Goals, Targets, and Timeframe?" *Center for Global Development Working Paper 297*, June 2012.
- Klasen, Stephan and Mark Misselhorn, "Determinants of the Growth Semi-Elasticity of Poverty Reduction" *EUDN Working Paper*, 2008.
- OECD DAC, "Fragile States 2013: Resource Flows and Trends in a Shifting World" OECD DAC, 2013.
- Palma, Jose Gabriel, "Globalizing Inequality: 'Centrifugal' and 'Centripetal' Forces at Work" *DESA Working Paper No. 35*, 2006.
- Quah, Danny, "One Third of the World's Growth and Inequality" *UN-WIDER Discussion Paper*, 2002.
- Ravallion, Martin, "Measuring Aggregate Welfare in Developing Countries: How Well do National Accounts and Surveys Agree?" *The Review of Economics and Statistics*, August 2003.
- Ravallion, Martin, "Benchmarking Global Poverty" *World Bank Policy Research Working Paper WPS 6205*, September 2012.
- Ravallion, Martin, "How long will it take to lift one billion people out of poverty?" *World Bank Policy Research Working Paper WPS 6325*, January 2013.
- Sala-i-Martin, Xavier and Maxim Pinkovskiy, "Parametric Estimation of the World Income Distribution" *NBER Working Paper No. 15433*, October 2009.
- World Bank, "Global Monitoring Report 2012: Food Prices, Nutrition, and the Millennium Development Goals" World Bank, 2012.