



Person-Equivalent Poverty: An Introduction

Abstract In its 2015-16 *Global Monitoring Report*, the World Bank introduced the “person-equivalent headcount ratio” of Castleman, Foster, and Smith (2015), an intuitively appealing indicator that incorporates both the breadth and the depth of poverty. This brief explains the new measure and some of its properties – including its simple relationship to the total monetary gap between the poverty line and the consumption of the poor – and shows how the global distribution of poverty shifts when examined through a person-equivalent lens. Not surprisingly, the prevalence of person-equivalent poverty is markedly higher than that of conventionally measured poverty in Sub-Saharan Africa, reflecting the relatively severe depth of poverty in that region. It is considerably lower than conventional poverty in South and East Asia. We argue that for purposes of monitoring the severity of the global poverty challenge, person-equivalent poverty provides a better summary measure than either of its components.

Why a new poverty indicator?

Efforts to summarize progress in reducing global poverty have long faced a dilemma. The poverty headcount ratio—the share of people living below the poverty line—clearly portrays the breadth of monetary poverty, but says nothing about its depth—how far below the poverty line the poor are living. In particular, the headcount ratio ignores improvements in income or consumption that bring poor people closer to the poverty line, unless they actually cross the line. Meanwhile, existing measures of the depth of poverty—notably the poverty gap index—are widely seen as unintuitive. As a result, public discussions of poverty trends rarely get beyond the headcount ratio.

This *Brief* discusses a new measure—the person-equivalent headcount ratio—that may finally solve this problem, by providing an intuitive sense of the depth of poverty. This concept was introduced by Castleman, Foster, and Smith (2015), and features prominently in the World Bank’s 2015-16 *Global Monitoring Report*. The advantage of the person-equivalent measure is that it incorporates both the breadth and the depth of poverty. For purposes of monitoring the severity of global poverty, the person-equivalent measure is arguably superior to either the conventional headcount ratio or the

poverty gap index. In what follows, we explain the person-equivalent measure, illustrate some of its properties, and use it to help summarize the current global distribution of poverty and where that distribution is going.

Much of the discussion will focus on Table 1, which shows poverty measures for developing countries in each region in 1990 and 2012, measured against the revised international poverty line of \$1.90 at 2011 international prices. Columns (1)-(3) show, for each region, the conventional poverty headcount ratio H ; the poverty headcount q —the total number of people living below the poverty line; and the poverty gap index FGT1. The remaining columns will be introduced below.

Person-equivalent poverty monitoring

The logic of person-equivalent poverty is similar to that of full-time equivalents in the labor market. Like employment, consumption poverty has a dichotomous interpretation (employed/not employed; poor/not poor), but is also subject to varying degrees of intensity. A natural measure of the intensity of any individual’s poverty is his or her *monetary shortfall*, defined as the difference between

Table 1: Poverty indicators by region, 1990 and 2012

Column	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Indicator	Poverty Headcount Ratio	Poverty Headcount ¹	Poverty Gap Index	Average Monetary Shortfall ²	Total Monetary Shortfall ²	Person- Equivalent Headcount Ratio	Person- Equivalent Headcount
Unit of measurement	%	millions	%	\$/day at 2011 PPP	\$million/day at 2011 PPP	%	millions
Symbol used in text	H	q	FGT1	A	TMS	H ^e	q ^e
1990							
East Asia and Pacific	60.56	995.5	21.52	\$0.68	\$672.15	61.61	1,012.9
Europe and Central Asia	1.89	8.8	0.57	\$0.57	\$5.03	1.63	7.6
Latin America and the Caribbean	15.47	68.0	6.22	\$0.76	\$51.98	17.81	78.3
Middle East and North Africa	5.98	13.5	1.15	\$0.37	\$4.94	3.29	7.4
South Asia	50.65	574.9	14.77	\$0.55	\$318.51	42.29	480.0
Sub-Saharan Africa	56.75	287.6	24.94	\$0.83	\$240.18	71.41	361.9
Developing world ³	44.12	1,948.4	15.41	\$0.6636	\$1,293.01	44.12	1,948.4
2012⁵							
East Asia and Pacific	7.21	147.2	1.47	\$0.39	\$57.28	4.22	86.3
Europe and Central Asia	2.11	10.1	0.58	\$0.53	\$5.34	1.68	8.0
Latin America and the Caribbean	5.58	33.7	2.64	\$0.90	\$30.23	7.55	45.6
South Asia	18.75	309.2	3.73	\$0.38	\$118.15	10.79	178.0
Sub-Saharan Africa	42.65	388.8	16.47	\$0.74	\$285.87	47.26	430.8
Developing world	14.88	896.7	4.35	\$0.55	\$496.80	12.42	748.7

¹Number of poor people

²Relative to \$1.90/day at 2011 PPP

³For clarity we apply the same notation used by Castleman *et al.*

⁴The benchmark monetary shortfall A_0 and the values used to calculate it are shown in bold.

⁵PovcalNet suppresses estimates for MENA in 2012 due to low survey coverage; implied $H=2.3\%$, $q=7.7$ million.

Source: Columns (1)-(3) from PovcalNet, the World Bank's online poverty tool; columns (4)-(7) calculated by authors. The set of countries included in each region and in the "developing world" total is the same in each year.

the poverty line and that person's consumption or income. We label the monetary shortfall for individual i as $A(i)$, and its average value among all poor people at a given point in time as A .¹ Person-equivalent poverty measures are scaled against a benchmark value of the monetary shortfall, just as full-time equivalent measures of employment are scaled against a 40-hour workweek. We denote this benchmark monetary shortfall as A_0 .

For a given choice of the benchmark monetary shortfall, the number of person-equivalents $q^e(i)$ represented by any poor individual i is just the ratio

of that person's monetary shortfall $A(i)$ to the benchmark shortfall A_0 :

$$q^e(i) = \frac{A(i)}{A_0}.$$

For example, assume for a moment that the benchmark monetary shortfall is set at 35 cents, relative to a poverty line of \$1.90 per day. In that case, a person whose average daily consumption is \$1.55 would have a monetary shortfall of 35 cents, and would therefore be counted as one person-equivalent. Similarly, a person with daily consumption of \$1.20 (monetary shortfall 70 cents) would be counted as two person-equivalents; while a person with daily consumption of \$1.63 (monetary shortfall 7 cents) would be counted as 0.2 person-equivalents. The monetary shortfall for anyone above the poverty line is 0 by definition, so the calculation refers only to the poor.

¹ Some authors refer to this shortfall as the "depth of poverty." However, this label can create confusion, because the "depth of poverty" is also used as a synonym for the poverty gap index. To avoid this confusion, we refer to A as the average monetary shortfall.

The *person-equivalent headcount* in any given population is then simply the sum of person-equivalents across all poor people:

$$q^e = \sum_i q^e(i) = \frac{\sum_i A(i)}{A_0}.$$

The corresponding *person-equivalent headcount ratio* is the person-equivalent headcount divided by the total population.

A moment's thought will confirm that these concepts are closely related to the *total monetary shortfall (TMS)* in the population being monitored, because the *TMS* is simply the sum of monetary shortfalls across all individuals in the population. Denoting the person-equivalent and conventional headcounts by q^e and q , and noting that the total monetary shortfall is equal to the conventional headcount multiplied by the average monetary shortfall ($TMS = q \cdot A$), we can relate the person-equivalent headcount to the total monetary shortfall:

$$q^e = \frac{TMS}{A_0} = q \cdot \frac{A}{A_0}.$$

Dividing these elements by population yields a similar expression relating headcount ratios.

$$H^e = H \cdot \frac{A}{A_0}.$$

The person-equivalent headcount ratio therefore combines the information in two existing poverty measures – the headcount ratio H and the average monetary shortfall A . In doing so it yields a much more revealing picture of poverty than do its individual components. The benchmark monetary shortfall, A_0 , operates as a scaling factor – it is held fixed when comparing poverty outcomes across countries, regions, or time, and therefore has no impact on proportional time trends of poverty, or on rankings across countries.

The benchmark monetary shortfall could, in principle, be set at any convenient value. However, once set, it must remain set. For person-equivalent measures to deliver on their promise of providing an intuitive sense of the depth of poverty among a given population, the benchmark monetary shortfall that underlies these measures must remain fixed

over time—or at least as long as the poverty line itself remains fixed.

In keeping with this logic, we base our own calculations on a benchmark monetary shortfall A_0 of 66.36 cents per day at 2011 international prices, the same value used by the World Bank in the 2015-16 *Global Monitoring Report*, the document that introduced person-equivalent measures to the non-specialist public.² This value of A_0 was adopted because it represents the average monetary shortfall in the world in 1990—the first year of the Millennium Development Goal period—relative to the World Bank's new extreme poverty line of \$1.90 per day at 2011 international prices.

As explained by Castleman *et al.* (2015), person-equivalent poverty is a rescaled version of the conventional *poverty gap index*, known as *FGT1* in the Foster-Greer-Thorbecke class of poverty measures. The poverty gap index is defined as the ratio of the total monetary shortfall to what total consumption would be if all persons were exactly at the poverty line. Equivalently, it is the product of the conventional headcount ratio and the average *relative* monetary shortfall: $FGT1 = H \cdot (A/z)$, where z is the poverty line.³ Using our earlier equations, the relationship between person-equivalent poverty and *FGT1* is

$$H^e = FGT1 \cdot (z/A_0).$$

For purposes of ranking countries by the relative magnitude of their poverty challenge – or tracking the progress of this challenge over time, the person-equivalent headcount ratio and the poverty gap index give identical results, as long as the same

² Oddly, the *Global Monitoring Report* never actually mentions the value of A_0 that underlies its discussion and charts on the subject. The next footnote explains the derivation of this value.

³ Rearranging this expression as $A = FGT1 \cdot (z/H)$ and inserting PovcalNet's estimates of the poverty gap index and poverty headcount ratio for each developing region allows us to calculate the average monetary shortfall A , shown in column (4) in Table 1. Inserting the values for all developing countries in 1990 (15.41% and 44.12% respectively), along with the \$1.90/day poverty line, allows us to calculate the benchmark monetary shortfall A_0 , \$0.6636 per day.

poverty line and benchmark monetary shortfall are used for all countries and periods.⁴ But the rescaling is useful: as emphasized above, the person-equivalent measures convey important information not embodied in their conventional counterparts.

The global geography of person-equivalent poverty

The upper-right section of Table 1 shows the person-equivalent headcount and headcount ratio for each developing region in 1990, calculated using the benchmark monetary shortfall of \$0.6636 per day. Because 1990 serves as the baseline year in computing this value, for the developing world as a whole, $q^e = q$ in 1990—that is, the conventional and person-equivalent headcounts are identical by construction in the baseline year.

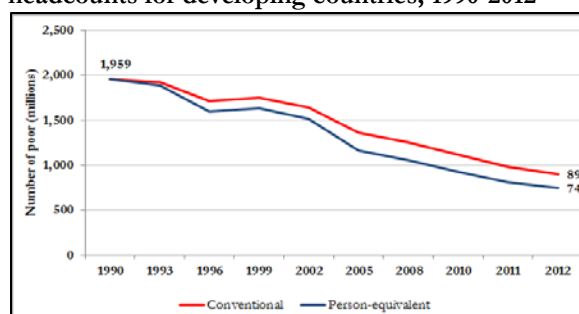
The lower panel of Table 1 illustrates how the person-equivalent lens informs our perceptions regarding the distribution of global poverty. The person-equivalent measures place the weight of global poverty much more squarely in Sub-Saharan Africa than do the conventional measures, reflecting the unusually high depth of poverty in that region, particularly in relation to the breadth of poverty. The same applies to Latin America and the Caribbean, though the absolute numbers of poor are much smaller than in Africa. By comparison, person-equivalent poverty is considerably lower than conventional poverty in South and East Asia, reflecting the relatively close proximity of poor households in those regions to the poverty line. These differences were evident in 1990, but have grown over time as person-equivalent poverty has declined more slowly in Sub-Saharan Africa than in other regions.

What can we say about the evolution of poverty to date? Figure 1 shows the conventional and person-equivalent headcounts since 1990, using a

⁴ The ratio of person-equivalent headcount ratios in any two countries and periods is equal to the ratio of poverty gap indexes (and the product of the ratios of conventional headcount ratios and average depths of poverty), because the constants drop out. Thus for example $H^e(i, t)/H^e(j, s) = FGT1(i, t)/FGT1(j, s)$ for any two countries i and j and years t and s . See Crosswell (2015) for further discussion.

benchmark monetary shortfall of \$0.6636/day. Since the person-equivalent headcount is proportional to the product of the conventional headcount and the current monetary shortfall⁵, its rate of decline over time is approximately equal to the sum of the rates of decline of its two components. Over the period covered in Figure 1, the average monetary shortfall in the developing world fell from the benchmark value of 66.36 cents/day to 55.41 cents/day. The person-equivalent headcount ratio therefore fell even faster than the already-impressive decline in the conventional poverty headcount ratio. In particular, although the conventional headcount ratio declined by 66 percent between 1990 and 2012, thus comfortably exceeding the rate of improvement needed to achieve Millennium Development Goal target 1a (cutting the prevalence of extreme poverty in half between 1990 and 2015), the person-equivalent headcount ratio fell 72 percent over the same period.

Figure 1: Conventional vs. person-equivalent poverty headcounts for developing countries, 1990-2012



Projecting person-equivalent poverty

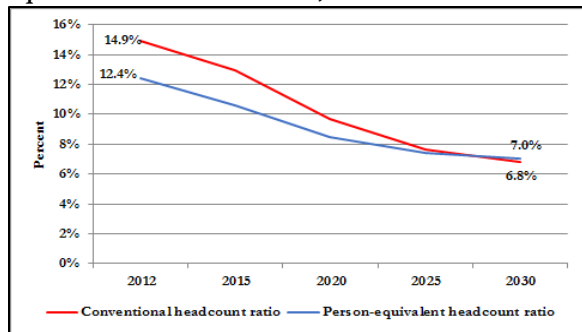
How will person-equivalent poverty evolve over the next 15 years? The most common approach to projecting future poverty headcounts uses a forecast of growth in GDP per capita to derive a projected growth rate for average consumption, along with an assumption about how the distribution of consumption will change over time. Taken together, these assumptions yield a projection of the country's entire distribution of consumption over time – allowing a straightforward calculation of the relevant poverty measures.

⁵ The constant of proportionality, as we saw earlier, is the benchmark average shortfall.

Following standard practice, our projections assume that future growth will be distribution-neutral, in the sense that all individuals in a given country have the same growth rate of consumption. This assumption embodies the key finding by Ravallion and Ferreira (2008) and others that past growth in developing countries has been roughly distribution-neutral on average, in that inequality has decreased with growth about as often as it has increased.⁶ The projections are derived by inserting growth forecasts from the IMF’s World Economic Outlook into USAID’s PovcalNet Projector. The Appendix explains the projection methodology in greater detail.

Figure 2 shows projections of the conventional and person-equivalent headcount ratios for the developing world from 2012 through 2030. The projection reflects the observation already registered by the World Bank (2015) and others, that at currently projected growth rates, distribution-neutral growth will not be sufficient to eliminate extreme poverty over the horizon of the 2030 Agenda. Our interest, however, is in the relationship between the conventional and person-equivalent headcounts. From this perspective, the most noteworthy aspect of this projection is that the conventional headcount ratio substantially exceeds the person-equivalent headcount ratio in 2012, but thereafter falls more rapidly, falling slightly below the person-equivalent headcount ratio by 2030.

Figure 2: Projection of conventional and person-equivalent headcount ratios, 2012-2030



⁶ In fact, inequality in the median developing country has modestly declined since 1990 (World Bank 2016). Holding growth projections unchanged, any decline in within-country inequality would reduce poverty more rapidly than our distribution-neutral projections indicate.

Figures 3 and 4 portray projected regional trends in person-equivalent poverty headcounts. Figure 3 suggests large regional differences in prospects for reducing person-equivalent poverty, with rapid growth in East and South Asia virtually eliminating extreme poverty in those regions by 2030. In contrast, the number of person-equivalent poor in Sub-Saharan Africa and in Latin America and the Caribbean is projected to increase gradually, in the former because population growth roughly offsets a reduction in the person-equivalent headcount ratio from 47 to 31 percent, in the latter because of projected slow growth or deterioration in the countries with the largest numbers of person-equivalent poor. The projected result, seen in Figure 4, is a dramatic concentration of the shares of those two regions in the total number of person-equivalent poor between 2012 and 2030. By 2030, Sub-Saharan Africa is projected to have 86.9 percent of the person-equivalent poor people in the developing world, while Latin America and the Caribbean have 9.4 percent.

Figure 3: Projection of person-equivalent poverty headcounts by region, 2012-2030

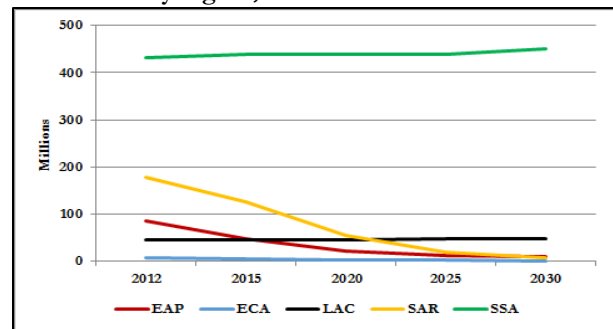
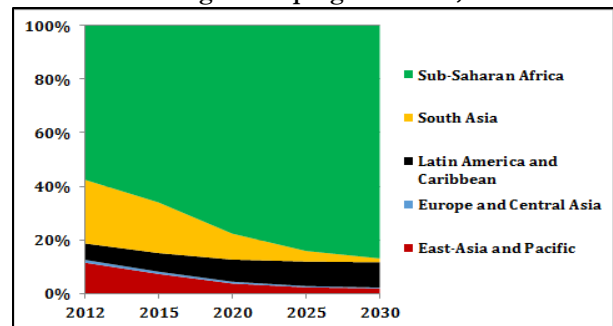


Figure 4: Regional shares of total person-equivalent headcount among developing countries, 2012-2030



These projected regional differences in the rate of poverty reduction account for the reversal between the conventional and person-equivalent headcount

ratios seen in Figure 2. The regions with the smallest average monetary shortfalls—East and South Asia—begin with 51 percent of all extreme poor in 2012, but thereafter achieve rapid reductions in the conventionally measured poverty headcount ratio, thereby pulling down the global average. In contrast, the two regions with the largest average monetary shortfalls—Sub-Saharan Africa and Latin America and the Caribbean—experience gradual increases in person-equivalent poverty, causing the global average person-equivalent headcount ratio to decline more slowly than the conventional ratio.

Quantifying a global basic income guarantee

The tight links between the person-equivalent poverty measures and the total monetary shortfall (TMS) relative to a particular poverty line suggests an additional use for these measures: calculating the cost of an ideal basic income guarantee (BIG) that would bring everyone in the world up to the \$1.90/day poverty line. We briefly illustrate this point as a way of underscoring the properties of person-equivalent poverty.

To qualify as ideal, a BIG would need to meet three criteria: (1) it would be perfectly targeted, generating a customized transfer to each poor household and paying no benefits to the non-poor; (2) it would involve no administrative costs, and no leakages of funds in the form of corruption; and (3) it would not induce any reduction in work effort or other productive behavior on the part of recipients. With these conditions in place, the cost of a BIG would simply be the product of the person-equivalent headcount and the benchmark shortfall.

The cost of an ideal BIG can be calculated using the numbers in the last row of Table 1, showing PovcalNet's estimates for 2012. An average daily shortfall of \$0.5541 implies an annual shortfall of \$202.24 per capita. Multiplied by the estimated 897 million persons living in extreme poverty in 2012, these numbers imply that an ideal BIG in 2012 would have cost just under half a billion dollars per day at 2011 international prices, or \$194 billion a year in 2016 international dollars. The dollar cost would be considerably lower if such a program were funded through budgetary support from donor countries, because consumption is relatively inexpensive in low-income countries. Using market

exchange rates rather than Purchasing Power Parity (PPP) exchange rates, the cost of an ideal BIG would be \$87 billion in 2012, considerably less than the \$138 billion in net Official Development Assistance (ODA) provided by traditional donor countries in 2012.⁷

However, the assumptions underlying this calculation represent best-case scenarios. The administrative costs of targeting are falling with the advance of mobile money platforms, social safety-net infrastructures, and biometric identity programs in low-income countries. But the adoption of such advances in administrative technology differs widely across countries, as does the quality of the bureaucracies charged with administering public services. And the program would of course generate strong incentives for poor households to under-report their income, because any income earned up to \$1.90 a day would be fully offset by an equivalent reduction in the household's grant. In essence, a "perfectly targeted" BIG would impose a 100-percent tax rate on earned income up to the poverty line. Allowing for imperfect targeting and various kinds of leakage could easily double the \$87 billion price tag cited above. This includes reduced international remittances, which would face the same 100% tax when directed to households with incomes below \$1.90.

Any adverse effects on labor supply (i.e., on actual earnings as opposed to reported earnings) would also increase the cost of a BIG. Economic theory suggests that such costs might be large, because any household lacking a secure prospect of earning more than \$1.90 a day would achieve a better combination

⁷ This is total net ODA from members of the Development Assistance Committee (DAC) of the OECD. See [DAC Financing for Sustainable Development](#) for statistics. The TMS would in some cases be large relative to the receiving economy, particularly in countries with large person-equivalent headcount ratios. If such transfers were additional to existing aid, they would be likely to produce a substantial real exchange rate appreciation, which would erode their purchasing power in the receiving economy and increase the baseline value of the person-equivalent headcount. A larger transfer would therefore be required. The so-called 'Dutch disease' impact of aid varies across countries, but tends to be larger in countries already receiving substantial inflows (Fielding 2013).

of consumption and non-market uses of time – at least in the short run – by stopping work and accepting a transfer of \$1.90, than by (in effect) contributing up to \$1.90 of uncompensated labor to the BIG program. In this worst-case scenario, the cost of the program would more than triple, to over \$620 billion per year in 2011 international dollars.⁸

In reality, of course, the labor-market impact of a BIG would presumably be context-specific and sensitive to a wide range of variables including not only the overall program design but also the social meaning of work and the demand-side impact of the transfer on the local economy. Recent research in behavioral economics stresses a set of potentially favorable impacts operating through the household's capacity for forward planning and risk-taking (Mullainathan and Shafir 2013). Very little is known, however, about the likely size of these effects and how they depend on the design of the program.

In principle, the surest way to eliminate the labor-market distortions discussed above (and limit leakages directly related to targeting) would be to provide a universal transfer rather than a basic income guarantee. In contrast to a BIG, a universal transfer would be paid to everyone regardless of their income level. Such a program would impose no offset against non-transfer income, and would therefore reduce labor effort only to the extent that recipients chose to use part of their increased income to enjoy more leisure time. On the other hand, such a program would be extremely expensive: guaranteeing a transfer of \$1.90 per day to all developing-country citizens would cost around \$1.9 trillion per year at market exchange rates. We hasten to emphasize that no serious proposal to support a universal transfer program is on the global political agenda. By contrast, national safety-net programs continue to be developed apace, employing a variety of designs to limit behavioral distortions and reconcile targeting objectives with administrative constraints.

Assessment

The new measures discussed in this *Brief* provide a useful means to monitor progress against poverty, combining the information on the breadth of poverty contained in the conventional headcount ratio with a more intuitive sense of the depth of poverty than is typically conveyed by the poverty gap index. In most cases, the person-equivalent headcount ratio is likely to be shown alongside the conventional headcount ratio, with changes in the gap between the two measures providing a sense of changes in the depth of poverty. Nevertheless, the person-equivalent headcount and headcount ratio can be used as stand-alone poverty measures, which arguably provide more insight into the poverty situation on the ground than their conventional counterparts. Meanwhile, the close link between the person-equivalent headcount measures and the average and total monetary shortfalls facilitates their use in examining important questions of development policy, such as the scale of resources needed to close the poverty gap through public transfers. In order to realize these benefits, it is necessary that potential producers and consumers of poverty data develop an intuitive sense of how the person-equivalent measures relate to their conventional counterparts—a goal to which this *Brief* seeks to contribute.

⁸ That is, \$1.90/day times 365 days per year times 896.7 million people living on less than \$1.90/day in 2012.

Appendix: data and methods

All of the data used in the calculations in this Brief come from PovcalNet, the World Bank's online poverty tool. The projections rely on USAID's PovcalNet Projector, developed in collaboration with the Department for International Development (DfID) of the United Kingdom, using data downloaded from PovcalNet. The Projector starts with PovcalNet's estimate of mean survey income or consumption in 2012, and then applies projected rates of growth in per capita real GDP for each country in PovcalNet, downloaded from the International Monetary Fund's *World Economic Outlook database, April 2016 Edition*. The Projector assumes that between 2012 and 2015, each country's real GDP per capita grew at the rate estimated or forecast by the IMF over the same period. For 2015-2030, the Projector applies the IMF's projected growth rate between 2015 and 2020.

In both cases, growth in real GDP per capita is converted into growth in mean consumption or income at the household level using elasticities suggested by Chandy et al. (2013): Thus, a 1 percent increase in real GDP per capita is assumed to translate into a 0.81 percent increase in household consumption per capita or into a 0.91 percent

increase in household income per capita, depending on the type of survey used in each country. In the special case of India, the elasticity is set to 0.54, based on evidence that the transmission of GDP growth into growth in measured household-level consumption has been especially limited in India.

Finally, the projected values of household income or consumption for each country are used to estimate the poverty headcount ratio and other poverty statistics, using the same parameters or distribution used by PovcalNet itself. In each case, the distribution is based on PovcalNet's 2012 Regional Aggregation, which is in turn based on the most recent household survey conducted in each country. Projected country population values are based on World Bank projections, downloaded via the World DataBank.

Acknowledgements

The authors gratefully acknowledge helpful comments from Michael J. Crosswell, Caroline Smith, and Greg Gangelhoff. All errors are the sole responsibility of the authors. Comments are welcome at soconne1@swarthmore.edu and dsillers@usaid.gov.

References

- Castleman, T., J. Foster, and S. C. Smith (2015) "[Person-equivalent Headcount Measures of Poverty](#)" *Institute for International Economic Policy Working Paper Series* 2015-10, Elliot School of International Affairs, George Washington University, Washington, DC
- Chandy, Laurence, Natasha Ledlie, and Veronika Penciakova (2013) "[The Final Countdown: Prospects for Ending Extreme Poverty by 2030](#)." Policy Paper 2013-04, Brookings Institution, Washington, DC
- Crosswell, Michael J. (2015) "[Country Selectivity and Extreme Poverty](#)," *Economics Brief* No. 2, U.S. Agency for International Development, December 11
- Fielding, David and Fred Gibson (2013) "[Aid and Dutch Disease in Sub-Saharan Africa](#)" *Journal of African Economies* 22(1): 1-21
- Mullainathan, Sendhil and Eldar Shafir (2013) *Scarcity: Why Having So Little Means So Much* (New York: Times Books)
- Ravallion, Martin and Francisco H. G. Ferreira (2008). "[Global Poverty and Inequality: a Review of the Evidence](#)," Policy Research Working Paper 4623, World Bank, Washington, DC
- World Bank (2015) [Global Monitoring Report 2015/2016: Development Goals in an Era of Demographic Change](#) (Washington, DC: The World Bank)