

Reaching Poor People

Sobering evidence from Africa illustrates how hard it is to target antipoverty efforts well

Caitlin Brown, Martin Ravallion, and Dominique van de Walle

It has often been said that the world's aggregate poverty gap—the total monetary amount by which all poor people fall below the poverty line—is modest when one uses poverty lines typical of low-income countries. For example, Annie Lowrey wrote in the *New York Times* magazine on Feb. 23, 2017, that “one estimate ... recently calculated that the global poverty gap ... is roughly what Americans spend on lottery tickets every year, and it is about half of what the world spends on foreign aid.”

The implication is sometimes drawn that only a modest sum of money is needed to eliminate global poverty—to bring all poor people up to the international poverty line that separates the poor from the nonpoor.

However, eliminating poverty is a lot harder than the size of the aggregate poverty gap might suggest. Identifying who is poor and by how much is particularly challenging. The calculation cited in the *New York Times* could therefore be way off the mark. Some who are truly poor go wanting while funds end up in the hands of others. Because of imperfect information about levels of living, the amount of money needed to eliminate poverty can quickly balloon.

We have tried to assess whether the data typically available and routinely used by policymakers in sub-Saharan Africa—the poorest region of the

world by most measures—are adequate to reliably determine who is poor.

Finding poor households

Identifying poor households is often complicated by a lack of reliable data. It is difficult, or even impossible, in many cases to assess the living standards of all individuals in the population. In higher-income countries, income tax records help. But tax records are not a feasible option in many developing economies, given that many households work in the informal sector or traditional agriculture. Governments are often hindered by constraints on record keeping for reliably measuring all incomes, and these constraints can be severe in poor countries. In addition, household-level data may not be a good indicator of the living standards of individuals within the household.

To try to overcome this obstacle, governments across the world have increasingly turned to some form of proxy means test to identify poor households. The idea is simple. A score is given to each household based on a (usually small) set of readily observable household characteristics that are suggestive of whether a household is poor. Such characteristics may include the size of the household; the gender of the head of the household; the demographic



composition of the household; the type of dwelling in which the family lives; what the dwelling is made of; and the assets a household has (for example, whether the household owns basic items such as a radio or phone). Each of the characteristics is given a weight based on its observed statistical relationship with household consumption based on nationally representative sample surveys.

Among researchers and practitioners there has been much debate about the efficacy of proxy means tests (that is, how well the characteristics substitute for direct evidence of income or consumption). Supporters claim that it is a reliable method; critics say that the approach gives unsatisfactory predictions about who is poor and who is not. Concerns have also been raised about lack of transparency and divisiveness within communities, whereby similar households are treated very differently based on some opaque score on a proxy means test.

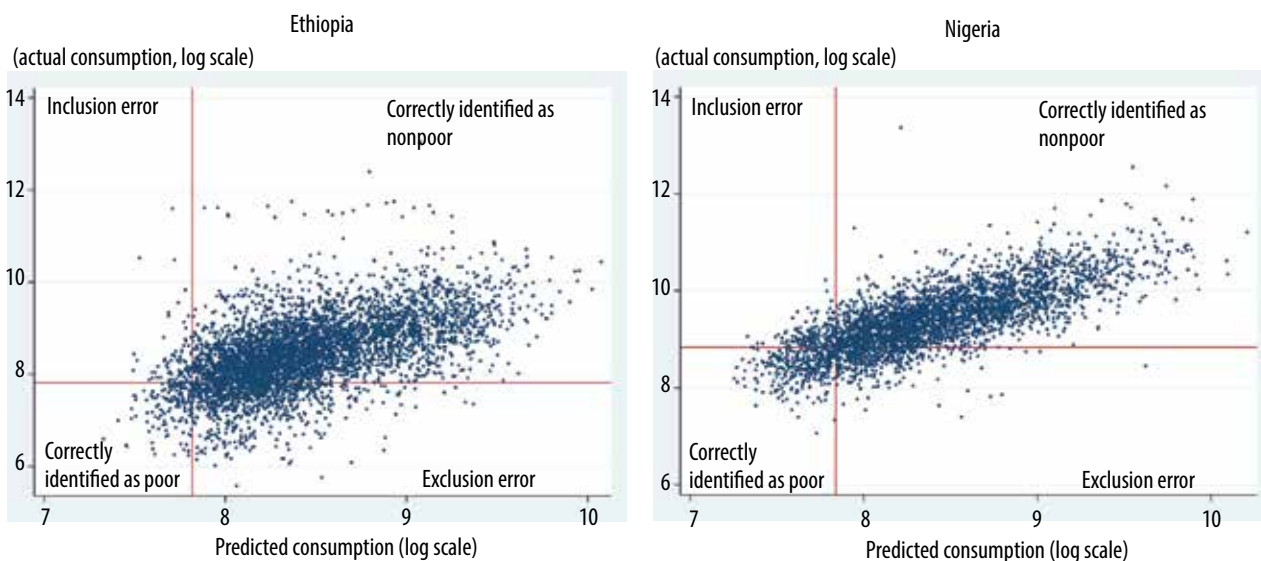
We study the performance of this popular method in a number of African countries. Our results point to both the strengths and weaknesses of the method. The

good news is that proxy means tests can substantially reduce the inclusion of nonpoor households in an antipoverty program; in most cases we studied, the inclusion error rate can be at least halved. The bad news is that this comes at the cost of substantial exclusion of the poor. And when the objective is to reduce poverty, policymakers should be worried about exclusions.

A key reason for the high exclusion error rates is that the standard proxy means test works less well near the extremes of the distribution of household consumption. The statistical properties of the method often lead it to overestimate living standards for the poorest (and underestimate them for the richest). When we compare actual household consumption to the predicted values from the proxy means test, it becomes clear just how much this overestimation matters. For the poorest 20 percent of households, in terms of actual consumption, proxy means tests yield predicted values that are between 50 percent and 100 percent higher than actual consumption. This means that the test misses many of the poorest households in almost all

Poor predictors

In Nigeria and Ethiopia, proxy means tests do a far better job of excluding from antipoverty programs households that are not poor than in identifying those that are poor. Similar errors occur in most low-income countries.



Source: Brown, Caitlin, Martin Ravallion, and Dominique van de Walle. 2016. "A Poor Means Test? Econometric Targeting in Africa." NBER Working Paper 22919.

Note: The red lines represent the poverty line that divides the households with the lowest 20 percent of consumption (poor) from those who are considered nonpoor. Inclusion error (upper left quadrant) occurs when a nonpoor household is identified as poor and exclusion error (lower right quadrant) when a poor household is incorrectly identified as nonpoor. Proxy means tests use household characteristics in lieu of income or consumption documentation to determine whether a household is considered poor. Consumption data are converted to a log base to improve the representation.

countries: on average, 80 percent of poor households are counted as nonpoor by the test, and 40 percent of nonpoor households are counted as poor.

For two of the countries in our study, Ethiopia and Nigeria, the chart shows the relationship between actual consumption and the scores using a standard proxy means test. In both countries, there is a strong positive relationship between the proxy means test scores and actual consumption; most of those deemed not to be poor based on the score are correctly classified. But there are substantial exclusion errors, strikingly so for Ethiopia, where 95 percent of the poor are identified as nonpoor (compared with 55 percent for Nigeria). But for both countries, and indeed for all those in our study, the commonly used proxy variables are clearly not doing a very good job of distinguishing poor households.

For a fixed budget, we find that a common form of the proxy means test reduces poverty only slightly more on average than a universal basic income, in which everyone gets the same transfer, whether they are rich, poor, or middle income. One can do about as well as the proxy means test by making a uniform transfer based on just a few household-level characteristics, such as gender of the household head or whether the household has young children. Indeed, once the often-long delays in implementing proxy means tests and households' changing circumstances are considered, these simpler targeting methods perform

better on average in bringing down the poverty rate. When the costs of constructing and implementing the proxy means test are considered, these simpler targeting methods may be preferable in terms of poverty reduction for a given budget.

Pinpointing poor individuals

Even if poor households could be correctly targeted, it is still unclear whether that ensures that poor individuals will be reached. Poverty is individual deprivation, but is almost invariably measured using household data. Typically, every member of a poor household is assumed to be poor, and every member of a nonpoor household is assumed not to be poor.

But the widely used household-based measures may not do a good job of identifying disadvantaged individuals who may share relatively little in the household's aggregate consumption or face impediments in accessing opportunities outside the household—including health, education, and financial services. Missing data on individual-level poverty present a significant hurdle to examining whether antipoverty programs targeted at poor households reach poor people. Individual-level consumption is not easily collected, and it is difficult to determine how income earned by individuals is shared with other household members. For example, in a household where one member works, income could be shared equally among all members or one of them may take a disproportionate share. Specific

We can find nonpoor individuals residing in poor households and poor individuals in nonpoor households.

members, such as the elderly or orphans, may be discriminated against. Thus we can find nonpoor individuals residing in poor households and poor individuals in nonpoor households.

One dimension of individual welfare that is suggestive of poverty and can be observed in many surveys is nutritional status. We undertook a comprehensive study of the relationship between household wealth (measured by either an index of assets owned or household per capita consumption) and individual nutritional status for 30 countries in sub-Saharan Africa using the Demographic and Health Surveys.

We find a reasonably robust relationship between household wealth and undernutrition indicators for women and children—that is, the incidence of undernutrition tends to fall as household wealth rises. Nonetheless, about three-quarters of underweight women and undernourished children are not found in the poorest 20 percent of households. And about half are not found in the poorest 40 percent. Moreover, countries with a higher overall incidence of undernutrition tend to be those in which a larger share of the undernourished are found in nonpoor families.

There are several potential explanations for these results. The demographic imbalance between poor and nonpoor households, such as poor households having more children than nonpoor households, does not turn out to be a major factor. While measurement errors are clearly present, our tests do not suggest that this is the main reason for our findings.

Intrahousehold inequality helps explain why such a large proportion of undernourished women and children reside in nonpoor households. We find that a sizable proportion of undernourished women and children live in households where a male head of household is not underweight—although sometimes the male head is underweight and other family members are not.

However, intrahousehold inequality is only part of the explanation. This is evident when we redo our calculations assuming that there is no intrahousehold inequality (every household member is assigned the average household nutritional status). Even then we find that a sizable share of undernourished women

and children are not found in poor households as identified in the survey data. That appears to be because both poor and nonpoor households living in impoverished areas often share the same health environment and are thus exposed to similar health risks. We find evidence consistent with this explanation using data on the incidence of illness in children across the household wealth distribution.

No easy solution

Information is not of course the only factor affecting antipoverty policies; government budget constraints (also reflecting the government's capacity to raise revenue), incentive effects (such as when nonpoor people change their behavior to receive benefits intended for the poor), and political economy (when some nonpoor people do not support efforts to help poor people) must also be considered. But information is undeniably an important constraint. Policymakers need to have realistic expectations of what can be accomplished given the reliability of the data available.

Our results suggest that the standard data sources on poverty are not very effective in identifying poor households or poor individuals. To reach undernourished women and children, policy interventions will require either much more individualized information or broader coverage than policies finely targeted to poor households. This is especially true in countries with a high incidence of undernutrition.

There is some potential for using better data and better methods. But the idea that we can easily eliminate poverty by finely targeted transfers is overly optimistic. This is true even before we start to think about the (potentially serious) adverse incentive effects that such a policy could generate. **FD**

CAITLIN BROWN is a PhD candidate and **MARTIN RAVALLION** is Edmond D. Villani Professor of Economics, both at Georgetown University. **DOMINIQUE VAN DE WALLE** is a lead economist in the Development Research Group at the World Bank.

This article draws on "A Poor Means Test? Econometric Targeting in Africa," a 2016 working paper for the National Bureau of Economic Research, and "Are Poor Individuals Mainly Found in Poor Households? Evidence Using Nutrition Data for Africa," a 2017 World Bank Policy Research Working Paper, both by the authors.