development initiatives

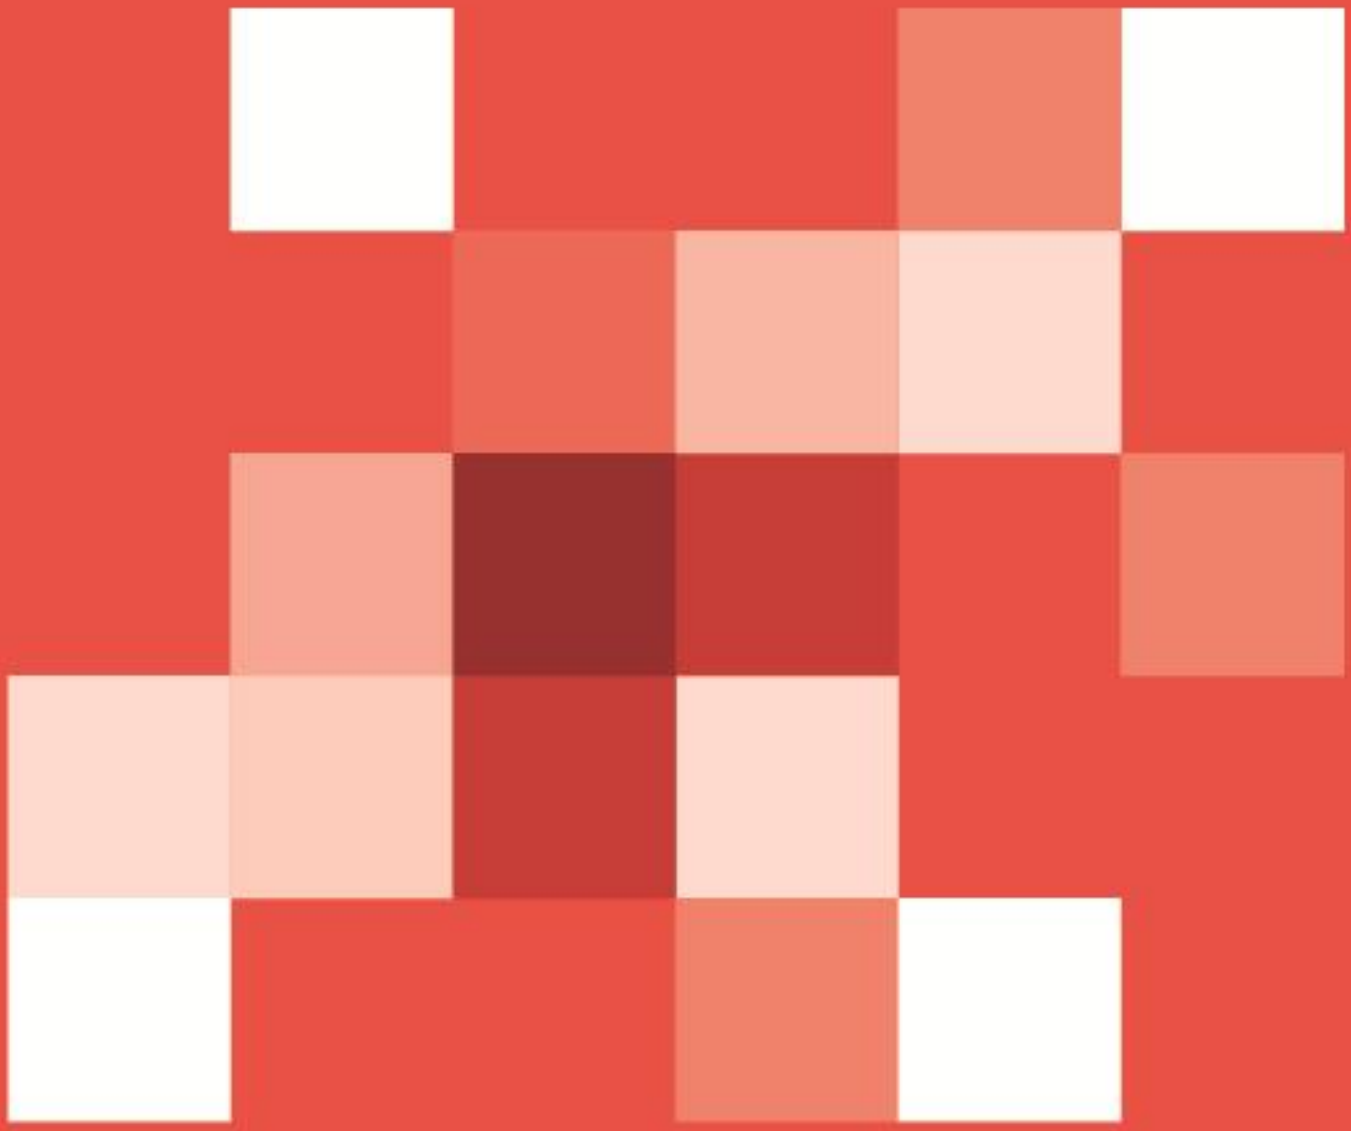

February 2018

coding the P20

how we developed and coded the P20 Initiative

background paper

Contents

Introduction

This briefing outlines the unification process we used in creating the statistics for the P20 Initiative.

In creating the statistics for the P20 Initiative, we used 99 different household surveys. Of these, 56 came from the Demographic and Health Surveys (DHS), 41 from Multiple Indicator Cluster Surveys (MICS), and the remaining 2 from the China Family Panel Studies (CFPS) and the Brazilian National Demographic and Health Survey (PNDS). Because DHS and MICS were both designed to assess the health of mothers and children, they ask generally the same set of questions across all countries. Despite common questions, the answers recorded by each survey differ greatly by country and across time. The surveys also did not systematically include questions about income or consumption, and so we needed to devise a method to compare economic well-being.

The Integrated Public Use Microdata Series (IPUMS)-DHS project[1] has attempted to unify the answers for a select set of indicators for women, children, and births but has only recoded 23 countries to date. In order to create an internationally comparable dataset, we needed to create a unified coding scheme for all of our 99 household surveys.

Languages, spellings, and cultural differences

The biggest coding differences in our surveys tended to occur when free-text fields were allowed in the survey questionnaires. For example, a question about wall materials resulted in at least 13 different answers that involve bamboo in some regard (e.g. "bamboo (guadua)", "bamboo/sticks/mud", "bambu with mud" etc.), and there were 29 different answers to questions where we would expect only a yes/no answer (e.g. "oui", "si", "sim" etc.). At the moment, these differences have been manually unified with a recode guide that re-defines responses for 8 categories of questions (school attendance, floor material, school levels, toilets, urban/rural, wall material, water, and yes/no). To avoid having to create unique codes for floors, walls, toilets, and water, we decided to reduce each into a binary variable that indicates whether the asset is "improved" according to UN definitions. For example, a hanging toilet is not improved, while a pit latrine with slab is improved. For water sources, we had to split the binary variable along the urban/rural divide, since the definition of improved depended on where the household was located. For urban households, only water piped into the household and bottled water were considered improved, while rural household requirements were much more lenient.

Even when everything is spelled correctly and in the same language, there are still some coding difficulties due to cultural differences between countries. While most DHS reduce schooling into broad categories like preschool, primary, secondary, and higher, some MICS questions code a vast variety of different types of schools (for example religious school, technical school, vocational school). To avoid making a value judgement, particularly on religious schools, we relied on the number of years of schooling. In other words, if a student attended a religious school for more than five years, they were considered to have completed primary education. We also needed to code whether a birth attendant is considered "skilled". In most cases, it's easy to decide that all professional medical staff are "skilled", but there are edge-cases like midwives and community health staff that can be classified in either direction. For the purposes of this project, we ended up marking all health volunteers, traditional midwives and non-medical people as "un-skilled". Toward the end of the P20 Initiative, we were still in the process of manually combing through national reports to derive country-specific definitions for skilled or un-skilled birth attendants.

Comparing wealth

With measures of education, stunting, assets, and home materials homogenised to the best of our ability, we needed a comparable measure of economic well-being to determine the relationship between poverty and these other statistics. The development of this particular feature of the methodology took months of research. Weeks were spent working on methods that we eventually abandoned altogether in favour of more parsimonious solutions. To better understand how we arrived at our present methodology, it is important to appreciate the lessons we learned from the failed methodologies.

The nature of the problem

Our objective was to ultimately determine the demographic details of those people who are in the poorest 20% of the global population. The best tools we had at our disposal to accomplish this were the global income distribution provided by the World Bank's PovCalNet[2] and the microdata from household surveys (mostly DHS and MICS). The issue arises when we need to apply the information from PovCalNet to the household survey microdata. Technically, PovCalNet is also based on household surveys (Living Standards Measurement Study (LSMS)); if we had been able to obtain this microdata, we could have circumvented this entire issue by simply measuring the demographics in the LSMS microdata. Unfortunately, a large majority of these LSMS surveys are behind pay-walls and data-use agreements that must be negotiated via email on a country-by-country basis (opposed to DHS and MICS, which provide global access for free with one automated request).

Without access to LSMS microdata, we were left with two pieces of information from PovCalNet to achieve our ultimate objective; firstly, the dollar threshold under which the poorest 20% live ($2.38 for 2012; or $2.56 for 2013), and, subsequently, the absolute number of people in each country that fall under that 20% threshold. Neither of these pieces of information were independently useful; we couldn't use our dollar threshold since the DHS and MICS only have a non-monetary wealth index based on asset ownership (generally arbitrarily running from –2 to 2).

The comparative wealth index

Our first attempted solution was inspired by an experimental methodology[3] created by Rutstein and Staveteig for the DHS. We called this methodology the Comparative Wealth Index (CWI), as opposed to each individual country's "relative wealth index". The idea is to find common points in each society, find the relative wealths at those points, and regress the points to find how to fit one society to another. In Rutstein and Staveteig's example, they took Vietnam as the base country because its gross domestic product

(GDP) per capita was roughly in the middle of the global distribution (so other countries would regress up and down to meet it in the middle). Their cut-points were defined by points of unmet basic-needs to capture the bottom of the distribution (inadequate wall material, crowded homes, inadequate sanitation, and high dependency on few income-earners) as well as the ownership of some key assets to capture the upper end of the distribution (television set, refrigerator, car/truck, and telephone).

Using this methodology, we took East Timor as our base as it was in the middle of the distribution for 2012. We were able to rank all 1.3 million households in our surveys by wealth. Our theory was that we could then take the poorest 27.7% (since our surveys cover about 5.1 billion people, or 72% of the global population) of those households and we would roughly end up with the poorest 20% globally. Unfortunately, in addition to being incredibly complex, this methodology was heavily dependent on the wealth indices of individual countries being well constructed. However, we found that this was not necessarily the case. In several surveys, poor data quality led to poor wealth indices; Jordan saw negative values as large as 10, and so our poorest family in the whole world ended up being a rural Jordanian family that owned a car. This methodology was also originally designed to make comparisons between small sets of countries, rather than attempting to understand the entire global distribution of wealth.

The one technique we retained from this methodology was the idea that you could apply relative percentiles to the poorest segment of society. For example, in the construction of the comparative wealth index cut-points, Rutstein and Staveteig didn't simply take the mean wealth score of individuals that owned a TV; they first found that 70% of the population owned a TV, and then used the relative wealth score of the family that sat at the 70th percentile in terms of wealth.

Monetisation of the relative wealth index

Our second idea was to attempt to monetise the CWI so we could compare it directly with the $2.38 cut-point established by PovCalNet. After rank-ordering every household in our set, we calculated the mean CWI for every survey and joined it with a measure of gross national income (GNI) per capita. The relationship was statistically significant with a p-value (probability of having no relationship) of less than 0.01, but the CWI was only able to explain about 51% of the observed variance in GNI per capita.

Figure 1: CWI and GNI per capita

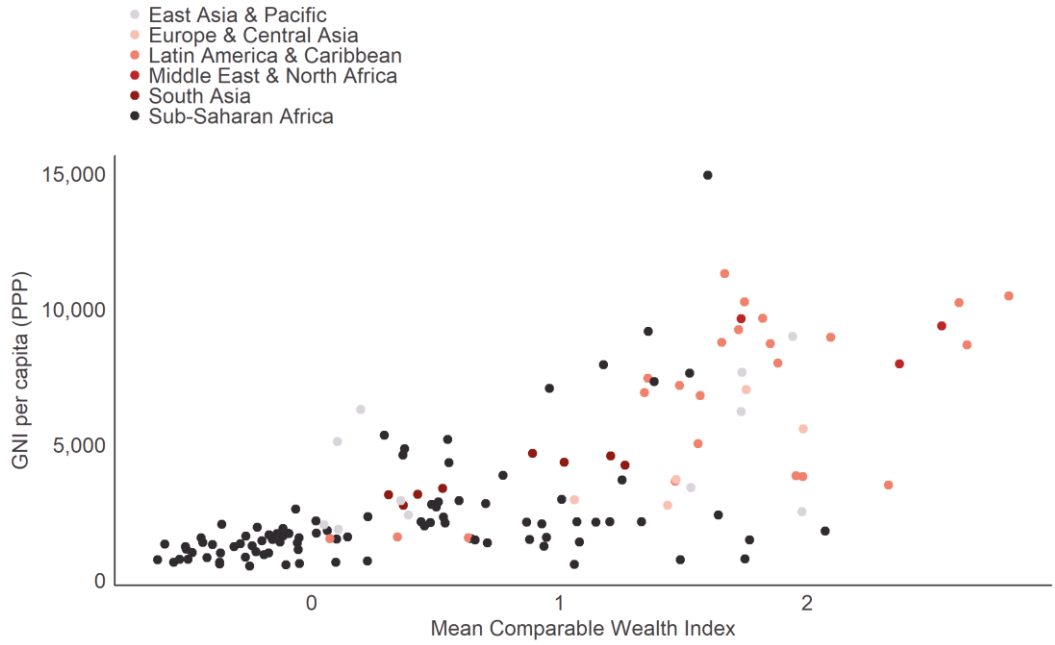

Source: DI based on DHS, MICS, and PovCalNet.

Notes: Mean CWI is along the X axis, while GNI per capita is along the Y. World Bank regional classifications are denoted via colour.

Once the relationship between our non-monetary CWI and GNI per capita was established, we used the regression coefficients to transform the CWI into an approximation of GNI per capita. We were able, therefore, to create poverty headcounts, as well as poverty gaps using only household surveys and our assumed relationship between wealth and GNI.

The advantage to this approach was that we could assign an approximate dollar value to each household based solely on the number and type of assets they owned. However, this method suffered many of the same pitfalls as the CWI itself. It added even further layers of complexity and fell short of approximating good poverty headcounts in countries where the wealth index itself was inadequate (especially for countries with large rural populations).

Equivalent percentiles

Working from the percentile method we learned from the CWI, we set out to make a simpler comparison method that would avoid regressions and more complicated transformations. Since we used PovCalNet to not only define the P20 poverty line, but also determine the P20 headcounts on a country-level, we knew roughly how many people in each country should be considered in the poorest 20%. We made the simple assumption that the poorest people as measured by relative wealth indices in any country

should be roughly the same group as the poorest as measured by income or consumption. With this assumption in place, we were able to measure the demographics of the P20 in each given country by assigning poverty status to the poorest percentile by wealth that matched PovCalNet's poverty headcount. For example, in 2012 the P20 headcount for Uganda was approximately 47%, so we assigned P20 status to the weighted 47th percentile of respondents in the Uganda DHS. This is similar to the method that Adam Wagstaff tested in his 2003 publication, *Child health on a dollar a day: some tentative cross-country comparisons*. To work around the lack of income or consumption data in the DHS, Wagstaff rank-ordered households by the wealth index, and then made the assumption "that a child's rank in the wealth distribution is close to his or her rank in the income or consumption distribution".[4]

While this method is rather simple and intuitive, the underlying assumption is potentially flawed. In a case study using an LSMS survey conducted in Russia, we were able to calculate poverty status by income, expenditure, and wealth. While income and expenditure both agreed upon the poverty headcounts (3.12% and 3.06% respectively), neither of those cohorts overlapped at all. Not a single household considered poor by income standards was also considered poor by expenditure standards. Only 0.14% of the population was considered poor by wealth standards, which also did not overlap with expenditure or income poverty.

Figure 2: Russia LSMS Euler Diagram

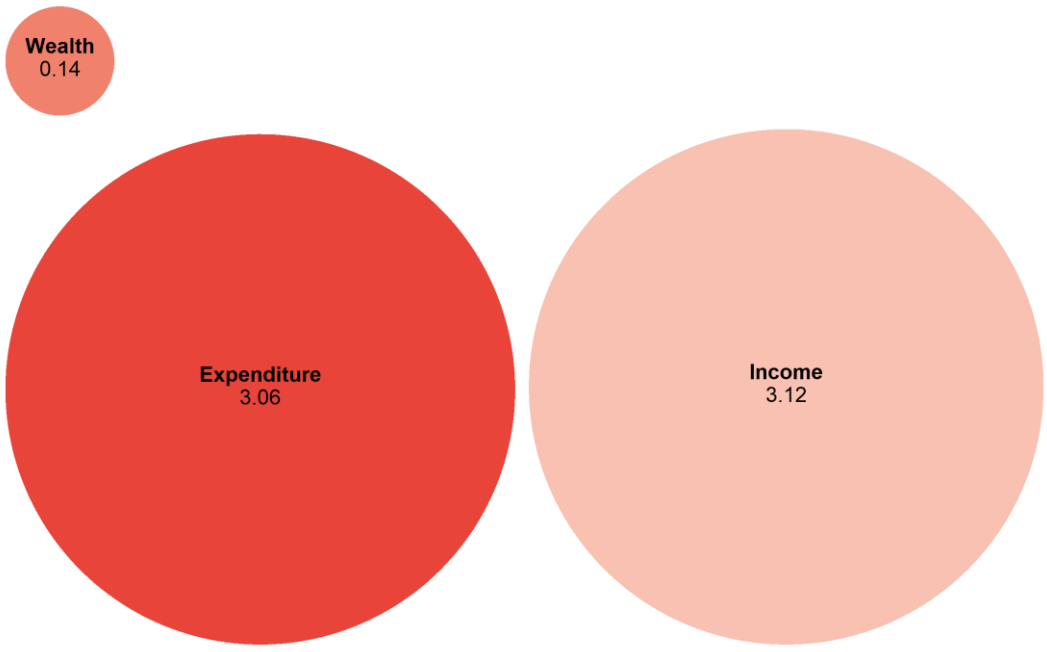

Source: Russia 1992–2005 LSMS.

Additionally, a case study using Ethiopia's LSMS survey showed a very weak relationship between the wealth index and per capita expenditure. The wealth index was a statistically significant component of expenditure (with a p-value less than 0.001), but it was only able to explain 0.7% of the variation in expenditure. Also, when our poverty-status methodology was applied to the wealth index, the wealth index was only able to predict a respondent's expenditure poverty status 52.3% of the time (slightly better than a coin flip at 50%).

Figure 3: Per capita expenditure vs. wealth index in Ethiopia

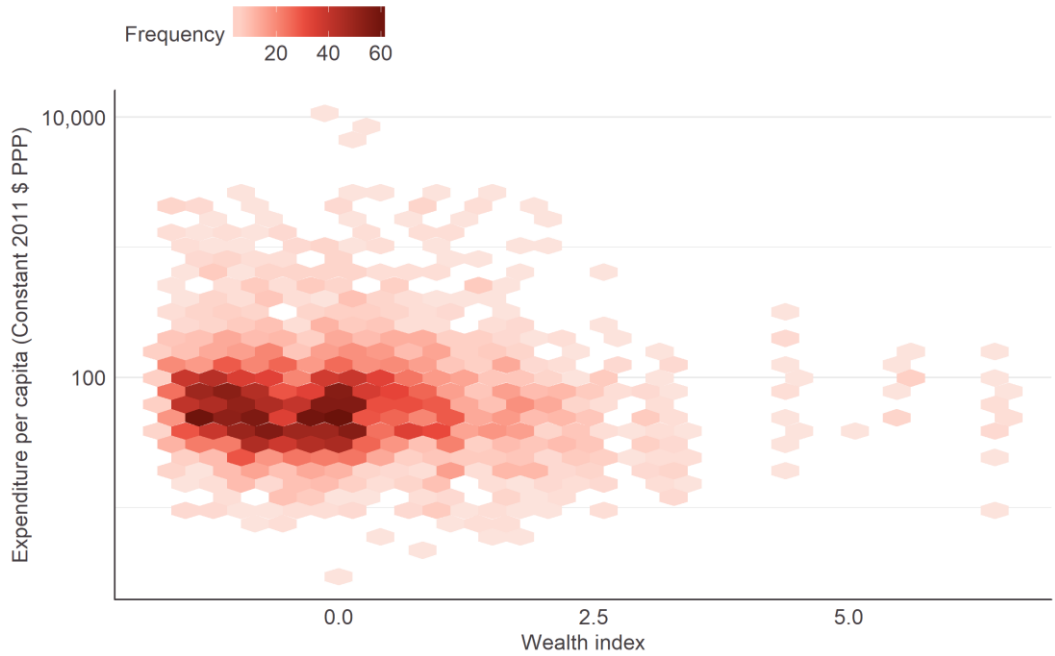

Source: Ethiopia LSMS 2011.

Conclusion

Poverty, like many aspects of humanity, has always been difficult to measure. Multidimensional and dashboard-type approaches may have the most potential for measuring poverty in the future, but we ultimately wanted to abstract poverty from the deprivations it creates.

The next-best solution involves precisely measuring income or consumption, like the LSMS surveys and PovCalNet, upon which our final analysis was based. Even in these ideal cases, measurements of poverty are not without their own sets of problems. Income tends to over-estimate poverty due to the inclusion of debt, while consumption can be hard to estimate with incomparable baskets of goods between the rich and the poor. Our methodology may represent just one more imperfect addition to the methods described above, but it's the next-best alternative given the current lack of open-access monetary data.

Perhaps if sources of monetary poverty were to open their data and methodologies to other researchers, and more surveys were able to incorporate better ways of estimating welfare, methods like the P20 could be rendered unnecessary. Until that point, however, we hope that our research acts as a good stop-gap solution to estimating the effects of poverty in the absence of better data.

Notes

1. IPUMS International https://www.idhsdata.org/idhs/
2. World Bank's PovCalNet http://iresearch.worldbank.org/PovcalNet/home.aspx
3. Making the Demographic and Health Surveys Wealth Index Comparable https://dhsprogram.com/pubs/pdf/MR9/MR9.pdf
4. Child health on a dollar a day: some tentative cross-country comparisons https://www.sciencedirect.com/science/article/pii/S0277953602005555

Development Initiatives (DI) is an independent international development organisation working on the use of data to drive poverty eradication and sustainable development. Our vision is a world without poverty that invests in human security and where everyone shares the benefits of opportunity and growth.

We work to ensure that decisions about the allocation of finance and resources result in an end to poverty, increase the resilience of the world's most vulnerable people, and ensure no one is left behind.

Contact
Alex Miller
Data scientist
+1 (571) 232 8935/Alex.Miller@devinit.org

To find out more about our work visit:
www.devinit.org
Twitter: @devinitorg
Email: info@devinit.org

Development Initiatives is the trading name of Development Initiatives Poverty Research Ltd, registered in England and Wales, Company No. 06368740, and DI International Ltd, registered in England and Wales, Company No. 5802543. Registered Office: North Quay House, Quay Side, Temple Back, Bristol, BS1 6FL, UK.

UK OFFICE
Development Initiatives
North Quay House
Quay Side, Temple Back
Bristol, BS1 6FL, UK
+44 (0) 1179 272 505

KENYA OFFICE
Development Initiatives
Shelter Afrique Building
4th Floor, Mamlaka Road
Nairobi, Kenya
PO Box 102802-00101
+254 (0) 20 272 5346

DEVELOPMENT RESEARCH AND TRAINING (DRT)
Ggaba Road, Mutesasira Zone, Kansanga
PO Box 22459
Kampala, Uganda
+256 (0) 312 – 263629/30
+256 (0) 414 – 269495
www.drt-ug.org

US OFFICE
Development Initiatives
1110 Vermont Ave NW,
Suite 500, Washington DC
20005, US