

### **Why Measure Poverty?**

To design effective policies and strategies aimed at reducing poverty, it is critical to understand the characteristics of poverty in a country or in a geographic area. Poverty measurement techniques can help shed light on whether or not poverty is increasing or decreasing, or whether general economic growth has benefited the poor. Poverty profiles of countries or geographic regions can help governments identify potential target groups for specific programs by region, employment, education and gender. With good poverty data, we can evaluate how the poor are affected by economy-wide policies, and in turn, how to develop strategies for economic growth that have a better chance of benefiting the poor.

### **The Concept of Well-being**

To measure poverty levels and characteristics of poverty, we begin with a definition of well-being, a concept which can be approached from many different angles. We can think of well-being as command over commodities or resources, or access to assets, or ability to function in society, and we can think of poverty as a lack of each of these things. Generally, poverty measurement focuses on assessing whether households or individuals have enough resources to meet their needs. Many techniques for measuring poverty focus on the monetary aspects of poverty, but we know that non-monetary dimensions are important. Poverty is associated not just with lack of resources that can be easily quantified, but also with lack of access to health services, education, and with low self-confidence and a sense of powerlessness.

There are other concepts of well-being that stretch beyond poverty, which can be thought of as monetary and non-monetary dimensions. For instance, inequality within a country is a measure of the distribution of resources across the population, and high levels of inequality can be viewed as related to well-being of those at the bottom of the resource scale. Vulnerability, which is generally defined as the risk of being poor, or of falling deeper into poverty, is another important factor related to well being.

### **Steps in Measuring Poverty**

The three main steps to be taking in measuring poverty are:

- 1) define an indicator of welfare
- 2) establish a minimum acceptable standard of that indicator to separate the poor and the non-poor (often known as the poverty line) and
- 3) generate a summary statistic which aggregates the information we get from looking at the distribution of the welfare indicator that we have chosen, and it's position relative minimum acceptable standards.

We have discussed the broader concept of well-being that forms that basis our of overall approach to poverty measurement, but the most common approach to measuring poverty focuses on economic welfare, which looks at monetary aspects of poverty.

Measures of economic welfare are usually based on household consumption expenditure or household income, to are which each resident in the household is assigned a share of the total amount.

### **Consumption-based Measures versus Income-Based Measures**

In general, consumption-based measures of living standards are generally preferred to income-based measures because the incomes of the poor often vary considerably over time, which might be the case, for example, in rural areas where income is generated primarily from rain-fed agriculture. Consumption measures reflect to a greater degree the level of resources households control, and may also reveal information about incomes in both the past and the future. Consumption includes both goods and services that are purchased, and those that are provided from ones' own production ("in-kind").

Consumption is often more practical to measure than income as households may be more able or willing to recall what they have spent rather than what they have earned.

It is possible to measure household welfare by looking at household income, but there are some complications that make it often a second choice for poverty analysts. For one, as mentioned above, households maybe more be more willing to recall what they have spent then what they have earned, since people may fear that a tax collector might hear of their total income, or because people may be reluctant to report income they have earned illegally, for instance from smuggling. In addition, certain types of income are not always easy to measure, such as farm income, or changes in the value of farm inventory or changes in the value of a house that one owns.

For these reasons, consumption expenditure is often the indicator used to measure economic well-being. We face several major challenges when we use consumption a our measure of well-being. In order to make comparisons across households using consumption/expenditure measures of poverty, it is necessary to calculate the value of durable goods owned by the household in the year of measurement, including both the depreciation of the item during the year, and the interest cost of having locked money up in the item, as opposed to in an interest-bearing account. It is important to measure durable goods because we know that even some households that cannot afford adequate quantities of food devote some expenditure to other items, (such as clothing and shelter) so we can assume that these items represent very basic needs of the household, and should be included in our poverty measures. This means that we need to do a calculation of the value and of the depreciation for each durable good that that household owns, and the potential for error in this process can be quite large. We also need to include housing services in consumption, and generally we value these services by asking if you paid rent instead of owning your home, how much would you need to pay? Another element of consumption is expenditures on weddings and funerals.

In brief, if we choose to assess poverty based on household consumption or expenditure per capita, as discussed above, we can use an expenditure function in our analysis. In simple terms, an expenditure function shows the minimum expense required to meet a given level of utility  $u$ , which is derived from a list of goods  $x$ , at prices  $p$ . Using an expenditure function, which incorporates a given level of prices of goods and services, and takes into consideration household characteristics (e.g. number of adults in the household, number of young children, etc.) we can measure the amount of spending that is needed to reach a certain level of utility. Typically, we compute the actual level of household spending from household survey data that include information on consumption, and then construct per capita household consumption for every individual in the household. This approach assumes that all individuals in the household have the same needs, which in reality is not true, since household members of different ages, for instance, will have different needs.

### **Accounting for household consumption differences**

Since households differ in size and composition, a simple comparison of aggregate welfare can be misleading about the well-being of individuals in a give household. The most straightforward way of dealing with this problem is to convert from a measure of household consumption to a measure a individual consumption by dividing household expenditures by the number of people in the household. This is the easiest procedure, but it does not take into consideration the obvious fact that different individuals have different needs (e.g. a young child typically needs less food than an adult) and also that there are economies of scale in consumption of non-food items. We deal with this problem by assigning a system of weights using an equivalence scale that measures the number of adult males that the household is equivalent to. This way, each member of the household counts as some fraction of an adult male. Selection of the scale is not straightforward because there is no generally consensus on an appropriate scale, or on its usefulness in practice, and there are two general methods for dealing with the problem: pick one of the scales that is already being used (for example the “OECD scale” which assigns a two-adult household an “Adult Equivalent” of 1.7, and a three adult household and AE of 2.4) or to estimate an equivalence scale for the population by looking at house aggregate household consumption of various goods during some survey period tends to vary with household size and composition.

### **Other measures of household welfare**

There are other measures of household welfare that we can use, but each also has it's own drawbacks. We can measure well-being through the number of calories consumed per person per day, for instance, or food consumption as a fraction of total expenditure. We can also measure outcome rather than inputs, such as rates of malnutrition rather than food intake. This can be problematic, however, as developing measures by which we can compare children by size and growth to assess their nutritional needs can be very difficult. We can evaluate welfare at the community level by looking at life expectancy, or infant mortality rates, or school enrollment rates, but these are not replacements for consumption per capita. They can be used to get a more complete and multidimensional

view of the well-being a population. In sum, there is not perfect measure of well-being. This does not mean we should avoid measuring poverty, but it does imply argue for approaching poverty measures with a certain degree of caution, and for looking carefully at how the measures were constructed.

## **Household surveys**

The main instruments for collecting data to support poverty analysis are household surveys, and it is important to have an understanding of the key issues associated with setting up and interpreting the data obtained from these surveys.

Sampling – Measures of poverty and inequality are always based on data that have been collected using sample surveys of households. This has two important implications 1) it means that actual measures of poverty and inequality are generated from a sample and not from the entire population. For these this reason, they estimate the true population parameters with some error. This should be noted with, for example, with confidence intervals, which indicate the chance that a statistics is accurate within a certain range. Survey data may also need to be weighted to get the right estimates of certain measures such as mean income, or poverty rates, which means that information about the sample frame (which is essentially the design of the sample selected) needs to be clearly documented.

Coverage of goods and income sources in the survey should be comprehensive, covering both food and non-food goods, and all income sources. Consumption should cover all monetary expenditures on goods and services consumed plus the monetary value of all consumption from income in kind, such as food produced on the family farm, and the rental value of owner-occupied housing. Similarly, the definition of income should include income in kind. Valuation of non-market goods can be quite complicated, and there is no widely preferred method, but there are a number of techniques to choose from.

Another issue is how to compare households at similar consumptions levels. Households size and composition vary across households, as do prices with which they are faced and wage rates. As a result it takes different resources to make ends meet for different households, that is to say, two households with the same level of expenditure may achieve different levels of well-being. The most general approach to measuring across households is to look at demand patterns to reveal consumer preferences over market goods, and to look at the minimum total expenditure that would be required for a consumer to achieve his or her actual utility level, but evaluated at pre-determined and arbitrary reference prices and demographics fixed over all households. Deflators are used in these calculations, which depend on preferences. There are many problems with this approach....

Ideally, we should not have to rely solely on a household level survey in making interpersonal comparisons of welfare. A separate community survey (done at the same times as the interviews, and possibly by the same interviewers) can provide useful supplementary data on the local prices of a range of goods and local public services. By

matching these to the household level data, one can greatly improve the accuracy and coverage of household welfare assessments. This has become common practice in the World Bank's Living Standards Measurement Study surveys.

### **Living Standard Measurement Study Surveys**

The LSMS surveys ask about a wide variety of topics, not just demographic characteristics or specific narrow issues. The LSMS survey includes a household questionnaire, which often runs 100 pages or more, and which can be adapted to the needs of each country. There is also a community questionnaire, which asks community leaders (teachers, health workers, village officials) for information about the whole community, such as the number of health clinics, access to schools, tax collection, demographic data, and agricultural patterns. The third part is a price questionnaire, which collects information about a large number of commodity prices in each community where the survey is undertaken. This is useful because it allows analysts to correct for differences in price levels by region, and over time.

### **Introduction to Poverty Lines**

Once we have chosen a measure of well-being (consumption expenditure, say) the next step is to choose a poverty line. Households whose consumption falls below this line are considered to be poor. A poverty line is obtained by specifying a consumption bundle considered adequate for basic consumption needs and then by estimating the costs of these basic needs. In other words, the poverty line represents a minimum standard required by an individual to fulfill his or her basic food and non-food needs.

Once we have computed a household's consumption, we need to evaluate if that amount places the household "in poverty", or defines the household as "poor". The threshold that we use for this is the poverty line. In reality, it makes sense to define more than one poverty line, since there are many ways to evaluate whether or not an individual or household is in or has come out of poverty. We can generate a poverty line for each household, and adjust it from household to household to take into consideration the differences in prices that they face as well as differences in the demographic composition of the households. For example, a small household in a rural area may face low housing costs and relatively modest food prices. Thus their poverty line may be low compared to a large household living in a city where housing more expensive and food prices are perhaps higher. This gives a different poverty line for each household. A second approach is to construct one per capita poverty line for all individuals, but adjust per capita income for differences in prices and household composition. The adjusted per capita income is then compared with the one poverty line to determine if the individual is living below the poverty line.

A poverty profile done in Cambodia in 1999 used an approach somewhere in between, by constructing one poverty line for each of the three major regions in the country based on the prices prevailing in those regions. Whether a household in any given region is poor,

is then determined by comparing its expenditure per capita with the appropriate regional poverty line.

Nominal poverty lines (as opposed to relative ones) can change over time, due to changes in prices (typically from inflation), or because the real poverty threshold is revised over time. This raises the question of whether we should look at relative, or absolute poverty lines.

### **Relative poverty Lines**

Sometimes we are interested in focusing on the poorest segment of the population, those that are relatively poor. It is often helpful to have a measure such as this in order to identify those who are poor today, and to help us design targeted programs that are geared toward the poor. Relative poverty lines need to be tailored to the overall level of development of a country, as countries using them will tend to revise them upwards with increases in per capita consumption. For instance, a \$1 per day poverty line might be useful in Vietnam, where 27% of the population would be considered poor by this standard in 1998, but this line would be of little relevance in the United States, where almost nobody would be poor by this standard.

### **Absolute Poverty Lines**

An absolute poverty line is fixed over time, which enables poverty analysts to judge the impact of anti-poverty policies over time. Comparisons of poverty rates between countries can only be made if the same absolute poverty line is used in both countries. The World Bank commonly uses two absolute poverty lines 1) people living on less than \$1 dollar per day, and 2) people living on less than \$2 dollars per day.

Several important conceptual problems come up when working with absolute poverty lines. For one, it is often difficult to agree on how to define a “standard of living” that the poverty line measures. In practice, almost all poverty lines are set in terms of the cost of buying a basket of goods (a “commodity-based poverty line”), but one can argue that individual “standards of living” or “utilities” are interdependent. That is, a household of four with an income of \$12,000 per year would not be considered poor in Indonesia, but when this household compares its position with average incomes in the US, it may feel very poor. This implies that a commodity-based poverty line would rise as a country becomes more affluent, because the minimum resources needed to participate fully in society probably rises over time. Thus the absolute poverty line is more of a real poverty line. In addition, it may be difficult to define the correct commodity value of the poverty line, because both size and demographic composition of households vary.

One way to determine poverty lines is to set “objective” poverty lines, that is, to set the poverty line at a level that enables individuals to achieve certain capabilities, including a healthy and active life and full participation in society. A common way of approaching this is to begin with nutritional requirements. Two common methods are the Food Energy Intake (FEI) method, and the Cost-of-Basic-Needs (CBN) approach.

The goal of the FEI is to find the level of consumption expenditure (or income) that allows the household to obtain enough food to meet its energy requirements. First one needs to determine the amount of food that is adequate (for example, Vietnam sets this value at 2,100 calories per day, while recognizing that different individuals may need more or less food than this). After the calorie level is set, food expenditure lines can be estimated by obtaining the cost of obtaining the 2100 calories of food per day. There are important weaknesses to this approach however, so it should only be used when other alternatives are feasible. For instance, important differences exist between urban and rural food consumers not just access to food, but also in prices, both of which will influence the type of calories that people consume, making comparisons difficult.

The CBN approach uses a consumption bundle that is deemed adequate, including both food and non-food components, and it estimates the cost of the bundle for each subgroup of a population (urban/rural, each region, etc.). The key difference in this approach is the poverty line is measured in money, so it does not insist (as with the FEI) that each basic need be met, only that it *could* be met. There are challenges with this approach, namely decisions on how to measure the non-food components of the poverty line (for example, countries in colder climates may put more emphasis on the cost of energy needed for home heating than countries in warmer climates).

## **Measures of Poverty**

Given information on per capita consumption, and a poverty line, there are several aggregate measures of poverty that can be computed.

### **Headcount index:**

By far the most widely-used measure is the headcount index, which simply measures the proportion of the population that is counted as poor, often denoted by  $P_o$ . The headcount index is simple to construct and easy to understand, which are important qualities. The measure has several weaknesses, however: the headcount index does not capture the depth of poverty, that is, if a somewhat poor household were to give to a very poor household, the headcount index would remain unchanged, even though poverty as a whole has dropped. In using the headcount index, we can see that the easiest way to reduce poverty would be to target benefits to people just below the poverty line because they are the ones who can be moved across the poverty line most cheaply. In addition, what is calculated using the headcount index is the percentage of *individuals* who are poor, and not the percentage of households, which means that we are making a critical assumption that all household members enjoy the same level of well-being, which may not be the case.

## **Poverty Gap Index:**

Another measure of poverty is the poverty gap index, which adds up the distances that poor people fall from the poverty line, and expresses that as a percentage of the poverty line. The poverty gap index is thought of as a way to measure the total cost of bringing each poor member of a society up to the poverty line, but it should be noted that this is not a precise measure because it requires exact information on each poor member of society, and this assumes that the government has a lot of information. The main drawback of the poverty gap index is that because it is an average measure over all people of the gaps between poor people's standard of living and the poverty line, it is not capable of capturing inequality among the poor. For instance the poverty gap index maybe the same for two countries, but one of the countries may have many more very poor people, and more people close to the poverty line, which would generate a poverty gap index that is similar to a country where most of the poor people are considered moderately poor. The main advantages of the poverty gap index are that it gives an idea of the minimum amount of financial resources that are needed to tackle poverty problems, and that it highlights the importance of identifying the characteristics of the poor, as it demonstrates the potential savings of well-targeted poverty alleviation programs.

## **Squared Poverty Gap**

To solve this problem of identifying inequality among the poor, some researchers use the squared poverty gap index, which is simply a weighted sum of poverty gaps (as a proportion of the poverty line), where the weights are proportionate poverty gaps themselves. For instance, a poverty gap of 10% of the poverty line is given a weight of 10%, while one of 50% is given a weight of 50%, which contrasts with the poverty gap index, where they are weighted equally. By squaring the poverty gap index, we put more weight on the observations that fall well below the poverty line.

## **Measures of Inequality**

Measures of poverty focus on the situation of individuals or households at the bottom of the income distribution, but sometimes we are interested in measuring the distribution of income over an entire population, and not just looking at the distribution below the poverty line. Evaluating inequality in a population, and in particular looking at how inequality changes over time, can provide useful information to analysts and policymakers designing policy interventions. For instance, it would be useful to look at the distribution of income in a certain economic sector or population group and how it changes over time relative to changes in income over the entire population, or within a sector. If rural incomes have increased at the same time as income inequality has increased in the rural sector, this effect could be linked to some policy reform, say agricultural price policy, with the benefits of the change accruing to wealthier farms and not to the poor and less-efficient farmers.

The simplest way to evaluate inequality over a population is to divide the population into fifths (quintiles) from poorest to richest, and reporting the levels or proportions income (or expenditure) for each level.

The most widely used measure of inequality is the Gini coefficient, a measure which is based on the Lorenz curve, a cumulative frequency curve that compares the distribution of a specific variable (for instance, income) with the uniform distribution that represents equality. One axis of the graph measures the cumulative percentage of a population, graphed from richest households to the poorest and the other axis measures the cumulative percentage of either income or expenditure, whatever it is that is being measured. The Lorenz curve would be a 45 degree line if all income recipients (or expenditure levels of individuals or households measured) had equal shares, that is if 10% of the population had a 10% share of income, and 20% had a 20% share of income. The extent to which the Lorenz curve deviates from the 45% line (of absolute equality) indicates the degree of inequality within the population. So the Gini coefficient varies from zero, which indicates perfect equality in a population, with every household earning exactly the same income, or spending exactly the same amount, to one, which implies absolute inequality, with a single household earning a country's entire income. The Gini coefficient can be used to compare the relative rate of inequality in one region of a country to another region in the same country, or of one country to another, but the Gini coefficients of various regions in one country cannot be added up. The Gini cannot be broken down by population groups or income sources, or in other dimensions.

**Poverty Profiles** – A poverty profile is a comprehensive poverty comparison, showing how poverty varies across various sub-groups of society, such as region of residence or sector of employment. It sets out the major facts on poverty and inequality, and then examines the pattern of poverty to see how it varies by geography (by region, urban/rural, mountain/plain, etc) by community characteristics (e.g. in communities with and without a school, etc.) and by household characteristics (e.g. by education of household head, by size of household). A well-presented poverty profile can be immensely informative and useful in assessing how sectoral or regional patterns of economic change are likely to affect aggregate poverty, although it typically uses basic techniques, such as graphs and tables. For example, regional poverty comparisons are important for targeting development programs to poorer areas. A recent poverty study for Cambodia showed that headcount poverty rates were highest in the rural sector and lowest in Phnom Penh in 1999. Approximately 40% of the rural population, 10% of Phnom Penh population, and 25% of other urban residents live in households that are below the poverty line. Poverty profiles can show access to services by households in different regions, and can shed light on the relationship between education levels in a household and the likelihood that a household will be poor. Key questions that can be included in a poverty profile are: -- Which are the most important goods in the consumption basket of the poor? To what public services do the poor have access? What is the quality of the service? What are the main sources of income for the poor? Can the poor access formal or informal credit markets? Are certain population groups in society at a higher risk of being poor than others are?

## **Divergence Between Household Surveys and National Accounts**

One issue in poverty measurement that has become a topic of debate in recent years is related to the question of how best to measure whether or not economic growth is benefiting the poor, how to achieve growth that benefits the poor if it not happening naturally. The central question is how to assess whether income or consumption of the poor increases by at least as much as for the economy as a whole. Studies from the 1990s showed that in many cases economic growth did not bring with it expected rates of poverty reduction, although there was no evidence of rising inequality. The main poverty measurement issues revolve around the use of Household Budget Survey (HBS) data versus data from National Accounts in assessing consumption levels, and how to deal with biases that may occur in the HBS data. There is strong evidence that rich households are less likely to comply with HBS reporting, for a variety of reasons, perhaps in an effort to conceal income or consumption. The result is that upper incomes are underrepresented in the HBS, but National Accounts will pick up the transactions of these upper incomes through key economic aggregates. Thus there is a tendency of “average” national accounts consumption to be higher than “average” consumption from the HBS. But because the HBS is better at capturing consumption that takes place in the informal sector, it is argued that while the National Accounts measure of consumption maybe a more accurate measure of the consumption of a population, the HBS is better measure of consumption for the poor. Research by Bhalla (get citation) in which he uses the National Accounts measure to evaluate poverty in India generated much lower levels of measured poverty than previous studies, and also predicts much more optimistic trends in poverty reduction than standard projections. This is debate is particularly important for countries undergoing high rates of economic growth, but the issues are relevant for other countries as well.