

Measuring Poverty

Summary

The first step in measuring poverty is defining an indicator of welfare such as income or consumption per capita. Information on welfare is derived from survey data. Good survey design is important. Although some surveys use simple random sampling, most use stratified random sampling. This requires the use of sampling weights in the subsequent analysis. Multistage cluster sampling is also standard; it is cost-effective and unbiased, but it lowers the precision of the results, which calls for some adjustments when analyzing the data.

The World Bank-inspired Living Standards Measurement Surveys (LSMS) feature multitopic questionnaires and strict quality control. The flexible LSMS template is widely used.

Income, defined in principle as *consumption + change in net worth*, is generally used as a measure of welfare in developed countries, but it tends to be seriously understated in less-developed countries. Consumption is less understated and comes closer to measuring *permanent income*. However, it requires one to value durable goods (by assessing the implicit rental cost) and housing (by estimating what it would have cost to rent).

While consumption per capita is the most commonly used measure of welfare, some analysts use consumption per adult equivalent, in order to capture differences in need by age, and economies of scale in consumption. The Organisation for Economic Co-operation and Development (OECD) scale ($= 1 + 0.7 \times (N_A - 1) + 0.5 \times N_C$) is popular, but such scales are controversial and cannot be estimated satisfactorily.

Other popular measures of welfare include calorie consumption per person per day, food consumption as a proportion of total expenditure, and nutritional status as measured by stunting or wasting. However, there is no ideal measure of well-being, and analysts need to be aware of the strengths and limitations of any measure they use.

Learning Objectives

After completing the chapter on *Measuring Poverty*, you should be able to

1. Summarize the three steps required to measure poverty.
2. Recognize the strengths and limitations arising from the need to use survey data in poverty analysis, including the choice of sample frame, unit of observation, time period, and choice of welfare indicators.
3. Describe the main problems that arise with survey data, including:
 - survey design (sampling frame/coverage, response bias)
 - stratification
 - multistage cluster sampling.
4. Explain why weighting is needed when surveys use stratified random sampling.
5. Describe and evaluate the use of equivalence scales, including the OECD scale.
6. Define consumption and income as measures of welfare, and evaluate the desirability of each in the context of measuring well-being in less-developed countries.
7. Summarize the problems that arise in measuring income and consumption, and explain how to value durable goods and housing services.
8. Identify measures of household welfare other than consumption and income, including calorie consumption per capita, nutritional status, health status, and food consumption, as a proportion of total expenditure.
9. Argue the case that there is no ideal measure of welfare.

Introduction: Steps in Measuring Poverty

The goal of this chapter is to set out a method for measuring poverty. Given the enormous literature available on the subject, we simply set out the main practical issues, with suggestions for further reading for those interested in pursuing the subject more.

Three steps need to be taken in measuring poverty (for further discussion, see Ravallion 1998):

- Defining an indicator of welfare
- Establishing a minimum acceptable standard of that indicator to separate the poor from the nonpoor (the poverty line)
- Generating a summary statistic to aggregate the information from the distribution of this welfare indicator relative to the poverty line.

This chapter defines an indicator of welfare; chapter 3 discusses the issues involved in setting a poverty line; chapter 4 deals with measuring aggregate welfare and its distribution.

Household Surveys

All measures of poverty rely on household survey data, so it is important to recognize the strengths and limitations of such data and to set up and interpret them with care.

Key Survey Issues

Ravallion (1992) lists a number of issues related to surveys that require attention before one even attempts to measure or analyze poverty:

- *The sample frame:* The survey may represent a whole country's population, or some more narrowly defined subset, such as workers or residents of one region. The appropriateness of a survey's particular sample frame will depend on the inferences one wants to draw from it. Thus, a survey of urban households would allow one to measure urban poverty, but not poverty in the country as a whole.
- *The unit of observation:* This is typically the household or occasionally the individuals within the household. A household is usually defined as a group of persons eating and living together.
- *The number of observations over time:* Most surveys are single cross-sections, covering a sample of households just once. *Longitudinal surveys*, in which the same households or individuals are resurveyed one or more times (also called *panel data sets*) are more difficult to do, but these have been undertaken in a several countries (for example, the Vietnam Living Standards Surveys of 1992–93 and 1997–98, or parts of the Thailand Socioeconomic Surveys of 2002 and 2004).
- *The principal living standard indicator collected:* Most measures of welfare are based on household consumption expenditure or household income. Many surveys collect both, although this typically requires two interviews per household: to save on costs, some surveys gather data on either income or expenditure. Given budget constraints, there are always tradeoffs: Since a more detailed and complex questionnaire takes longer to administer, the sample size will have to be smaller, which reduces the precision of the statistics based on these data and limits the amount of disaggregation (for example, to the provincial level) that is possible.

Common Survey Problems

Several common problems arise when using and interpreting household survey data. We review these, organizing our thoughts largely along the lines set out in Ravallion (1992).

Survey Design. If the sample on which a survey is based is not random, then the resulting estimates of poverty are almost impossible to interpret. They are likely to be biased, but we do not know by how much.

A simple national random sample would create a list of everyone in the country and then randomly choose subjects to be interviewed, with each person having an equal chance of being selected. In practice, sampling always falls short of this ideal for three reasons. First, some people or households may be hard to find; for instance, most surveys interview people at their homes, but this completely overlooks homeless persons, a group that is likely to be poor.

Second, some of the surveys that have been used to measure poverty were not designed for this purpose in that their sample frames were not intended to span the entire population.

Examples: This is true of labor force surveys, which have been widely used for poverty assessments in Latin America; the sample frame is typically restricted to the “economically active population,” which precludes certain subgroups of the poor. To take another example, household surveys in the Republic of Korea have typically excluded one-person households from the sample frame, which renders the results unrepresentative.

Key questions to ask about any survey are the following:

- Does the sample frame (the initial listing of the population from which the sample was drawn) span the entire population?
- Is there likely to be a response bias? This may take one of two forms: *unit nonresponse*, which occurs when some households do not participate in the survey, and *item nonresponse*, which occurs when some households do not respond fully to all the questions in the survey.

Third, it is very often cost-effective deliberately to oversample some small groups (for example, minority households in remote areas) and to undersample large and homogeneous groups. Such *stratified random sampling*—whereby different subgroups of the population have different (but known) chances of being selected but all have an equal chance in any given subgroup—can increase the precision in poverty measurement obtainable with a given number of interviews. When done, it is necessary to use weights when analyzing the data, as explained more fully in the following section.

Sampling. Two important implications flow from the fact that measures of poverty and inequality are always based on survey data.

First, it means that actual measures of poverty and inequality are *sample statistics*, and so estimate the true population parameters with some error. Although it is standard practice to say that, for instance, “the poverty rate is 15.2 percent,” it would be more accurate to say something like “We are 99 percent confident that the true poverty rate is between 13.5 percent and 16.9 percent; our best point estimate is that it is 15.2 percent.” Outside of academic publications, such caution is, unfortunately, rather rare.

The second implication is that it is essential to know how the sampling was done, because the survey data may need to be weighted in order to get the right estimates of such measures as mean income or poverty rates. In practice, most household surveys oversample some areas (such as low-density mountainous areas, or regions with small populations), to get adequately large samples to compute tolerably accurate statistics for those areas. Conversely, areas with dense, homogeneous populations tend to be undersampled. For instance, the Vietnam Living Standards Survey of 1997/98 (VLSS98) oversampled the sparsely populated central highlands and undersampled the dense and populous Red River delta (Vietnam 2000).

In cases such as this, it is not legitimate to compute simple averages of the sample observations such as per capita income to make inferences about the whole population. Instead, weights must be used, as the following example shows.

Example: Consider the case of a country with 10 million people who have a mean annual per capita income of \$1,200. Region A is mountainous and has 2 million people with average per capita incomes of \$500; region B is lowland and fertile and has 8 million people with an average per capita income of \$1,375.

Now suppose that a household survey samples 2,000 households, picked randomly from throughout the country. The mean income per capita of this sample is the best available estimator of the per capita income of the population, and so we may calculate this and other statistics using the simplest available formulae (which are generally the ones shown in this *Handbook*). For example, the Vietnam Living Standards Survey of 1992–93 (VLSS93) essentially chose households using a simple random sample, using the census data from 1989 to determine where people lived; thus, the data from the VLSS93 are easy to work with, because no special weighting procedure is required.

Further details are set out in table 2.1. If 400 households are surveyed in Region A (one household per 5,000 people) and 1,600 in Region B (one household per 5,000 people), then each household surveyed effectively “represents” 5,000 people; a simple average of per capita income (\$1,215.60), based on the survey data, would then generally serve as the best estimator of per capita income in the population at large, as shown in the “case 1” panel in table 2.1.

Table 2.1 Illustration of Why Weights Are Needed to Compute Statistics Based on Stratified Samples

	Region A	Region B	Whole country
Population (million)	2.0	8.0	10.0
True income/capita (\$/year)	500	1,375	1,200
Case 1. Simple random sampling. Use simple average.			
Sample size (given initially)	400	1,600	2,000
Estimated total income, \$	196,000	2,235,200	2,431,200
	= 400 × 490	= 1,600 × 1,397	= 196,000 + 2,235,200
Estimated income/capita, (\$/year)*	490	1,397	1,215.6
			= 2,431,200/2,000
Case 2. Stratified sampling.			
Sample size (given initially)	1,000	1,000	2,000
Estimated total income, \$	490,000	1,397,000	1,887,000
	= 1,000 × 490	= 1,000 × 1,397	= 490,000 + 1,397,000
Case 2a. Stratified sample, using simple average. This is incorrect, so don't do this!			
Estimated income/capita (\$/year)	490	1,397	943.5
			= 1,887,000/2,000
Case 2b. Stratified sampling, using weighted average. This is the correct approach.			
Weight (Based on population)	0.2	0.8	
	= 2.0/10.0=	= 8.0/10.0	
Estimated income/capita (\$/year)	490	1,397	1,215.6
			= (0.2 × 490) + (0.8 × 1,397)

Source: Example created by the authors.

* Estimated income per capita is likely to differ from true income per capita, due both to sampling error (only a moderate number of households were surveyed) and nonsampling error (for example, underreporting, poorly worded questions, and the like).

But now suppose that 1,000 households were surveyed in Region A (one per 2,000 people) and another 1,000 in Region B (one per 8,000 people). If weights were not used, the estimated income per capita would be \$943.50 (see the “case 2a” panel in table 2.1), but this would be incorrect. Here, a weighted average of observed income per capita is needed in order to compute the national average. Intuitively, each household sampled in Region A should get a weight of 2,000 and each household in Region B should be given a weight of 8,000 (see table 2.1). The mechanics are set out in the “case 2b” panel in table 2.1 and yield an estimated per capita income of \$1,215.60.

In picking a sample, most surveys use the most recent population census numbers as the sample frame. Typically, the country is divided into regions, and a sample is picked from each region (referred to as a *stratum* in the sampling context). Within each region, subregional units such as towns, counties, districts, and communes are

usually chosen randomly, with the probability of being picked being in proportion to population size. Such multistage sampling may even break down the units further, for example, to villages within a district.

At the basic level (the *primary sampling unit* such as a village, hamlet, or city ward), it is standard to sample households in clusters. Rather than picking individual households randomly throughout a whole district, the procedure is typically to pick several villages and then randomly sample 15 to 20 households within each chosen village. The reason for doing cluster sampling, instead of simple random sampling, is that it is far cheaper to survey several households in a small area than to have to find households scattered widely over a potentially very large area.

But the use of cluster sampling, which is now almost ubiquitous, has an important corollary: The information provided by sampling clusters is less reliable as a guide to conditions in the overall area than pure random sampling would be. To see this, compare figure 2.1.a (simple random sampling) with figure 2.1.b (cluster sampling). Although, on average, cluster sampling will give the correct results (for per capita income, for instance), so the expected mean values are unaffected—it is less reliable because we might, by chance, have chosen two particularly poor clusters, or two rich ones. Thus, cluster sampling produces larger standard errors for the estimates of population parameters. This needs to be taken into account when programming the statistical results of sample surveys. Not all statistical packages handle clustering; however, Stata deals with it well using the *svyset* commands (see appendix 3 for details).

Most living standards surveys sample households rather than individuals. If the variable of interest is household-based—for instance, the value of land owned per household or the educational level of the household head—then the statistics should be computed using household weights.

But a survey that samples households will give too little weight to *individuals* in large households. To see this, consider the realistic case of a survey that, at the village

Figure 2.1a Simple Random Sample

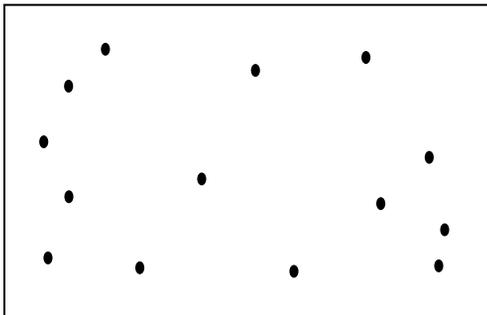
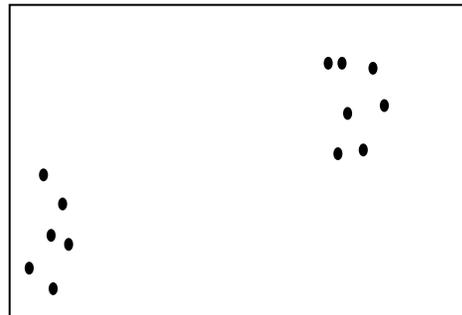


Figure 2.1b Cluster Sampling



Source: Authors.

level, randomly chooses 15 households for its sample. Perhaps one of these households consists of five persons; and another has just a single individual. In effect, the large household represents five times as many people as the small household, and if we are interested in individual-level measures—such as income per capita, or the age or gender of individuals—then we should put five times as much weight on the large household as the small. This calls for the use of *individual weights*, which are usually computed as the household weights times the size of the household. Most, but not all, statistical packages handle this easily, but the analyst still has to provide the appropriate instructions.

Goods Coverage and Valuation. It has been widely observed that the more detailed the questions about income and expenditure, the higher are the reported levels of income and expenditure. It follows that if economic welfare is to be measured satisfactorily, these questions must be comprehensive; to ensure comparability, they should not change over time. It is important to collect information on the volume and value of “own consumption”—such as food from the family farm that the household eats—since this is a component both of income and of expenditure (“in kind”). Such in-kind income/expenditure will typically have to be valued at local prices. It is also essential to collect enough information on housing (rent or current capital value if the household owns its residence), and the main durable goods (age, purchase price, current value), in order to be able to quantify these important components of expenditure and income.

Variability and the Time Period of Measurement. Income and consumption vary from month to month, year to year, and over a lifetime. But income typically varies more significantly than consumption. This is because households try to smooth their consumption over time, for instance by managing their savings, or through risk-sharing arrangements such as using remittances. In less-developed countries, most analysts prefer to use current consumption than current income as an indicator of living standards in poor countries for the following reasons:

- In the short run it reflects more accurately the resources that households control.
- Over the long term, it reveals information about incomes at other dates, in the past and future.
- In poor countries, income is particularly difficult to measure accurately, especially in agriculture and for pastoralists.

This does not mean that consumption is a perfect measure of well-being. Any household that is credit-constrained—and this is likely to be especially true of poorer households—will be limited in the extent to which it can smooth consumption over its lifecycle.

Comparisons across Households at Similar Consumption Levels. Households vary not only in their income or expenditure levels, but in size, in the prices they face, in the publicly provided goods (such as roads and schools) to which they have access, in the amount of leisure time they enjoy, and in the agreeableness of the environment in which they live (some areas are too hot or too cold or too dry or too flood-prone). This makes it difficult to compare household welfare across households. Thus an annual income of US\$1,000 might suffice for a couple living in a rural area, where food and housing are cheap, but it would be utterly inadequate for a family of four in an urban setting.

In practice, it is impossible to take all such factors fully into account, so all comparisons across households remain imperfect. However, some corrections are easier than others: It is relatively straightforward, data permitting, to correct for differences in the cost of living faced by households; and income or expenditure can be expressed in per capita (or per adult equivalent) terms, an issue to which we return in the following section. But researchers rarely include the value of publicly provided goods and services, mainly because these are hard to value and it is difficult to attribute usage. However, an attempt is made in benefit-incidence analysis, a topic addressed in chapter 15.

Review Questions

1. Which of the following is *not* one of the three steps involved in measuring poverty?

- A. Generating a summary statistic to reflect the degree of poverty.
- B. Computing a Gini coefficient.
- C. Defining an indicator of well-being.
- D. Establishing a minimum acceptable standard of well-being.

2. To measure poverty, one needs data based on surveys of individuals or households, and data from all such surveys are useful in measuring poverty.

- True.
- False.

3. You have information based on a regionally stratified random sample of households chosen with clustering.

- A. This means that regions in the country were first grouped together and then a simple random sample of households was chosen from each of these groupings.
- B. This implies that any measures of poverty will need to be computed using household or individual weights.
- C. This means that the standard error of measures of income will be smaller than would be the case with simple random sampling.
- D. The result is that one cannot generally break down poverty rates by region.

Key features of Living Standards Measurement Surveys

Motivated by the need to measure poverty more accurately, the World Bank has taken the lead in the development of relatively standard, reliable household surveys, under its Living Standards Measurement Surveys (LSMS) project. The electronic version of the books edited by Grosh and Glewwe (2000) includes an extensive sample questionnaire—best thought of as a template, not an off-the-shelf survey—and detailed chapters that deal with the design and implementation of such surveys. The LSMS surveys have two key features: multitopic questionnaires and considerable attention to quality control.

Multitopic Questionnaires. The LSMS surveys ask about a wide variety of topics, not simply demographic characteristics or health experience or some other narrow issue.

- The most important single questionnaire is the *household questionnaire*, which often runs to 100 pages or more. Although there is an LSMS template, each country needs to adapt and test its own version. The questionnaire is designed to ask questions of the best informed household member. The household questionnaire asks about household composition; consumption patterns, including food and nonfood; assets including housing; landholding and other durables; income and employment in agriculture/nonagriculture and wage/self-employment; sociodemographic variables, including education, health, migration, and fertility; and anthropometric information, especially the height and weight of each household member.
- 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 collections, demographic data, and agricultural patterns. Sometimes there are separate community questionnaires for health and education.
- The third part is the *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.

Quality Control. The LSMS surveys are distinguished by their attention to quality control. Key features include the following:

- Most important, they devote a lot of attention to obtaining a representative national sample (or regional sample, in a few cases). Thus, the results can usually be taken as nationally representative. It is surprising how many other surveys are undertaken with less attention to sampling, so one does not know how well they really represent conditions in the country.

- The surveys make extensive use of screening questions and associated skip patterns. For instance, a question might ask whether a family member is currently attending school; if yes, one jumps to page x and asks for details; if no, then the interviewer jumps to page y and asks other questions. This reduces interviewer errors.
- Numbered response codes are printed on the questionnaire, so the interviewer can write a numerical answer directly on the questionnaire. This makes subsequent computer entry easier, more accurate, and faster.
- The questionnaires are designed to be easy to change and to translate, which makes it straightforward to modify them in the light of field tests.
- The data are collected by decentralized teams. Typically each team has a supervisor, two interviewers, a driver/cook, an anthropometrist, and someone who does the data entry onto a laptop computer. The household questionnaire is so long that it requires two visits for collecting the data. After the first visit, the data are entered; if errors arise, they can be corrected on the second visit, which is typically two weeks after the first visit. In most cases the data are entered onto printed questionnaires and then typed into a computer, but some surveys now enter the information directly into computers.
- The data entered are subject to a series of range checks. For instance, if an age variable is greater than 100, then it is likely that there is an error that needs to be corrected.

This concern with quality has some important implications, notably the following:

- The LSMS data are usually of high quality, with accurate entries and few missing values.
- Since it is expensive to maintain high quality, the surveys are usually quite small; the median LSMS survey covers just 4,200 households. This is a large enough sample for accurate information at the national level, and at the level of half a dozen regions, but not at a lower level of disaggregation, such as a province, department, or county.
- The LSMS data have a fairly rapid turnaround time, with some leading to a statistical abstract (at least in draft form) within two to six months of the last interview.

Even when surveys are based on the LSMS template, it is still difficult to compare measures of consumption or income or poverty either across countries or within a country over time. This is because small differences in the way questions are phrased, or in the detail requested in the household interview, can have a substantial impact on the reported results. Seemingly minor changes—such as adding a few questions on tobacco use, or asking for more details about durable goods—can have effects on the

measurement of the variables that are large enough to swamp the underlying trends. At a minimum, before presenting comparative results, the analyst should examine the underlying questionnaires for comparability and should be cautious about the way in which any comparative results are presented and interpreted.

Review Questions

4. Living Standards Measurement Surveys are generally characterized by all of the following *except*:

- A. Large sample sizes.
- B. Multiple questionnaires.
- C. Close attention to quality control.
- D. Extensive questions.

Measuring Poverty: Choosing an Indicator of Welfare

There are a number of ways to measure well-being. The welfarist approach (Sen 1979) seeks to measure household *utility*, which in turn is usually assumed to be approximated by household consumption expenditure or household income; these may be considered as *inputs* into generating utility. Given enough income, the household is assumed to know best how to deploy these resources, whether on food, clothing, housing, or the like. When divided by the number of household members, this gives a per capita measure of consumption expenditure or income. Of course, even household expenditure or income is an imperfect proxy for utility; for instance, it excludes potentially important contributors to utility such as publicly provided goods or leisure.

A more paternalistic, or nonwelfarist, approach might focus on whether households have attained certain minimal levels of, say, nutrition or health. Thus, while the welfarist approach focuses on per capita consumption expenditure or income, other (nonwelfarist) measures of individual welfare might include indicators such as infant mortality rates in the region, life expectancy, the proportion of spending devoted to food, housing conditions, or child schooling; these may be thought of as measures of *output*, reflections of utility rather than inputs into the generation of utility.

Such measures are useful in fleshing out a multidimensional portrait of poverty, but they rest on a somewhat different philosophical foundation from the welfarist approach, and this can make interpretation difficult. For instance, if people have enough income to feed, clothe, and house themselves adequately, how concerned should we then be if they do not in fact do so? In some cases there may be informational problems—perhaps no one in the household knows how to cook—but in the absence of such imperfections, to what extent are we justified in trying to save people from themselves? This age-old dilemma does not have a simple solution.

If we choose to assess poverty based on household consumption or expenditure per capita, it is helpful to think in terms of an expenditure function, which shows the minimum expense required to meet a given level of utility u , which is derived from a vector of goods x , at prices p . It can be obtained from an optimization problem in which the objective function (expenditure) is minimized subject to a set level of utility, in a framework where prices are fixed.

Let the consumption measure for the household i be denoted by y_i . Then an expenditure measure of welfare may be denoted by:

$$y_i = p \cdot q = e(p, x, u), \quad (2.1)$$

where p is a vector of prices of goods and services, q is a vector of quantities of goods and services consumed, $e(\cdot)$ is an expenditure function, x is a vector of household characteristics (number of adults, number of young children, and so on), and u is the level of “utility” or well-being achieved by the household. Put another way, given the prices (p) that it faces, and its demographic characteristics (x), y_i measures the spending that is needed to reach utility level u .

Typically, we compute the actual level of y_i from household survey data that include information on consumption. Once we have computed y_i , we can construct *per capita* household consumption for every individual in the household, which implicitly assumes that consumption is shared equally among household members. For this approach to make sense, we must also assume that all individuals in the household have the same needs. This is a strong assumption, for in reality, different individuals have different needs based on their individual characteristics.

Several factors complicate the process of estimating per capita consumption. Table 2.2 reports estimates of both nominal and inflation-adjusted (“real”) per capita consumption from three different household surveys in Cambodia. Using the 1997 Cambodia Socioeconomic Survey (CSES), for example, nominal and real per

Table 2.2 Summary of per Capita Consumption from Cambodian Surveys

Surveys	Nominal	Real (inflation adjusted)
SESC 1993–94	1,833	2,262
CSES 1997 (adjusted)	2,223	2,530
CSES 1997 (unadjusted)	1,887	2,153
CSES 1999 (Round 1)	2,037	1,630
CSES 1999 (Round 2)	2,432	1,964
CSES 1999 (both rounds)	2,238	1,799

Source: Gibson 1999.

Note: CSES = Cambodia Socio-Economic Survey. SESC = Socio-Economic Survey of Cambodia. All values are in riels per person per day. Real values are estimated in 1993–94 Phnom Penh prices, as deflated by the value of the food poverty lines. Adjusted figures from 1997 incorporate corrections for possible underestimation of certain types of consumption (see Knowles [1998] and Gibson [1999] for details). Differences between Rounds 1 and 2 in 1999 are detailed in Gibson (1999).

capita consumption were 2,223 and 2,530 riels, respectively (Gibson 1999). However, across years the estimates of consumption in real terms for 1993–94 may not be directly comparable with the 1999 estimates because the surveys did not have exactly the same set of questions regarding consumption. For example, real consumption per capita was measured as 2,262 riels for 1993–94, but was only 1,799 in 1999, despite robust economic growth during the interval; this may merely be an artifact of the different ways in which questions were asked.

Traditionally, we use a monetary measure to value household welfare. The two most obvious candidates are income and expenditure.

Candidate 1: Income

It is tempting to measure household welfare by looking at household income. Practical problems arise immediately: What is income? Can it be measured accurately? The most generally accepted measure of income is the one formulated by Haig and Simons (Haig 1921, Simons 1938):

Income \equiv consumption + change in net worth.

Example: Suppose I had assets of \$10,000 at the beginning of the year. During the year I spent \$3,000 on consumption. And at the end of the year I had \$11,000 in assets. Then my income was \$4,000, of which \$3,000 was spent, and the remaining \$1,000 added to my assets.

The first problem with this definition is that it is not clear what time period is appropriate. Should we look at someone's income over a year? Five years? A lifetime? Many students are poor now, but have good lifetime prospects, and we may not want to consider them as being truly poor. On the other hand, if we wait until we have information about someone's lifetime income, it will be too late to help him or her in moments of poverty.

The second problem is measurement. It is easy enough to measure components of income such as wages and salaries. It may be possible to get adequate (if understated) information on interest, dividends, and income from some types of self-employment. But it is likely to be hard to get an accurate measure of farm income; or of the value of housing services; or of capital gains (for example, the increase in the value of animals on a farm, or the change in the value of a house that one owns).

For instance, the Vietnam Living Standards Surveys undertaken in 1992–93 (Vietnam 1994) and again in 1997–98 (Vietnam 2000) collected information on the value of farm animals at the time of the survey, but not the value a year before. Thus, it was not possible to measure the change in the value of animal assets. Many farmers that reported negative cash income may in fact have been building up assets, and they actually had positive income.

It is typically the case, particularly in societies with large agricultural or self-employed populations, that income is seriously understated. This certainly appears to be the case for Vietnam. Table 2.3 shows income per capita for households in 1993 for each of five expenditure quintiles: a quintile is a fifth of the sample, and quintile 1 contains the poorest fifth of individuals, and so on. For every quintile, households on average reported less income than expenditure, which is simply not plausible. This would imply that households must be running down their assets, or taking on much more debt, which was unlikely in a boom year like 1993.

Income tends to be understated for several reasons:

- People forget, particularly when asked in a single interview, about items they may have sold, or money they may have received, up to a year before.
- People may be reluctant to disclose the full extent of their income, lest the tax collector or a neighbor get wind of the details.
- People may be reluctant to report income earned illegally, for instance, from smuggling, corruption, poppy cultivation, or prostitution.
- Some parts of income are difficult to calculate, for example, the extent to which the family buffalo has risen in value.

Research based on the 1969–70 socioeconomic survey in Sri Lanka estimated that wages were understated by 28 percent; business income by 39 percent; and rent, interest, and dividends by 78 percent (Visaria 1980, 18). It is not clear how much these figures are applicable elsewhere, but they do give a sense of the potential magnitude of the understatement problem.

Candidate 2: Consumption Expenditure

Consumption includes both goods and services that are purchased and those that are provided from one's own production (in-kind).

Table 2.3 Income and Expenditure by per Capita Expenditure Quintiles, Vietnam

(In thousands of dong per capita per year, 1992/93)

	Lowest	Lower-mid	Middle	Mid-upper	Highest	Overall
Income/capita	494	694	956	1,191	2,190	1,105
Expenditure/capita	518	756	984	1,338	2,540	1,227
Memo						
Food spending/capita	378	526	643	807	1,382	747
Food as % of expenditure	73	70	65	60	54	61

Source: Vietnam 1994.

Note: In 1993, exchange rate was about 10,000 dong/US\$.

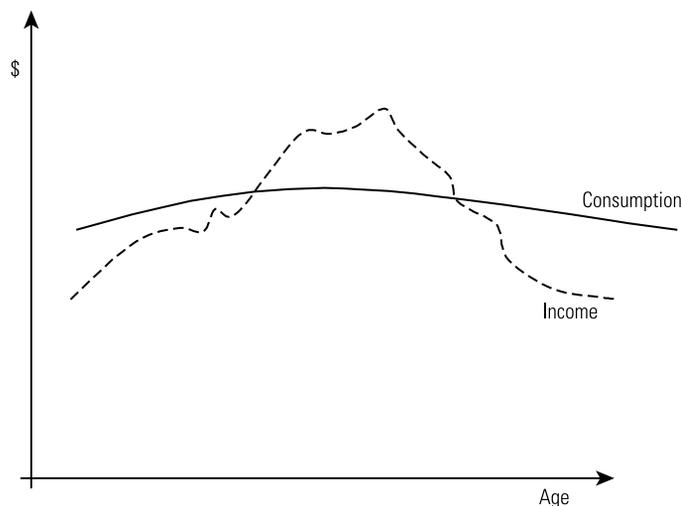
In developed countries, a strong case can be made that consumption is a better indicator of lifetime welfare than is income. Income typically rises and then falls in the course of one's lifetime, in addition to fluctuating somewhat from year to year, whereas consumption remains relatively stable. This smoothing of short-term fluctuations in income is predicted by the permanent income hypothesis, under which transitory income is saved while long-term (permanent) income is largely consumed.

The lifecycle of income and consumption is captured graphically in figure 2.2. While the available evidence does not provide strong support for this *lifecycle hypothesis* in the context of less-developed countries, households there do appear to smooth out the very substantial seasonal fluctuations in income that they typically face during the year (see Alderman and Paxson 1994; Paxson 1993). Thus, information on consumption over a relatively short period, such as one a month, as typically collected by a household survey is more likely to be representative of a household's general level of welfare than equivalent information on income, which is more volatile.

A more practical case for using consumption, rather than income, is that households may be more able, or willing, to recall what they have spent rather than what they earned. Even so, consumption is likely to be systematically understated for the following reasons:

- Households tend to underdeclare what they spend on luxuries or illicit items. For instance, the amount that households said they spent on alcohol, according to the 1972–73 household budget survey in the United States was just half the amount that companies said they sold (Carlson 1974).

Figure 2.2 Lifecycle Hypothesis: Income and Consumption Profile over Time



Source: Authors.

- Questions matter. According to VLSS93, Vietnamese households devoted 1.7 percent of their expenditure to tobacco; the VLSS98 figures showed that this had risen to 3 percent. An increase of this magnitude is simply implausible and not in line with sales reported by the cigarette and tobacco companies. A more plausible explanation is that VLSS98 had more detailed questions about tobacco use. When the questions are more detailed, respondents are likely to remember in more detail and to report higher spending.

The understatement of both income and consumption means that poverty rates are overstated. It also means that the estimates of total income and consumption that are based on the survey data invariably fall short of the levels observed in national accounts data—and in some countries this gap is growing. It is tempting to gross up the survey results—for instance, raising everybody’s income by 10 percent if this is the size of the gap between the survey data and national accounts—before computing poverty rates. Some countries do this, but it is not a satisfactory solution, since understatement seems to be a smaller problem for the poor than the rich, at least in absolute terms. We address this issue in more detail in chapter 10.

Measuring Durable Goods. In measuring poverty it might be argued that only food, the ultimate basic need (which constitutes seven-tenths of the spending of poor households), should be included. On the other hand, even households that cannot afford adequate quantities of food devote some expenditures to other items, such as clothing, and shelter. It is reasonable to suppose that if these items are getting priority over food purchases, then they must represent very basic needs of the household, and so should be included in the poverty line. This argument also applies to durable goods.

The problem is that durable goods, such as bicycles and televisions, are bought at a point in time, and then consumed over a period of several years. Consumption should only include the amount of a durable good that is eaten up during the year, which can be measured by the change in the value of the asset during the year, plus the cost of locking up one’s money in the asset.

Example: For instance, if my watch was worth \$25 a year ago, and is worth \$19 now, then I used \$6 worth of watch during the year; I also tied up \$25 worth of assets in the watch, money that could have earned me \$2.50 in interest (assuming 10 percent) during the year. Thus, the true cost of the watch during the year was \$8.50. This is essentially the amount that I would have to pay if I were to try to rent the watch for a year.

A comparable calculation needs to be done for each durable good that the household owns. Clearly the margin of potential measurement error is large, since the price of each asset may not be known with much accuracy, and the interest rate used

is somewhat arbitrary. The Vietnamese VLSS surveys asked for information about when each good was acquired, and at what price, and the estimated current value of the good.¹ This suffices to compute the current consumption of the durable item, as the illustration in box 2.1 shows.

One might wonder why attention needs to be paid to calculating the value of durable goods consumption when the focus is on poverty—in practice, first and foremost the ability to acquire enough food. The answer is that when expenditure is used as a yardstick of welfare, it is important to achieve comparability across households. If the value of durable goods were not included, one might have the impression that a household that spends \$100 on food and \$5 on renting a bicycle is better off than a household that spends \$100 on food and owns a bicycle (that it could rent out for \$5), when in fact both households are equally well off, all else being equal.

Box 2.1 Calculating the Value of Durable Goods Consumption: An Illustration

A Vietnamese household surveyed in April 1998 says that it bought a television two years earlier for 1.1 m dong (about US\$100). The television is now believed to be worth 1 m dong. Over all prices rose by a total of 10 percent over the past two years. How much of the television was consumed over the year prior to the survey?

- Recompute the values in today's prices. Thus, the TV, purchased for 1.1 m dong in 1996, would have cost 1.21 m dong ($= 1.1 \text{ m dong} \times (1+10 \text{ percent})$) now.
- Compute the depreciation. The television lost 0.21 m dong in value in two years, or 0.105 m dong per year (about US\$7).
- Compute the interest cost. At today's prices, the television was worth 1.105 m dong a year ago (1.21 m dong less this past year's depreciation of 0.105 m dong), and this represents the value of funds locked up during the year prior to the survey. At a real (inflation-adjusted) interest rate of 3 percent, the cost of locking up these resources was 0.03315 m dong over the course of the year.

Thus the total consumption cost of the television was 0.138 m dong ($= 0.105 + 0.033$), or about US\$10.

This computation is only possible if the survey collects information on the past prices of all the durables used by the household. Where historical price data are not available, researchers in practice typically apply a depreciation + interest rate to the reported value of the goods; so if a television is worth 1 m dong now, is expected to depreciate by 10 percent per annum, and the real interest rate is 3 percent, then the imputed consumption of the durable good is measured as 1 m ($(10 \text{ percent} + 3 \text{ percent}) = 0.13 \text{ m dong}$). Deaton and Zaidi (1998) recommend that one use average depreciation rates derived from the sample, rather than the rates reported by each individual household.

Measure the Value of Housing Services. If you own your house or apartment, it provides housing services that should be considered as part of consumption. The most satisfactory way to measure the values of these services is to ask how much you would have to pay if, instead of owning your home, you had to rent it, although this question is seldom asked in practice.

The standard procedure is to estimate, for those households that rent their dwellings, a function that relates the rental payment to such housing characteristics as the size of the house (in square feet of floor space), the year in which it was built, the type of roof, and whether there is running water. This gives the following:

$$\text{Rent} = f(\text{area, running water, year built, type of roof, location, number of bathrooms...})$$

The estimates based on this “hedonic” regression then are used to impute the value of rent for those households that own, rather than rent, their housing. For all households that own their housing, this imputed rental, along with the costs of maintenance and minor repairs, represents the annual consumption of housing services.² In the case of households that pay interest on a mortgage, it is appropriate to count the imputed rental and costs of maintenance and minor repairs in measuring consumption, but not the mortgage interest payments as well, because this would represent double-counting.³

In the case of Vietnam there is a problem with this approach: almost nobody rents housing. Of those that do, most pay a nominal rent for a government apartment. Only 13 of the 5,999 households surveyed in VLSS98 paid private sector rental rates.⁴ On the other hand the VLSS surveys did ask each household to put a capital value on their house or apartment. In computing consumption expenditure, the Vietnam General Statistics Office assumed that the rental value of housing was 3 percent of the capital value of the housing. This is, of course, a somewhat arbitrary procedure.

Weddings and Funerals. Families spend money on weddings. Such spending is often excluded when measuring household consumption expenditure. The logic is that the money spent on weddings mainly gives utility to the guests, not the spender. Of course if one were to be strictly correct, then expenditure should include the value of the food and drink that one enjoys as a guest at other people’s weddings, although in practice this is rarely included. Alternatively, one might think of wedding expenditures as rare and exceptional events, which shed little light on the living standard of the household. Similar considerations apply to other large and irregular spending, on items such as funerals and dowries.

Accounting for Household Composition Differences. Households differ in size and composition, and so a simple comparison of aggregate household consumption can be quite misleading about the well-being of individuals in a given household. Most researchers recognize this problem and use some form of normalization. The

most straightforward method is to convert from household consumption to individual consumption by dividing household expenditures by the number of people in the household. Then, total household expenditure per capita is the measure of welfare assigned to each member of the household. Although this is by far the most common procedure, it is not very satisfactory for two reasons:

- First, different individuals have different needs. A young child typically needs less food than an adult, and a manual laborer requires more food than an office worker.
- Second, there are economies of scale in consumption, at least for such items as housing. It costs less to house a couple than to house two individuals separately.

Example: Suppose we have a household with two members and monthly expenditure of \$150 in total. We would then assign each individual \$75 as their monthly per capita expenditure. If we have another household with three members, it would appear that each member is worse off, with only \$50 per capita per month. However, suppose we know that the two-person household contains two adult males age 35, whereas the second household contains one adult female and two young children. This added information may change our interpretation of the level of well-being in the second household, since we suppose that young children may have much lower costs (at least for food) than adults.

In principle, the solution to this problem is to apply a system of weights. For a household of any given size and demographic composition (such as one male adult, one female adult, and two children), an equivalence scale typically measures the number of adult males to which that household is deemed to be equivalent. Each member of the household counts as some fraction of an adult male. Effectively, household size is the sum of these fractions and is not measured in numbers of persons but in numbers of *adult equivalents*. Economies of scale can be allowed for by transforming the number of adult equivalents into effective adult equivalents.

In the abstract, the notion of equivalence scale is compelling. It is much less persuasive in practice, because of the problem of picking an appropriate scale. How these weights should be calculated and whether it makes sense even to try is still subject to debate, and there is no consensus on the matter. However, equivalence scales are not necessarily unimportant. For example, take the observation that in most household surveys, per capita consumption decreases with household size. It is probably more appropriate to interpret this as evidence that there are economies of scale to expenditure, and not necessarily as proof that large households have a lower standard of living.

The commonest solution to this problem is to pick a scale that seems reasonable, on the grounds that even a bad equivalence scale is better than none at all, and explore the robustness of the results (for example, estimates of the poverty rate) to different equivalence scales. Often the equivalence scales are based on the different calorie needs of individuals of different ages.

OECD scale The OECD scale is widely used, and may be written as

$$AE = 1 + 0.7 (N_{adults} - 1) + 0.5 N_{children} \quad (2.2)$$

where *AE* refers to “adult equivalent.” A one-adult household would have an adult equivalent of 1, a two-adult household would have an *AE* of 1.7, and a three-adult household would have an *AE* of 2.4. Thus the 0.7 reflects economies of scale; the smaller this parameter, the more important economies of scale are considered to be. In developing countries, where food constitutes a larger part of the budget, economies of scale are likely to be less pronounced than in rich countries. The 0.5 is the weight given to children, and presumably reflects the lower needs (for food, housing space, and so forth) of children. Despite the elegance of the formulation, there are real problems in obtaining satisfactory measures of the degree of economies of scale and even of the weight to attach to children.

Other scales Many other scales have been used. For instance, a number of researchers used the following scale in analyzing the results of the LSMS that were undertaken in the Côte d’Ivoire, Ghana, and Peru (Glewwe and Twum-Baah 1991):

Age (years)	0–6	7–12	13–17	>17
Weight (i.e., adult equivalences)	0.2	0.3	0.5	1.0

An elegant formulation is as follows:

$$AE = (N_{adults} + \alpha N_{children})^{\theta}$$

where α measures the cost of a child relative to an adult, and $\theta \leq 1$ is a parameter that captures the effects of economies of scale. Consider a family with two parents and two children. For $\alpha = \theta = 1$, $AE = 4$ and our welfare measure becomes expenditure per capita. But if $\alpha = 0.7$ and $\theta = 0.8$, then $AE = 2.67$, and the measure of expenditure per adult equivalent will be considerably larger.

Estimate an equivalence scale Several researchers have tried to estimate the extent to which there are economies of scale in consumption, essentially by looking at how aggregate household consumption of goods, such as food, varies with household size and composition (for example, Pendakur 1999; see, too, Jenkins and

Cowell 1994). There are a number of problems here. First, it assumes that resources like food are equitably distributed within the household, although in practice the intrahousehold allocation is likely to reflect the distribution of power among household members. And second, there is a very basic identification problem: If children provide utility to parents, then how can we say that a couple with a child, earning \$10,000, is necessarily worse off than a childless couple with the same income? (Pollack and Wales 1979). The consensus view is summed up by Deaton and Zaidi (1998), who argue, “there are so far no satisfactory methods for estimating economies of scale.”

Income or Expenditure? Most rich countries measure poverty using income, while most poor countries use expenditure. There is a logic to this; in rich countries, income is comparatively easy to measure (much of it comes from wages and salaries), while expenditure is complex and hard to quantify. On the other hand, in less-developed countries income is hard to measure (much of it comes from self employment), while expenditure is more straightforward and hence easier to estimate. The arguments for and against income and consumption as the appropriate welfare measures for poverty analysis are summarized in table 2.4; further discussion may be found in Hentschel and Lanjouw (1996), Blundell and Preston (1998), and Donaldson (1992).

Table 2.4 Which Indicator of Welfare: Income or Consumption?

Income (“potential”)

Pro:

- Easy to measure, given the limited number of sources of income.
- Measures degree of household “command” over resources (which they could use if they so wish).
- Costs only a fifth as much to collect as expenditure data, so sample can be larger.

Con:

- Likely to be underreported.
- May be affected by short-term fluctuations (for example, the seasonal pattern of agriculture).
- Some parts of income are hard to observe (for example, informal sector income, home agricultural production, self-employment income).
- Link between income and welfare is not always clear.
- Reporting period might not capture the “average” income of the household.

Consumption (“achievement”)

Pro:

- Shows current actual material standard of living.
- Smooths out irregularities, and so reflects long-term average well-being.
- Less understated than income, because expenditure is easier to recall.

Con:

- Households may not be able to smooth consumption (for example, via borrowing, social networks).
 - Consumption choices made by households may be misleading (for example, if a rich household chooses to live simply, that does not mean it is poor).
 - Some expenses are not incurred regularly, so data may be noisy.
 - Difficult to measure some components of consumption, including durable goods.
-

Review Questions

5. One commonly used measure of household welfare is income, which ideally is defined as

- A. Wages plus salaries.
- B. Earnings plus remittances.
- C. Wages plus profits plus transfer income.
- D. Consumption plus change in net worth.

6. Especially in less-developed countries, income as reported in household surveys is typically understated, even by poor people, because

- A. People forget how much income they made over the past year or even month.
- B. Income taxes are high.
- C. Most households rent their homes in the informal market.
- D. Illegal income is a large proportion (typically at least a quarter) of household income.

7. Expenditure is widely believed to reflect welfare better than incomes, in less-developed countries, because

- A. Annual expenditure is more closely related to lifetime (“permanent”) income than is annual income.
- B. Households understate their spending on luxuries.
- C. Household surveys tend to ask more questions about expenditure than about income.
- D. The imputed value of durables (including housing) is included in expenditure but not income.

8. A household owns a bicycle that it bought two years ago for \$40 and is now worth \$28. It faces an interest rate of 10%. Then the true economic value of the services of the bicycle over the past year was closest to

- A. \$12.
- B. \$9.40.
- C. \$6.60.
- D. \$4.00.

9. For households that rent their homes, expenditure on housing is given by the rent they pay; and for households that own their homes, there is no need to count expenditure on housing because it is already counted as part of income.

- True.
- False.

10. The OECD scale measures adult equivalents as follows:

$$AE = 1 + 0.7 (N_{\text{adults}} - 1) + 0.5 N_{\text{children}}$$

Based on this scale, a household consisting of two parents, a grandmother, two teenagers, and an infant would:

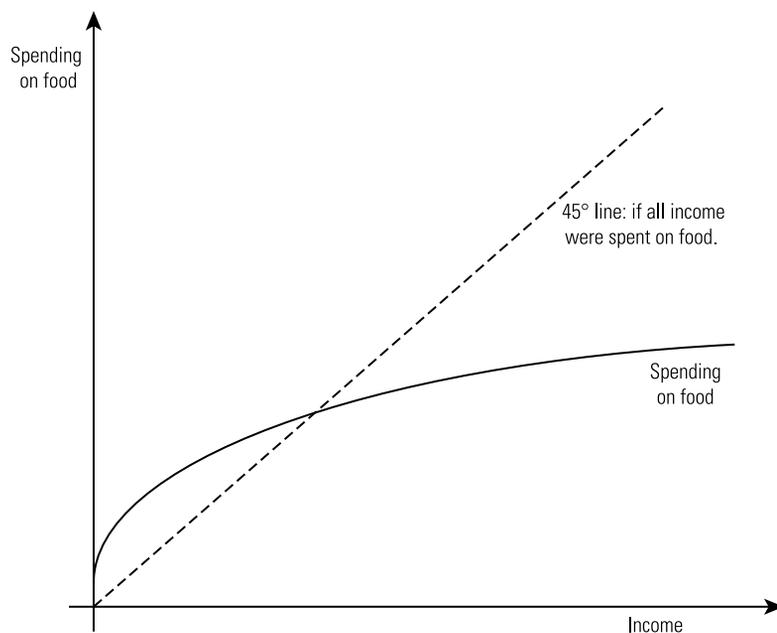
- A. Have a lower adult equivalent expenditure than it would have expenditure per capita.
- B. Have an adult equivalent of 3.4.
- C. Have a higher adult equivalent than a household with five adults.
- D. Become better off when the eldest teenager becomes an adult.

Candidate 3. Other Measures of Household Welfare

Even if they were measured perfectly, neither income nor expenditure would be an ideal measure of household well-being. For instance, neither measure puts a value on the leisure time enjoyed by the household; neither measures the value of publicly provided goods (such as education, or public health services); and neither values intangibles such as peace and security.

Other possible measures of well-being include the following:

- *Calories consumed per person per day.* If one accepts the (nonwelfarist) notion that adequate nutrition is a prerequisite for a decent level of well-being, then we could just look at the quantity of calories consumed per person. Anyone consuming less than a reasonable minimum—often set at 2,100 Calories per person per day—would be considered poor.⁵ Superficially, this is an attractive idea, and we will return to it in chapter 3. However, at this point we just note that it is not always easy to measure calorie intake, particularly if one wants to distinguish between different members of a given household. Nor is it easy to establish the appropriate minimum number of calories per person, as this will depend on the age, gender, and working activities of the individual.
- *Food consumption as a fraction of total expenditure.* Over a century ago, Ernst Engel observed in Germany that as household income per capita rises, spending on food rises too, **but less quickly**. This relationship is shown in figure 2.3. As a result, the proportion of expenditure devoted to food falls as per capita income rises. One could use this finding, which is quite robust, to come up with a measure of well-being and hence a measure of poverty. For instance, households that devote more than (say) 60 percent of their expenditures to food might be considered to be poor. The main problem with this measure is that the share of spending going to food also depends on the proportion of young to old family members (more children indicates a higher proportion of spending on food), and on the relative price of food (if food is relatively expensive, the proportion of spending going to food will tend to be higher).

Figure 2.3 Engel Curve: Food Spending Rises Less Quickly Than Income

- *Measures of outcomes rather than inputs.* Food is an input, but nutritional status (being underweight, stunting or wasting) is an output. So one could measure poverty by looking at malnutrition.

Of course, this requires establishing a baseline anthropometric standard against which to judge whether someone is malnourished. This is a controversial issue; generally, less stunting (as measured by height for age) is found in Sub-Saharan Africa than in Southeast Asia, although there is no reason to believe that the latter region is poorer than Africa. On the other hand, anthropometric indicators have the advantage that they can reveal living conditions within the household (rather than assigning the overall household consumption measure across all members of the household without really knowing how consumption expenditure is divided among household members).

- *Peer or observer assessments.* In Vietnam, very poor households are eligible for some subsidies, to cover health care and educational fees, for instance. The decision about who qualifies as being sufficiently poor is taken at the village level, where the local People's Committee typically knows enough about individual households to make the determination. Krishna et al. (2004) have made use of villager assessments in rural Kenya to validate observations about the degree of household poverty and its evolution over a generation. And Lanjouw and Stern (1991) classified villagers in a village in northern India into categories of poverty and wealth, based on discussions with villagers themselves.

The main problem with such assessments is that, because they are based on perceptions formed on imperfect information, they may be biased. Lanjouw and Stern found that in their survey village, landless agricultural laborers were almost all deemed by their peers to be poor; yet based on income information, only about half fell below the poverty line used.

When one is looking at a community (province, region) rather than individual households, it might make sense to judge the poverty of the community by life expectancy, or the infant mortality rate, although these are not always measured very accurately. School enrollments (a measure of investing in the future generation) represent another outcome that might indicate the relative well-being of the population. Certainly, none of these other measures of well-being are replacements for consumption per capita; nor does consumption per capita fully replace these measures. Rather, when taken together they allow us to get a more complete and multidimensional view of the well-being of a population, although this does not guarantee greater clarity. Consider the statistics in table 2.5, which refer to 11 different countries. How countries are ranked in terms of living standards clearly depends on which measure or indicator is considered.

In sum, there is no ideal measure of well-being: all measures of poverty are imperfect. That is not an argument for avoiding measuring poverty, but rather for approaching all measures of poverty with a degree of caution, and for asking in some detail about how the measures are constructed.

Table 2.5 Poverty and Quality of Life Indicators

Countries	GNP per capita (1999 dollars)	% population below poverty line	Female life expectancy at birth, years (1998)	Prevalence of child malnutrition, % children <5 years (1992–98)	Female adult illiteracy rate, % of people 15+ years, (1998)
Algeria	1,550	22.6 (1995)	72	13	46
Bangladesh	370	35.6 (1995/96)	59	56	71
Cambodia	260	36.1 (1997)	55	—	80
Colombia	2,250	17.7 (1992)	73	8	9
Indonesia	580	20.3 (1998)	67	34	20
Jordan	1,500	11.7 (1997)	73	5	17
Morocco	1,200	19.0 (1998/99)	69	10	66
Nigeria	310	34.1 (1992/93)	55	39	48
Peru	2,390	49.0 (1997)	71	8	16
Sri Lanka	820	35.3 (1990/91)	76	38	12
Tunisia	2,100	14.1 (1990)	74	9	42

Source: World Bank 2000.

Note: — = Not available.

Review Question

11. Which of the following is *not* a popular measure of welfare at the household level:
- A. The percentage of its spending that a household devotes to food.
 - B. The average height of children given their age, relative to a well-fed population.
 - C. The consumption of calories per capita.
 - D. Life expectancy at birth.

Notes

1. The questionnaires for the VLSS surveys are available from the Environmental and Social Statistics Department of the General Statistics Office, Hanoi.
2. This assumes that renters are responsible for maintenance and repair costs, so that the rental paid does not include a provision for these items. In some countries the owner, rather than the renter, would bear these costs, in which case the imputed rental also includes the costs, and no further adjustment would be called for.
3. However, if we want to measure income (rather than consumption), then we should use the imputed rental for households that own their property free and clear, and rental less mortgage interest payments for those who have borrowed against their housing.
4. Computation by the authors using the data from the Vietnam Living Standards Survey of 1997/98. The data are available from the General Statistics Office, Hanoi.
5. A calorie is the energy required to heat 1 gram of water by 1°C. A kilocalorie, also referred to as a Calorie, represents 1,000 calories.

References

- Albert, José Ramon. 2004. "Measuring Poverty." Presentation (PowerPoint) at course on Introduction to Poverty Measurement and Analysis, Ministry of Economic Planning and Development, Blantyre, Malawi, January 19–20.
- Alderman, Harold, and Chris Paxson. 1994. "Do the Poor Insure? A Synthesis of the Literature on Risk and Consumption in Developing Countries." In *Economics in a Changing World, Vol. 4: Development, Trade, and Environment*, ed. E. L. Bacha. London: Macmillan.
- Blundell, R., and I. Preston. 1998. "Consumption Inequality and Income Uncertainty." *Quarterly Journal of Economics* 113: 603–40.
- Carlson, Michael D. 1974. "The 1972–73 Consumer Expenditure Survey." *Monthly Labor Review* December, 16–23.
- Deaton, Angus. 1997. *The Analysis of Household Surveys: A Macroeconomic Approach to Development Policy*. Baltimore: Johns Hopkins University Press.
An essential volume for any graduate student working with survey data.
- Deaton, Angus, and Salman Zaidi. 1999. "Guidelines for Constructing Consumption Aggregates for Welfare Analysis." Woodrow Wilson School of Public and International Affairs, Princeton University, Princeton, NJ. http://www.wws.princeton.edu/%7Erpds/downloads/deaton_zaidi_consumption.pdf Subsequently issued in 2002 as Living Standards Measurement Study Working Paper No. 135, World Bank, Washington, DC.

- A clear, sensible discussion of the practical issues that arise in measuring a consumption indicator of welfare. Includes a sample questionnaire and some useful Stata code.
- Donaldson, D. 1992. "On the Aggregation of Money Measures of Well-Being in Applied Welfare Economics." *Journal of Agricultural and Resource Economics* 17: 88–102.
- Duclos, Jean-Yves, and Abdelkrim Araar. 2006. *Poverty and Equity: Measurement, Policy and Estimation with DAD*. New York: Springer, and Ottawa: International Development Research Centre.
- An excellent, but technically rather advanced, treatment of poverty and inequality, with a very complete discussion of measurement and estimation. The authors have developed their own (free) software, DAD, that is one of the most accurate and comprehensive tools for reporting and graphing measures of poverty and inequality.
- Gibson, John. 1999. "A Poverty Profile of Cambodia, 1999." Report to the World Bank and the Ministry of Planning, Phnom Penh.
- Glewwe, Paul, and Kwaku A. Twum-Baah. 1991. "The Distribution of Welfare in Ghana, 1987–88." Living Standards Measurement Study Paper No. 75, World Bank, Washington, DC.
- Grosh, Margaret, and Paul Glewwe, eds. 2000. *Designing Household Survey Questionnaires for Developing Countries: Lessons from 15 Years of Living Standard Measurement Study*. Volumes 1, 2, and 3. Washington, DC: World Bank and Oxford University Press.
- Every statistics office should have a copy of this pair of volumes, or better still the CD-ROM version. This reference work includes sample questionnaires as well as detailed chapters on all aspects of designing, implementing, and using living standard measurement surveys.
- Haig, Robert M. 1921. "The Concept of Income: Economic and Legal Aspects." In *The Federal Income Tax*, 1–28. New York: Columbia University Press.
- Hentschel, J., and P. Lanjouw. 1996. "Constructing an Indicator of Consumption for the Analysis of Poverty: Principles and Illustrations with Reference to Ecuador." Working Paper No. 124, Living Standards Measurement Study, World Bank, Washington, DC.
- Jenkins, S., and F. Cowell. 1994. "Parametric Equivalence Scales and Scale Relativities." *Economic Journal* 104: 891–900.
- Knowles, James C. 1998. "An Updated Poverty Profile for Cambodia–1997: Technical Report." Ministry of Planning, Phnom Penh.
- Krishna, Anirudh, Patricia Kristjanson, Maren Radeny, and Wilson Nindo. 2004. "Escaping Poverty and Becoming Poor in Twenty Kenyan Villages." *Journal of Human Development*, 5 (2): 211–26.
- Lanjouw, Peter, and Nicholas Stern. 1991. "Poverty in Palanpur." *World Bank Economic Review* 5 (1): 23–55.
- Paxson, Christina. 1993. "Consumption and Income Seasonality in Thailand." *Journal of Political Economy* 101 (1): 39–72.
- Pendakur, K. 1999. "Semiparametric Estimates and Tests of Base-Independent Equivalence Scales." *Journal of Econometrics* 88: 1–40.
- Pollack, R., and T. Wales. 1979. "Welfare Comparisons and Equivalence Scales." *American Economic Review* 69: 216–21.
- Ravallion, Martin. 1998. "Poverty Lines in Theory and Practice." LSMS Working Paper No. 133, World Bank, Washington, DC.
- . 1992. "Poverty Comparisons: A Guide to Concepts and Methods." LSMS Working Paper No. 88, World Bank, Washington, DC.

- A very useful survey, with an emphasis on comparing poverty across time and space. In writing chapters 2 and 3, we have drawn heavily on the ideas in this paper.
- Rowntree, S. 1901. *Poverty: A Study of Town Life*. London: Macmillan.
- Sen, Amartya. 1979. "Personal Utilities and Public Judgements: Or What's Wrong with Welfare Economics?" *The Economic Journal* 89: 537–58.
- Simons, Henry. 1938. *Personal Income Taxation: The Definition of Income as a Problem of Fiscal Policy*. Chicago: University of Chicago Press.
- Vietnam, General Statistical Office. 2000. *Vietnam Living Standards Survey 1997–1998*. Hanoi: Statistical Publishing House.
- Vietnam, State Planning Committee and General Statistical Office. 1994. *Vietnam Living Standards Survey 1992–1993*. Hanoi.
- Visaria, Pravin (with Shyamalendu Pal). 1980. "Poverty and Living Standards in Asia." Living Standards Measurement Study Working Paper No. 2, World Bank, Washington, DC.
- World Bank. 2000. *World Development Report 2000/2001: Attacking Poverty*. Washington, DC: World Bank.

